Synthese (2007) 158:41–60
DOI 10.1007/s11229-006-9049-5

ORIGINAL PAPER

The undecidability of propositional adaptive logic

Leon Horsten · **Philip Welch**

Received: 27 October 2005 / Accepted: 8 May 2006 / Published online: 16 September 2006

Abstract We investigate and classify the notion of final derivability of two basic inconsistency-adaptive logics. Specifically, the maximal complexity of the set of final consequences of decidable sets of premises formulated in the language of propositional logic is described. Our results show that taking the consequences of a decidable propositional theory is a complicated operation. The set of final consequences according to either the Reliability Calculus or the Minimal Abnormality Calculus of a decidable propositional premise set is in general undecidable, and can be Σ^0_3-complete. These classifications are exact. For first order theories even finite sets of premises can generate such consequence sets in either calculus.

Keywords Adaptive logic · Paraconsistent logic · Dynamic logic · Undecidability

1 Introduction

Adaptive logics have been proposed as systems for reasoning sensibly from inconsistent premise sets. When an inconsistent set of premises is given, the rules of adaptive logic allow one to derive sound information concerning the class of those models that are no more inconsistent than is required by the premises.

The distinguishing feature of adaptive logic is that it involves a *revision rule*. In general, consequences that are drawn from a premise set are provisional: it occasionally happens that the reasoner is forced in the course of a reasoning process to withdraw earlier conclusions. By thus inferring and occasionally revising according to the rules of adaptive logic, the reasoner gradually zooms in on the structure of

L. Horsten (✉)
University of Leuven, Leuven,
Brabant, B-3000, Belgium
e-mail: Leon.Horsten@hiw.kuleuven.be

P. Welch
Bristol University, Bristol, UK
e-mail: p.welch@bristol.ac.uk

the minimally inconsistent models, i.e., the models that verify no more contradictions than are necessary to make a given inconsistent set of premises true.

We shall investigate complexity aspects of adaptive consequence relations. As our point of departure, we consider two very basic inconsistency-adaptive logics, developed by Diderik Batens. In order to keep these questions manageable, we will work in a simplified setting. For this reason, we concentrate on the *propositional* fragment of Batens' systems of adaptive logic. We will be interested in Batens' two main systems of adaptive logic, called **ACLuN1** and **ACLuN2** by him, or also the *reliability reasoning strategy* and the *minimal abnormality reasoning strategy*. Both strategies will be investigated in this paper. We will be mostly concerned with propositional adaptive logics, but we shall also make some remarks about adaptive predicate logics.

Even though in general a consequence that is adaptively drawn from a set of premises is revisable by further extending the argument, there are also situations in which a reasoner has adaptively arrived at a consequence for which no reasons can be adduced for retracting them in a later stage. Such a consequence is said to be an unrevisable or *final* consequence of the set of premises. The focus of this paper is on the complexity of the collection of final consequences of a recursive (but possibly infinite) set of premises. Propositional derivability relations are usually recursive, or at worst recursively enumerable operations. But it will turn out that the final consequence operation is substantially more complicated than that.

We shall argue that these complexity results for adaptive logic have philosophical implications. They cast some doubt on Batens' philosophical thesis that adaptive consequence closely reflect how people *actually* reason on the basis of inconsistent theories and that provisional, finite adaptive proofs provide an ever improving approximation of the final consequences of a set of premises. And these results have consequences for humanly attainable convergence to the truth in the infinite limit concerning adaptive derivability questions.

2 Propositional adaptive logic

The system which we present first is close to the basic inconsistency-adaptive logic called **ACLuN1**, which was formulated in Batens (1999) and which is referred to in the literature as the *reliability strategy*. We shall refer to it as the *Reliability Calculus*. Here we concentrate on the *propositional fragment* of **ACLuN1**.

In order to facilitate metamathematical analysis, we describe a system which is equivalent to the propositional fragment of **ACLuN1**. We shall call our system **AL1**. The (non-essential) differences with **ACLuN1** are highlighted as we go along. In our discussion, we also refer to general features of architectures for adaptive logic.[1]

2.1 The architecture

We work in a propositional language $\mathcal{L}$ which contains the connectives $\neg, \wedge, \vee, \rightarrow$. As is typical for adaptive logics, our logical system **AL1** contains an *upper limit logic*, a *lower limit logic* and a collection of *abnormalities*. As upper limit logic, we take a Hilbert-style formalization of classical propositional logic. As lower limit logic, we take a Hilbert-style formalization of propositional paraconsistent logic. For definiteness,

[1] On this score, a useful article is Batens (1999). We will have occasion to refer to it several times.

let the lower limit logic be the propositional fragment of the system **CLuN**, as presented in Batens (1999, p. 452). We shall describe **CLuN** in more detail in the next subsection. As our *abnormalities*, we take disjunctions of outright contradictions.

An adaptive proof is a *finite* sequence of 6-tuples $\langle l_{i_1}, \ldots l_{i_6}\rangle$ such that:

1. l_{i_1} is the line number;
2. l_{i_2} is a sentence of $\mathcal{L}$: it is the sentence which is *derived* at line l.
3. l_{i_3} is a rule of the upper limit logic or of the lower limit logic;
4. l_{i_4} indicates the lines on which the inference depends;
5. l_{i_5} lists the formulas which must be assumed to behave consistently for the inference to be reliable.[2]
6. l_{i_6} contains the line numbers of derived formulas which cause the inference to be *marked*, i.e., judged to be unreliable.

This account is exactly like that of Batens (1999), except that the marking situation is explicitly taken to be part of a line in a proof.[3]

2.2 A weak paraconsistent logic

As the lower limit logic, we take the simple paraconsistent system **CLuN**, or, to be absolutely precise, its propositional fragment.

Proof-theoretically, **CLuN** is characterized as follows. The axioms and rules of **CLuN** are those of *positive logic* plus the principle of excluded third $(\phi \vee \neg\phi)$.[4] The axioms and rules of positive logic are exactly those of classical propositional logic, except for those axioms and rules which govern the behavior of the negation operator: those are omitted. So the idea behind **CLuN** is that the meaning of all the propositional logical connectives is exactly like that in classical propositional logic, except for negation. All that is postulated for negation is that every formula must have *at least one* of the truth-values *True* and *False*.

Semantically, **CLuN** can best be characterized as follows. Consider the variant on truth-tables, called *-tables, where the rules for evaluating the logical connectives $\vee, \wedge, \rightarrow$ are just as usual, except for negation ($\neg$). In *-tables, *all* negated formulas (so *not* just negations of proposition letters!) are treated as if they were atomic proposition letters, except that of each pair $\phi, \neg\phi$ at least one has to be assigned the value *True*. Then we say that a formula ϕ follows semantically from a finite set $\phi_1, \ldots, \phi_n$ (denoted as $\phi_1, \ldots, \phi_n \models_{CLuN} \phi$)) if and only if in a *-table for $\phi_1, \ldots, \phi_n, \phi$, on all lines where $\phi_1, \ldots, \phi_n$ are true, ϕ also comes out true.

It may be useful at this point to present a simple example of a *-table. Suppose we want to know hether $p \vee q, \neg p \models_{CLuN} q$. Then we write a *-table as follows:

$p \vee q$	$\neg p$	p	q
1	1	1	1
1	0	1	1
1	**1**	**1**	0
1	0	1	0
1	1	0	1

[2] We are using the words ‘consistently’ and ‘reliable’ in an informal sense here. What this precisely amounts to is specified below.

[3] Compare the above with Batens (1999, pp. 456–457).

[4] A detailed description of **CLuN** is given in Batens (1999, Sect. 3).

First, observe that in this *-table $\neg p$ is treated as if it were an *atomic* proposition letter. Second, this *-table shows that the inference is not semantically valid in **CLuN**. For on the third row, both $p \vee q$ and $\neg p$ are true, whereas q is false. But on this row, both p and $\neg p$ receive the same truth value, namely *True*. The adaptive logician would summarize the situation as follows: q does follow from $p \vee q$ and $\neg p$ *as long as* we assume that p does not behave in an abnormal way—which it does on the third line of the *-table. So it is not the case that $p \vee q, \neg p \models_{CLuN} q$.

There is a completeness property which connects the *-tables with the axioms and rules of **CLuN**. This is implicit in (Batens, 1999, p. 452):

Proposition 1 *For all* $\phi_1, \ldots, \phi_n, \phi$*:*

$$\phi_1, \ldots, \phi_n \vdash_{CLuN} \phi \Leftrightarrow \phi_1, \ldots, \phi_n \models_{CLuN} \phi$$

Since the semantic consequence relation based on the *-tables is obviously a decision method, the proof system for **CLuN** is a decidable calculus.

There also exists an effective relation between derivability in the classical propositional calculus **CL** and derivability in **CLuN**.

Proposition 2 *Suppose we have* $\phi_1, \ldots, \phi_n \vdash_{CL} \phi$*. Then in its *-table, there may be lines on which the premises are true while the conclusion is false. But according to such lines, there must be formulas* ψ_i *such that both* ψ_i *and* $\neg\psi_i$ *receive truth value* 1*. Then if* $\psi_1, \ldots, \psi_k$ *are all such formulas, it follows that*

$$\phi_1, \ldots, \phi_n \vdash_{CLuN} \phi \vee (\psi_1 \wedge \neg\psi_1) \vee \ldots \vee (\psi_k \wedge \neg\psi_k).$$

2.3 Writing down a line in an adaptive proof

A theory Γ is, as usual, a collection of sentences. On any line in an **AL1**-proof from a theory Γ, one is (of course) allowed to write down a sentence of Γ. It is then written down after a line number, annotated as a *Premise* (entry 3), dependent on no earlier lines (entry 4 is empty), assuming no sentences to behave consistently (entry 5 is empty). The description of the marking instructions (entry 6) is deferred to the next subsection.

We are also allowed to derive a sentence from earlier lines. The idea is to write down a sentence on a line of a proof in accordance with a rule of the upper limit logic, i.e., classical logic, while keeping track (in the fifth entry) of the consistency of the formulas on which this depends. This idea is implemented in the following way. Suppose ϕ follows *classically* from earlier derived sentences $\phi_1, \ldots, \phi_n$. Then according to proposition 2 above, we effectively find a set of formulas $\psi_1, \ldots, \psi_k$ such that

$$\phi \vee (\psi_1 \wedge \neg\psi_1) \vee \ldots \vee (\psi_k \wedge \neg\psi_k)$$

follows in **CLuN** from $\phi_1, \ldots, \phi_n$. In such a situation we may write down, on a new line, ϕ. We indicate which rule is used (entry 3), from which lines it is derived (entry 4), and *which formulas are assumed to behave consistently* (entry 5). Proposition 2 guarantees that there is an effective procedure for doing this. In this fifth entry we collect the fifth entries of the lines on which $\phi_1, \ldots, \phi_n$ are derived *plus* $\psi_1, \ldots, \psi_k$.

Let us consider a simple example of a derivation in our propositional adaptive logic. Suppose we define our theory Γ_1 thus:

$$\Gamma_1 = \{p \vee q, \neg q\}$$

Then a simple derivation from Γ_1 looks like this:

1	$p \vee q$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
2	$\neg q$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
3	p	DS	1,2	q	$\emptyset$

Here DS stands for 'Disjunctive Syllogism'. However, in the sequel we will not be concerned with exactly which principle of classical propositional logic is used: such details need not detain us. Note that on line 3, the derivation of p depends on the consistency of q. For if both q and $\neg q$ were true, then the first two lines could be true while line 3 is false. Incidentally, it is easily verified that the sentence $p \vee (q \wedge \neg q)$ can be derived in the weak lower limit logic **CLuN** from lines 1 and 2.

2.4 The marking rules of the Reliability Calculus

The marking rules which we shall now describe are distinctive of the *reliability* strategy.[5]

Definition 1 $A(p_1, \ldots, p_n)$ abbreviates $(p_1 \wedge \neg p_1) \vee \ldots \vee (p_n \wedge \neg p_n)$, and is called an abnormality. The abnormality set of $A(p_1, \ldots, p_n)$ is $\{p_1, \ldots, p_n\}$.

Sometimes we shall write $A(\overrightarrow{\phi})$ instead of $A(\phi_1, \ldots, \phi_n)$. We shall not always be careful to distinguish an abnormality from its associated abnormality set, but such confusions will be innocuous.

Suppose an abnormality $A(p_1, \ldots, p_n)$ is derived *unconditionally* at a line l. Then $\{p_1, \ldots, p_n\}$ is the abnormality set associated with $A(p_1, \ldots, p_n)$. As long as no unconditional abnormality $A(q_1, \ldots, q_k)$ is derived on a line k before or after line l such that $\{q_1, \ldots, q_k\} \subset \{p_1, \ldots, p_n\}$, the abnormality $A(p_1, \ldots, p_n)$ is said to be a *minimal abnormality*. And if a derived formula on a line m depends on formulas some of which belong to an abnormality set associated with a *minimal* abnormality derived on a line l, line number l is inserted as a mark in the sixth entry of line m, indicating unreliability of this inference step.[6] And *markings are never removed*. So even if, e.g., the abnormality on line l later becomes non-minimal, line m remains marked.

This brings the marking behavior in **AL1** close to that of **ACLuN1**. With our conventions, we strictly speaking depart from Batens' marking conventions: he decrees that markings are to be *erased* once the abnormality responsible for the marking becomes non-minimal.[7] In our set-up, the markings are not erased. But it is, in such a situation, possible to simply again re-derive the sentence on line m on another line in a further stage of the proof, in such a way that it is not marked by line l, for it no longer is a *minimal* abnormality. In general, extending a proof leads to revision upward in the proof, but only in one sense: markings are possibly added but no markings are deleted. So the difference between our marking rules and those of **ACLuN1** is inessential.

We emphasize here (firstly) that minimal abnormalities only cause markings if they are derived unconditionally or *categorically*, i.e., when the fifth entry on the line is

[5] See Batens (1999, p. 457).

[6] Again we are using 'unreliability' in the informal sense here, as opposed to the technical sense that this word has in the literature on adaptive logics.

[7] See Batens (2001, p. 60).

empty.[8] Secondly, note that adaptive logics are not a system of "belief revision" in the sense of Gärdenfors,[9] for *premises* are never retracted (or "marked") in the process of adaptive reasoning. And thirdly, we note here that adaptive proofs are allowed to be of transfinite length. We will come back to this last point later in the paper.

Let us extend the example from the previous subsection a bit. We first extend our theory Γ_1 to:

$$\Gamma_2 = \{p \vee q, \neg q, (q \wedge \neg q) \vee (r \wedge \neg r), r \wedge \neg r\}.$$

Now we can continue the adaptive proof of the previous subsection in a way that brings the marking rules into play:

1	$p \vee q$	*Pr*	$\emptyset$	$\emptyset$	$\emptyset$
2	$\neg q$	*Pr*	$\emptyset$	$\emptyset$	$\emptyset$
3	p	*DS*	1,2	q	4
4	$(q \wedge \neg q) \vee (r \wedge \neg r)$	*Pr*	$\emptyset$	$\emptyset$	$\emptyset$
5	$r \wedge \neg r$	*Pr*	$\emptyset$	$\emptyset$	$\emptyset$
6	p	*DS*	1,2	q	$\emptyset$

Line 4 causes line 3 to be marked, for it gives reason to doubt that q is reliable. But the more specific contradiction on line 5 removes these doubts about q, thus making abnormality on line 4 non-minimal. This implies that on line 6 we can derive p again. And this time we cannot derive any sentence which causes line 6 to be marked.

3 Final derivability

In adaptive proofs, conclusions are usually *provisional*, in the sense that by extending the proof one may come to know that they are unreliable after all. But there is also a notion of conclusive or *final derivability*. This will be the proof-theoretic consequence relation of the Reliability Calculus, to which we shall now turn.

3.1 Final derivability in the Reliability Calculus

We define the notion of *extension* of a proof (denoted $\mathcal{P} \sqsubseteq \mathcal{Q}$) in the way that one would expect. If $\mathcal{P}$ is a proof, then $\mathcal{Q}$ is an extension of $\mathcal{P}$ if and only if $\mathcal{Q}$ can be obtained by continuing the proof $\mathcal{P}$ by adding new lines in accordance with the rules of the adaptive logic. This induces a difference with the Batens architecture. For Batens, when a line is *inserted* in a proof $\mathcal{P}$, say between lines l and $l+1$, the resulting proof also counts as an *extension* of $\mathcal{P}$.[10] We will briefly return to this point in the next section.

We define the notion of final derivability for **AL1** as follows:

Definition 2 A formula ϕ is finally derivable from a set of premises Γ according to the Reliability Calculus if and only if there is a proof $\mathcal{P}$ of ϕ from Γ on a certain line l, and this proof cannot be extended to a proof $\mathcal{Q}$ in which the sixth entry of line l is nonempty.

---

8 See Batens (2001, p. 50–51).

9 See Gärdenfors (1988).

10 See Batens (1999, p. 466, fn 20), which is situated in the context of the system **ACLuN2**.

This differs from the concept of final derivability that is formulated by Batens (1999, p. 458–459, 466; 2001, p. 61). For Batens, a formula ϕ is finally derivable if and only if there is a derivation of ϕ, and every extension of this proof in which ϕ is marked can be extended to one in which it occurs unmarked. But this definition is not even open to us, since we do not allow marks to be erased. However Propositions 4 and 5 below will show that the two notions of final derivability coincide extensionally.

Note that for a formula to be finally derived on a line in an adaptive proof, it is not necessary that it has an empty dependency set, i.e., it is not necessary that the fourth entry of the line on which it is derived is empty.

This concludes the description of the system **AL1**.

3.2 Comparison with Baten's infinite proofs

We allow proofs in **AL1** to be of infinite or even transfinite length. In Batens' reliability system **ACLuN1** (and in **ACLuN2**) adaptive proofs have order type at most ω.[11] There are cases, in Batens' way of setting things up, in which extensions of infinite proofs need to be considered.[12] But in such cases, the extension is constructed by *inserting* a new line in the infinite proof (and renumbering), not by appending it.[13] And this of course generates a new proof of order type ω.

Our notion of transfinite proofs seems to be more closely associated with Batens' notion of *stages* of a proof:

> There is a 'deeper' account of the notion of proof. On this account, a stage of a proof is a sequence S of lines and a proof is a sequence or chain Σ of stages. In all cases that interest us here, proofs start from stage zero, which is the empty sequence... [A] stage is obtained by extending the previous stage, but possibly with the marks [i.e., numbers] of its lines changed, with exactly one line. (Batens, 2001, p. 62)

These chains can be of transfinite ordinal type, even though the proofs of which the chain is composed are at most of order type ω. Our suggestion is that the transfinite nature of the generation procedure of a proof be reflected in the ordinal type of the length of the proof. This is not to say that there is anything wrong with Batens' way of setting things up. We merely reformulate Batens' notion of adaptive proof in this way for diagnostic purposes.

Despite the superficial differences mentioned above, Batens' notion of final derivability from a set of premises extensionally coincides with our notion of final transfinite derivability, i.e., final derivability in **AL1**, as we shall now show.

When we speak about B-proofs, we mean proofs according to the definitions of Batens. When we mean proofs in our sense, we will just speak about proofs.

Let us start by listing the definitions that agree with Baten's approach:

Definition 3 A B-proof $\mathcal{P}$ of ψ from Γ can be seen a sequence of length $\leq \omega$ in which ψ occurs on a line in accordance with the requirements for being an **AL1**-proof except for the marking rule: in a B-proof, a mark l is associated with a line k if and only if at line l a categorical abnormality appears which is minimal, in the sense that no abnormality occurs anywhere in $\mathcal{P}$ which is more specific.

[11] See Batens (2001, p. 62).

[12] See Batens (1999, p. 466).

[13] See Batens (2001, p. 62, fn 33).

So if at some line $l + i$ a more specific abnormality appears but not before that, then there is no mark on line k in the B-proof, whereas there would be a mark in an **AL1**-proof.

Definition 4 A B-extension $\mathcal{P}'$ of a B-proof $\mathcal{P}$ is a B-proof that results from appending or inserting lines to/into $\mathcal{P}$.

Inserting lines may necessitate us to renumber lines and marking labels, and remove and add marking labels, which is somewhat awkward.

Definition 5 A formula ψ is finally B-derivable from Γ if and only if there is a B-proof $\mathcal{P}$ of ψ from Γ where ψ occurs unmarked on a line l such that every B-extension $\mathcal{P}'$ of $\mathcal{P}$ in which this occurrence of ψ is marked can be further B-extended so that it becomes unmarked again.

Batens' notion of final consequence can be brought closer to ours:

Proposition 3 *A sentence ψ is finally B-derivable from Γ if and only if there is a B-proof $\mathcal{P}$ of ψ from Γ in which ψ occurs unmarked and which cannot become marked by B-extending $\mathcal{P}$.*

Proof The right-to-left direction is immediate. For the left-to-right direction, let a B-proof be given that is a witness of the truth of the B-derivability of ψ. Derive all categorical abnormalities and insert and / or append them to $\mathcal{P}$. Call the resulting proof $\mathcal{P}'$. This is the proof that we are looking for: it will contain ψ unmarked and cannot be extended in such a way that it becomes marked. □

In virtue of this proposition, let us call a witness of final B-*derivability of ψ from Γ a* final B-proof of ψ: so this is a proof in which ψ occurs unmarked and which cannot be extended so that it becomes marked.

Proposition 4 *Every final B-proof of ψ from Γ can be transformed into a final **AL1**-proof.*

Proof Let there be given a final B-proof of ψ from Γ of length α. It must be an **AL1**-proof except that some sentences may be unduly unmarked. So we add the appropriate marks in accordance with the marking rule of **AL1**. It is possible that this results in the occurrences of ψ in the proof becoming marked. But then we can re-derive ψ after line α in such a way that it is unmarked and cannot be marked by extending. (Note that this may require us to introduce transfinite lines, but this is allowed by **AL1**.) □

Proposition 5 *Every final **AL1**-proof of ψ from Γ can be transformed into a final B-proof of ψ from Γ.*

Proof Let the final **AL1**-proof $\mathcal{P}$ be given. First we derive all categorical disjunctions of abnormalities in ω steps. Then we interleave this derivation with the proof $\mathcal{P}$, removing marks where this needs to be done. The result will be a final B-proof of ψ. □

For consistent theories, no disjunctions of abnormalities can ever be derived. Hence in such situations retraction of derived sentences (i.e., marking) is never necessary. Therefore for consistent theories, theorems can always be considered to be finally established after a finite number of stages. For some inconsistent (but recursive) theories, only an infinite derivation can witness that a sentence is finally derived.

3.3 An example

Batens asserts that in the definition of the notion of final **ACLuN1**-consequence, there is no need to refer to infinite proofs.[14] But this is not so. We will construct an example of a sentence which becomes finally derivable only at stage ω.

Definition 6 The derivability ordinal for a theory Γ of a sentence ϕ is the minimal ordinal stage (line) on which ϕ can be finally derived from Γ.

Consider the recursive theory Γ_3:

$$\Gamma_3 = \{p \vee q, \neg q, ((q \wedge \neg q) \vee (r_i \wedge \neg r_i)), ((q \wedge \neg q) \vee (r_i \wedge \neg r_i)) \rightarrow (r_i \wedge \neg r_i)\}_{i \in \omega}.$$

The shortest proof from Γ_3 in which p is finally derived is of the following type:

1	$p \vee q$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
2	$\neg q$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
3	p	DS	1,2	q	4
4	$(q \wedge \neg q) \vee (r_1 \wedge \neg r_1)$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
5	$((q \wedge \neg q) \vee (r_1 \wedge \neg r_1)) \rightarrow (r_1 \wedge \neg r_1)$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
6	$r_1 \wedge \neg r_1$	MP	4,5	$\emptyset$	$\emptyset$
...	...	...	...	...	...
k	p	DS	1,2	q	$k+1$
$k+1$	$(q \wedge \neg q) \vee (r_i \wedge \neg r_i)$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
$k+2$	$((q \wedge \neg q) \vee (r_i \wedge \neg r_i)) \rightarrow (r_i \wedge \neg r_i)$	Pr	$\emptyset$	$\emptyset$	$\emptyset$
$k+3$	$r_i \wedge \neg r_i$	MP	$k+1, k+2$	$\emptyset$	$\emptyset$
...	...	...	...	...	...
ω	p	DS	1,2	q	$\emptyset$

Note that it is impossible to derive p from Γ finally at any finite stage: at finite stages there are always minimal abnormalities with which the line on which p is derived can be marked. But at stage ω, all abnormalities have become non-minimal. So line ω is unmarked and cannot be marked by extending the proof. We conclude from this that p is finally derived only at line ω.

Our example highlights a way in which the notion of final derivability in adaptive logics differs essentially from the notion of derivability in standard logical systems. For standard logical systems, it is clear that the derivability ordinal of a final consequence of a theory is always smaller than ω, whereas we now see that even for recursive theories the derivability ordinals can be transfinite in the case of adaptive logics. However, the derivability ordinals of theories are never highly transfinite:

Proposition 6 *For all Γ, ϕ, if ϕ can be finally derived from Γ, it can be so in a proof of length $\leq \omega$.*

Proof Consider any formula ϕ which is derivable from Γ, possibly only conditionally. Begin by deriving, at finite stages, all categorical disjunctions of abnormalities, and make sure that ϕ is derived at stage ω. Then ϕ is unmarked at line ω if and only if it is finally derivable. □

---

14 See Batens (2004, p. 479).

3.4 The Minimal Abnormality Calculus

Until now, we have investigated **ACluN1**, or the reliability calculus, in a reformulated version. But Batens has formulated a second strategy of adaptive reasoning, resulting in a second notion of final consequence. This second strategy is more complicated, and it is called **ACluN2**, or *the Minimal Abnormality Calculus.*

An example will convey the idea of the Minimal Abnormality Calculus. Consider the following theory:

$$\Gamma_4 = \{A(p,q,r), p \wedge q \wedge r, \neg p \vee \neg q \vee s, \neg q \vee \neg r \vee s, \neg r \vee \neg p \vee s\}$$

Consider the following derivation, and consider it from the point of view of the rules of adaptive logic that we have applied so far:

1.	$A(p,q,r)$	$\emptyset$	$\emptyset$
...	...	...	...
k.	s	p,q	1
l.	s	q,r	1
m.	s	r,p	1

$A(p,q,r)$ is a minimal abnormality, so it causes lines k,l,m to be marked. So according to the adaptive logic that we have studied so far ("reliability strategy") s is not finally derivable from Γ_4.

$A(p,q,r)$ requires that *at a minimum*, one of p,q,r is involved in an inconsistency. But take any *minimally inconsistent situation*. Suppose, for instance, that p "behaves inconsistently", but q,r behave consistently. On line l, s is derived without relying on the unreliable p. Similarly if q is the inconsistent one (then line m does it) and if r is the inconsistent one (then line k does it). So one would say that in all minimally inconsistent situations, s holds. So if the final consequences are to describe these minimally inconsistent situations, then s ought to be finally derivable.

This motivation leads to a new proof system, associated with which is a new concept of final derivability. We shall express this new concept of final derivability without getting into the details of "provisional derivability", i.e., without explaining the new rules for writing down a line in a proof (which are rather unwieldy[15]).

Definition 7 Given a set S of finite sets S_i of proposition letters, a $\subseteq$-minimal set which contains at least one proposition letter from each S_i is called a selection set.

So suppose one has a set of abnormalities. Then one can consider a *selection set over this set of abnormalities*: such a selection set will select at least one formula ϕ_i from each abnormality $A(\phi_1, \phi_2, \ldots, \phi_i, \ldots, \phi_k)$ in the set, and this selection set will be minimal in the partial ordering induced by $\subseteq$.

Now we are ready to express the concept of final derivability for minimal abnormality. We call the resulting *Minimal Abnormality Calculus* **AL2**:

Definition 8 A formula ϕ is finally derivable from a theory Γ according to the minimal abnormality calculus if and only if for every selection set Ψ over the set of all minimal abnormalities that can be derived categorically from Γ, there exists a derivation of ϕ on a condition $\phi_1, \phi_2, \ldots, \phi_k$ such that $\Psi \cap \{\phi_1, \phi_2, \ldots, \phi_k\} = \emptyset$.

15 Batens acknowledges this. See, e.g., Batens (2001, p. 60).

In other words, the idea is that ϕ is finally derivable if no matter which *minimal* way the model is inconsistent, ϕ comes out true. So, when applied to the example above, s will in this new sense be finally derivable from Γ_4 even though there is no way that the derivation can be extended so that s occurs unmarked.

As with our discussion of the reliability strategy, our definition of final derivability for minimal abnormality is not exactly the same as it is given by Batens' definition which can be found, e.g., in Batens (2001, p. 60–61). But our definition is equivalent to Baten's definition. To verify this is tedious but straightforward.[16]

4 The complexity of final derivability

We now turn to the question of the complexity of the final consequence relations for the calculi **AL1** and **AL2**.

Evidently, the question whether a sentence is finally derivable in the calculi **AL1** and **AL2** from a given finite premise set is always decidable.[17] (We shall see however that this is false for predicate versions of these calculi.) But, as we have seen, Batens' calculi are also intended for infinite premise sets. And that is just as well, for many of our most fundamental real-world theories (Peano Arithmetic, Tarski's theory of truth, ...) are not finitely axiomatizable. So let us concentrate on infinite premise sets.

For the Reliability Calculus, Batens formulates a *conjecture* concerning the decidability of the relation of final consequence in a propositional context:[18]

> [...] there is also a decision method for finite (*and most plausibly also for infinite*) Γ in the propositional fragment of [**ACLuN1**] [...] (Batens, 1995, p. 316, our emphasis)

Obviously not every infinite set of sentences will do here. What is meant, is a set of premises that expresses a *theory*. Classically, a theory can be seen as the class of propositions that can be derived (using the laws of classical logic) from a recursive set of axioms. So a theory can be taken to be recursively enumerable set of sentences. But in adaptive logic, the derivability relation is more complicated. So we should focus on the axioms and insist that a theory is a (possibly infinite but) *recursive* set of sentences. So Batens may be taken to conjecture here that the final derivability relation of **ACLuN1** between recursive sets of premises and formulas is decidable.

But already the collections of classical propositional consequences of an infinite but recursive set of propositional formulas can form a complete Σ^0_1 set.[19] The argument showing this can serve as a warm-up for the technical results that will be proved in the following sections. Let C be a complete Σ^0_1 set: say it is the set of all natural numbers y such that $\exists x P(x,y)$, with P a recursive two-place predicate. Now consider the theory Γ_5 which consists of an infinite sequence of proposition letters $s_{x,y}$ where x and y range over all natural numbers, plus an infinite sequence of propositional axioms of the form $s_{x,y} \rightarrow p_y$ for all x and y such that $P(x,y)$. Clearly Γ_5 forms a recursive set of formulas. And it is easy to see that for every natural number y, p_y follows propositionally from

[16] Thanks to Kristof De Clercq for checking this for us.

[17] This was first proved in Batens (1995).

[18] See also Batens (2004, p. 480; 2005a, p. 85, footnote 5).

[19] Thanks to a referee for drawing our attention to this fact.

Γ_5 if and only if $y \in C$. So the collection of classical propositional consequences of Γ_5 indeed is a complete Σ^0_1 set.

Given this fact, it is clear that Batens' conjecture must fail. The theory Γ_5 is consistent. Therefore its final **ACLuN1**-consequences are just its classical propositional consequences. So the collection of final consequences according to the Reliability Calculus of Γ_5 already form a counterexample to the conjecture.

Then the question arises *how badly* Batens' conjecture fails. At the outset, it is not obvious that the final consequence set of a recursive set of premises can at worst be complete Σ^0_1. Steffo Weber, for instance, has shown that for *some* propositional non-monotonic logics, the consequence set of a recursive set of premises can be Σ^0_2-hard (Weber, 2000, p. 311).

In fact, we shall demonstrate that the collection of final consequences of a recursive propositional theory can be rather computationally complex. First, we establish upper bounds by looking at final derivability according to the Reliability and the Minimal Abnormality Calculus. Then we shall show that these bounds are best possible by showing that membership in Σ^0_3-complete sets can be derived from certain final derivability sets (using either calculus).

4.1 The reliability strategy: upper bounds

In the following, we shall make use of the concept of a *universal derivation* from a recursive collection of propositional premises:

Definition 9 The universal derivation $UD(\Gamma)$ from Γ is a list of lines of length ω such that for every ϕ, $\psi_1, \ldots, \psi_n$, if ϕ is classically derivable from Γ on the condition that $\psi_1, \ldots, \psi_n$ behave consistently, there is a line in $UD(\Gamma)$ witnessing this, i.e., there is a line in $UD(\Gamma)$ on which ϕ is derived on the condition $\{\psi_1, \ldots, \psi_n\}$.

$UD(\Gamma)$ can be taken to be given canonically in Γ. Since Γ shall be assumed to be a recursive set of premises, $UD(\Gamma)$ will be a recursively enumerable list of lines.

Theorem 1 *The notion of final consequence according to the reliability strategy is a Σ^0_3 concept.*

Proof Let a set of premises Γ be given. So, if Γ is a recursive set of sentences, then the question whether $UD(\Gamma)$ contains a line with number k on which ϕ is derived on condition $\psi_1, \ldots, \psi_n$ is decidable. Then the question whether a sentence ϕ is *finally derivable* can be expressed in terms of a search in $UD(\Gamma)$. Somewhat more specifically, φ is finally derivable if and only if:

$\exists n \phi_1, \ldots, \phi_k$
[(line $l_n \in UD(\Gamma)$ is of the form $\varphi \quad \{\phi_1, \ldots, \phi_k\}$)
$\wedge \forall j$(if line j of $UD(\Gamma)$ is the abnormality σ then [(σ is not unconditional)
or (if σ contains one of the $\phi_1, \ldots, \phi_k \longrightarrow$
$\exists m$ ($l_m \in UD(\Gamma)$ is an unconditional abnormality showing that
σ is non-minimal))])] .

As "$l_n \in UD(\Gamma)$" is Σ^0_1, the result follows. □

4.2 The minimal abnormality strategy: upper bounds

Computing the complexity of the final derivability relation for Minmal Abnormality is more complicated. We first show that the notion of being a final minimal abnormality consequence can be expressed as a Σ^0_3 operator. Then we show that the set of final minimal abnormality consequences of a decidable propositional theory can be a complete Σ^0_3 set. To conclude, we show that in a predicate logical setting, this can already happen from a *finite* set of premises.

Let S be a family of non-empty finite sets of propositional atoms. Recall that a *selection set* Ψ is a choice set for S which is $\subseteq$-minimal amongst all possible such choice sets. As before, $UD(\Gamma)$ is the universal derivation of any line that is derivable from the axioms Γ.

Definition 10 If ϕ is derived on a line of $UD(\Gamma)$, conditional on $\phi_1, \ldots \phi_k$, then we write $\phi \; \{\phi_1, \ldots \phi_k\}$.

If ϕ is derived unconditionally from Γ, i.e., conditional on the empty set, then we say that ϕ is *categorically derived*, and write $\Gamma \vdash_{CLuN} \phi$.

In the light of what was said before, we can reformulate the notion of final derivability on the minimal abnormality strategy.

Definition 11 $S_\Gamma = \{\{\phi_1, \ldots \phi_k\} \mid \Gamma \vdash_{CLuN} A(\phi_1, \ldots \phi_k)\}$.

Definition 12 We denote the fact that ϕ is finally derivable from Γ as $\Gamma \vdash_{AL2} \phi$, where:

$$\begin{aligned}\Gamma \vdash_{AL2} \phi \quad &\Leftrightarrow \forall \text{ selection set } \Psi \text{ over } S_\Gamma \exists \text{ derivation from } \Gamma \text{ of } \phi \; \{\phi_1, \ldots \phi_l\} \\ &\qquad \text{with } \Psi \cap \{\phi_1, \ldots \phi_l\} = \emptyset \\ &\Leftrightarrow \forall \text{ selection set } \Psi \text{ over } S_\Gamma \exists \text{ line of } UD(\Gamma) \text{ of the form } \phi \; \{\phi_1, \ldots \phi_l\} \\ &\qquad \text{with } \Psi \cap \{\phi_1, \ldots \phi_l\} = \emptyset\end{aligned}$$

Now we shall first show that the relation of final derivability is Π^1_1. Then we shall reduce this relation to being even Σ^0_3.

We shall from now on assume that Γ is recursive, i.e., that the premise set is decidable.

Proposition 7 *"$t \in S_\Gamma$" $\in \Sigma^0_1$*

Proof $t \in S_\Gamma$ if and only if there exist formulas $\phi_1, \ldots, \phi_k$ such that $t = \{\phi_1, \ldots \phi_k\}$ and there exists a finite list of lines from $UD(\Gamma)$ witnessing $\Gamma \vdash A(\overrightarrow{\phi})$. This is a Σ_1 search through $UD(\Gamma)$. □

Definition 13 We write $A(\overrightarrow{\phi}) \in CDA$ ("A is a categorical derived abnormality") if $A(\overrightarrow{\phi}) \in S_\Gamma$.

So we have shown that "$A(\overrightarrow{\phi}) \in CDA$" is Σ^0_1.

Definition 14 We say that $A(\overrightarrow{p}) \in MIN(CDA)$ if for all $q \in \overrightarrow{p}$ it is the case that $A(\overrightarrow{p} \backslash \{q\}) \notin CDA$

Then we immediately see that:

Proposition 8 *"$A(\overrightarrow{\phi}) \in MIN(CDA)$" $\in \Pi^0_1$*

Definition 15 Let $\mathcal{C}_\phi$ be the family of conditional sets on lines in which ϕ is derived in $UD(\Gamma)$.

Then we immediately see that:

Proposition 9 *"$C \in \mathcal{C}_\phi$" $\in \Sigma^0_1$*

Therefore:

Proposition 10 *"Ψ is a choice set for S" is Π^0_1 in S and Ψ.*

Proposition 11 *"Ψ is a selection set for S" is Π^0_2 in S and Ψ.*

Proof Ψ is a selection set for S if and only if Ψ is a choice set for S and all proper subsets of Ψ are not choice sets for S. This last conjunct is equivalent to $\forall t \in \Psi(\exists s \in S(\Psi \backslash \{t\} \cap s = \emptyset))$ □

Putting all these facts together, we see that:

Proposition 12 *"$\Gamma \vdash_{AL2} \phi$" $\in \Pi^1_1$*

Proof $\Gamma \vdash_{AL2} \phi$ if and only if:

$$\forall\Psi\forall S[S = S_\Gamma \wedge \Psi \text{ a selection set for } S \rightarrow$$
$$\exists C(\text{“ } \varphi \text{ occurs on a line of } UD(\Gamma) \text{ conditional on } C \text{ and } \Psi \cap C = \emptyset\text{”})]$$

□

Now we shall reduce this to Σ^0_3. The following is an important observation:

Proposition 13 *If Ψ is a choice set for S_Γ, then:*

$$\Psi \text{ is a selection set for } S_\Gamma \Leftrightarrow \Psi \subseteq \bigcup MIN(CDA)$$

Proof Let Ψ be a choice set. It suffices to prove that if it is also a selection set for S_Γ, then the right hand side holds. Then $\Psi \cap \bigcup MIN(CDA)$ is a selection set over $MIN(CDA)$. But that is sufficient for $\Psi \cap \bigcup MIN(CDA)$ to be a selection set over S_Γ. □

We now define a finitely branching tree, $\mathcal{T}$, of sequences of propositional atoms. The propositional atoms in the sequence are intended as initial segments of a putative selection set Ψ over $MIN(CDA)$, which will, if defined, demonstrate that $\Gamma \nvdash_{AL2} \phi$. For this to be the case we must have $\forall C \in \mathcal{C}_\phi, C \cap \Psi \neq \emptyset$.

Note, incidentally, that if for some $C \in \mathcal{C}_\phi$, $C \cap \bigcup MIN(CDA) = \emptyset$, then any selection set for S_Γ is, by our last proposition, disjoint from C. Hence in this case $\Gamma \vdash_{AL2} \phi$.

We define $\overrightarrow{\phi} = \langle \phi_0, \ldots, \phi_l \rangle \in \mathcal{T}$ by induction on l. We first fix an enumeration $C_0, \ldots, C_k, \ldots$ of $\mathcal{C}_\phi$.

Definition 16
- If $l = 0$, then $\langle \phi_0 \rangle \in \mathcal{T}$ if and only if $\phi_0 \in C_0 \cap \bigcup MIN(CDA)$
- If $l = k + 1$ and $\overrightarrow{\phi} = \langle \phi_0, \ldots, \phi_k \rangle \in \mathcal{T}$, then $\langle \phi_0, \ldots, \phi_k, \phi_{k+1} \rangle \in \mathcal{T}$ if $\phi_{k+1} \in C_{k+1} \wedge (\exists l \leq k(\phi_{k+1} = \phi_l)$ **or** $\exists A \in MIN(CDA)[\phi_{k+1} \in A \wedge \{\phi_0, \ldots, \phi_k\} \cap A = \emptyset \wedge p_{k+1} \in C_{k+1}])$.

Proposition 14 "$\langle\phi_0,\dots,\phi_k\rangle \in \mathcal{T}$" $\in \Sigma^0_2$

Proof

$$\langle\phi_0,\dots,\phi_k\rangle \in \mathcal{T} \Leftrightarrow$$
$$\exists g[g \text{ is a function } \wedge dom(g) = k+1 \wedge \forall l < k+1(g(l) \in MIN(CDA)) \wedge$$
$$\forall l < k+1(\phi_l \in C_l \wedge (\exists j < l(\phi_l = \phi_j) \vee (\phi \in g(l) \wedge \{\phi_0,\dots,\phi_{l-1}\} \cap g(l) = \emptyset)))]$$

The first line of the existentially quantified formula has complexity Π^0_1 due to "$g(l) \in MIN(CDA)$". The rest is bounded Σ^0_1, so Σ^0_1. □

As $\bigcup_{j\leq k} C_j$ is finite there are at most finitely many possibilities to extend $\langle\phi_0,\dots,\phi_k\rangle$. Hence the tree $\mathcal{T}$ is finitely branching.

Proposition 15 *Suppose* $\langle\phi_0,\dots,\phi_k\rangle \in \mathcal{T}$ *is maximal. Then there is no selection set* Ψ *such that (i)* $\{\phi_0,\dots,\phi_k\} \subseteq \Psi$ *and (ii)* $\Psi \cap C_{k+1} \neq \emptyset$.

Proof This is just what maximality means: for any $\phi \in C_{k+1}$ we must have

$$\phi \notin \bigcup_{i\leq k} C_i \wedge \forall A \in MIN(CDA)[\phi \in A \rightarrow \psi \cap A \neq \emptyset]$$

Hence no such ϕ can be in Ψ without contradicting the minimality of Ψ. □

Lemma 1 $\Gamma \not\vdash_{AL2} \phi \Leftrightarrow \mathcal{T}$ *contains an infinite branch.*

Proof ($\Leftarrow$) Let $\langle\phi_0,\dots\rangle$ be an infinite branch through $\mathcal{T}$. Then any selection set $\Psi \supseteq \{\phi_i\}_{i\in\omega}$ satisfies $\forall i(C_i \cap \Psi \neq \emptyset)$. Hence $\Gamma \not\vdash_{AL2} \phi$.

($\Rightarrow$) Suppose $\mathcal{T}$ has no infinite branch. Then $\mathcal{T}$ is finite. Suppose the maximal length of any sequence in $\mathcal{T}$ is $k_0 + 1 < \omega$. Then for any selection set Ψ there is a maximal $k = k(\Psi) \leq k$ so that for some choice of $\phi_i \in C_i$ $(i \leq k)$ we have (i) $\langle\phi_0,\dots,\phi_k\rangle \in \mathcal{T}$; (ii) $\{\phi_0,\dots,\phi_k\} \subseteq \Psi$.

Fix a selection set Ψ with $k = k(\Psi)$ defined as above, with witnessing $\phi_0,\dots,\phi_k$ satisfying (i) and (ii). □

Claim $\Psi \cap C_{k+1} = \emptyset$.

Proof Suppose first that $\langle\phi_0,\dots,\phi_k\rangle$ is not maximal in $\mathcal{T}$. Let $\psi \in C_{k+1}$. If $\langle\phi_0,\dots,\phi_k,\psi\rangle \in \mathcal{T}$ then by definition of $k\{\phi_0,\dots,\phi_k,\psi\} \not\subseteq \Psi$; thus $\psi \notin \Psi$. However if $\langle\phi_0,\dots,\phi_k,\psi\rangle \notin \mathcal{T}$ then $\psi \notin \{\phi_0,\dots,\phi_k\}$ and

$$\forall A \in MIN(CDA)(\psi \in A \longrightarrow A \cap \{\phi_0,\dots,\phi_k\} \neq \emptyset).$$

So if there does exist $A \in MIN(CDA)$ with $\psi \in A$, we have that $\Psi \cap A\backslash\{\psi\} \neq \emptyset$, *i.e.* $\psi \notin \Psi$ (by the minimality of the selection set Ψ). Either way $\Psi \cap C_{k+1} = \emptyset$.

Lastly, if $\langle\phi_0,\dots,\phi_k\rangle$ is maximal in $\mathcal{T}$, the Claim follows from (ii) of the last Lemma. QED (Claim)

As Ψ was an arbitrary selection set, the claim shows that $\Gamma \vdash_{AL2} \phi$. □

To say that $\mathcal{T}$ has an infinite branch, it suffices, again by König's lemma, to say that $\mathcal{T}$ is itself infinite: $\forall n \exists \overrightarrow{\phi} \in \mathcal{T}(l(\overrightarrow{\phi}) \geq n)$. By a previous proposition, $\overrightarrow{\phi} \in \mathcal{T}$ is Σ^0_2. This is thus Π^0_3. Doing this uniformly over all ϕ then yields our desired result:

Theorem 2 *For recursive* Γ, $\{\phi \mid \Gamma \vdash_{AL2} \phi\}$ *is* Σ^0_3.

4.3 Lower bounds

We now show that this classification cannot be improved to anything simpler.

Theorem 3 *The final consequences according to either the reliability strategy, or the minimal abnormality strategy, of a recursive theory can form a complete Σ_3^0 set.*

Proof Let C be a complete Σ_3^0 set, say

$$C(n) \equiv \exists w \forall v \exists u P(w, v, u, n)$$

with P recursive.

Let Γ comprise the following axioms:

- For all w, v, u, n: $s^n_{v,w,u}, A(q^n_{w,v}, r^n_w), p_n \vee r^n_w, \neg r^n_w$
- For all w, v, u, n such that $P(w, v, u, n)$ holds: $s^n_{v,w,u} \rightarrow A(q^n_{w,v})$

Clearly Γ is a recursive set of axioms in the given propositional letters. □

Claim 1 p_n is finally derivable according to the reliability strategy if and only if $C(n)$ holds.

Proof ($\Rightarrow$) If $C(n)$ holds, then $\exists w_0 \forall v \exists u P(w_0, v, u, n)$. Thus there is a w_0 such that for all v, $A(q^n_{w_0,v})$ is derived by virtue of the last axiom group, and so determines a minimal abnormality set. But then $A(q^n_{w_0,v}, r^n_{w_0})$ is not minimal, so p_n is finally derivable.

($\Leftarrow$) If $\neg C(n)$ holds, then

$$\forall w \exists v(w) \forall u \neg P(w, v(w), u, n).$$

So for all w, $A(q^n_{w,v(w)}, r^n_w)$ determines a minimal abnormality set in the universal derivation, as $A(q^n_{w,v(w)})$ is not derived. So p_n is never derived on any unmarked line. Hence p_n is not finally derived. □

Claim 2 p_n is finally derivable according to the minimal abnormality strategy if and only if $C(n)$ holds.

Proof ($\Rightarrow$) With the notation of the previous claim, if $C(n)$ holds, $A(q^n_{w_0,v})$ is a minimal abnormality set, whereas $A(q^n_{w_0,v}, r^n_{w_0})$ is not minimal.

So any selection set Ψ must contain $q^n_{w_0,v}$ for any v. Therefore $r^n_{w_0}$ does not belong to the selection set. However in the universal derivation from Γ we have a line with p_n depending on $r^n_{w_0}$ only. Hence p_n is derived, conditional on $r^n_{w_0}$, which is not in any selection set. Hence p_n is finally derived.

($\Leftarrow$) Again if $\neg C(n)$ holds, then using the same notation, for all w, $A(q^n_{w,v(w)}, r^n_w)$ is a minimal abnormality set in the universal derivation, as $A(q^n_{w,v(w)})$ is not derived. Let Ψ be a selection set that for all w chooses r^n_w from this set. Thus on every line of the universal derivation containing p_n depending solely on r^n_w, Ψ hits the only condition. It is thus a selection set witnessing that p_n is not finally derivable. □

Batens recognizes that there is no positive test for final derivability from finite sets of premises in predicate logical versions of inconsistency-adaptive logic.[20] But with the techniques developed in this paper, precise complexity results can be obtained. For predicate logical derivability using the minimal abnormality strategy, it can be shown that a *finite* set of axioms can already have a complete Σ_3^0 final consequence set.

[20] For this reason, he has drawn up 'criteria' for final derivability. See Batens (2005b).

Let **Q** be the set of axioms of Robinson arithmetic (or any other finite set of axioms of a theory in which recursive functions can be represented). Thus, if $P(v_0, v_1, v_2, v_3)$ is recursive we have for some formula ϕ_P:

- if $P(v, w, u, n)$ then $\phi_P(\underline{v}, \underline{w}, \underline{u}, \underline{n})$ is derivable from **Q**;
- if $\neg P(v, w, u, n)$ then $\neg\phi_P(\underline{v}, \underline{w}, \underline{u}, \underline{n})$ is derivable from **Q**;

Let P be chosen so that

$$C(n) \equiv \exists w \forall w \exists u P(w, v, u, n)$$

be a complete Σ_3^0 set. Add three predicate symbols $S(v_0), R(v_0, v_1), Q(v_0, v_1, v_2)$ to the language of Q. With P recursive, choose ϕ_P as above. Now add to **Q** the following four axioms:

- $\forall n \forall w \neg R(n, w)$;
- $\forall n \forall w (S(n) \vee R(n, w))$;
- $\forall w \forall v \forall u \forall n ((Q(n, w, v) \wedge \neg Q(n, w, v)) \vee (R(n, w) \wedge \neg R(n, w)))$;
- $\phi_P \rightarrow (Q(n, w, v) \wedge \neg Q(n, w, v))$.

Then the argument of the propositional case minimal abnormality derivability can be re-run in this finite extension of **Q**. (By considering recursive partitions of the natural numbers, one can do this in a finite extension of **Q** that adds just a single monadic predicate symbol.)

In propositional inconsistency-adaptive logic, the abnormalities are all of a particularly simple form. But in other adaptive logics, the abnormalities are more complicated. The result may be that some of the corresponding final derivability relations may be even more computationally intractable than those for the reliability calculus and the minimal abnormality calculus.

5 Philosophical reflections on infinite proofs and final derivability

Batens asserts that in the definition the notion of final **ACLuN1**-consequence, there is no need to refer to infinite proofs. [21] For final **ACLuN2**-consequence, Batens recognizes that even in the propositional case, final derivability of a formula can in some instances only be witnessed by infinite proofs.[22] We have seen that this description of the situation is not quite correct. Even for the propositional fragment of **ACLuN1**, final derivability can in some instances only be established at a transfinite stage.

But this is philosophically worrisome. When Zermelo in 1932 formulated a logical system in which infinite proofs were allowed, his proposal was met with scepticism. A typical reaction was that of Church (Sinaceur, 2000, p. 33–34):

> [W]hile such systems might have considerable interest of one kind or another, they could not properly be considered *logics*, insofar as logics explicate the concept of *proof*. For what we mean by a proof is something which carries finality of conviction to any one who admits the assumptions (axioms and rules) on which the proof is based; and this requires that there be an effective (finitary, recursive) syntactical test of the validity of proposed proofs.

[21] Batens (2004, p. 479).

[22] Batens gives a propositional example on Batens (1999, p. 466).

This comment applies equally to propositional inconsistency-adaptive logic: their final proofs do not carry finality of conviction.

Batens thinks that these infinite proofs do not play a fundamental role in adaptive logic:

> However, there are some some weird cases where we have to consider infinite proofs... (Batens, 1999, p. 62, fn 33)

But we have seen that the notion of final consequence is *fundamentally* transfinite in nature, in the sense that many a (propositional) theory is such that sentences can be found such that the fact that they are a final consequence of the theory depends on the existence of certain infinitely long proofs.

Undecidability runs much deeper in adaptive logic than its proponents have realized. Taking the set of final consequences from a recursive *propositional* theory is already a complicated operation: the set of final consequences of a propositional recursive theory need not be recursively axiomatizable. For the Reliability Calculus (*pace* Batens' assertions on the matter), the final consequence set can be Σ^0_3 complete, as it is for the Minimal Abnormality Calculus. To try and get some perspective on these results, we can paraphrase, and say that were the final derivability consequences of the propositional Reliability Calculus really decidable we should have then a finitary algorithm, so a pencil and paper method, for determining *in a finite time*, given $e \in \mathbb{N}$, not just the classical Halting Problem for e, but moreover whether e codes a Turing machine that halts on all but finitely many inputs. [23]

So it is not an exaggeration to say that there exist no complete proof procedures for propositional adaptive logic, at least not if "proof" is understood in the usual (finitary) sense of the word.

Truth is complicated.[24] Derivability should be a comparatively simple relation. But even the propositional adaptive derivability relation is not simple. So we have a situation that is similar to the one in relevant logic. When it was shown that some of the most basic systems of propositional relevant logic are undecidable,[25] this was taken by some as evidence against the claim that relevant implication expresses the common sense notion of propositional implication, for it seems improbable that our common sense notion of propositional implication is so complicated. Concerning adaptive logic, one can also wonder whether such a complicated operation really is what we carry out when we propositionally reason from inconsistent theories. Nevertheless it is claimed that adaptive logic explicates how people *actually* reason from inconsistent premises:

> The dynamic proofs not only provide the [adaptive] logics with a proof theory. With their conditions and marking definitions, they explicate the actual reasoning in terms of such consequence relations. This is extremely important because they thus provide a clear and transparent conceptual analysis for forms of reasoning that were often qualified as mere tinkering or even as logically flawed. (Batens, 2004, p. 480–481)

[23] Here we are saying no more than the well known fact that $\{e | W_e$ is co-finite$\}$, is an example of Σ^0_3-complete sets (where W_e denotes the domain of the e'th Turing machine) See Rogers (1968, Chap. 14 Theorem XVI).

[24] See Burgess (1986).

[25] See Urquhart (1984).

Actually the situation is worse than in the case of relevance logic. For relevance logic intends to describe an intuitive *conditional operator* which was recognized to be more complex than the most elementary propositional logical connectives. But the clauses governing the implication symbol $\rightarrow$ of inconsistency-adaptive logic are just the ordinary classical truth-table clauses.

Adopting a learning-theoretic perspective,[26] it can be seen that our computational complexity results have implications for *convergence to the truth* in the infinite limit. Suppose we are confronted with the problem whether for a given set of natural numbers A, x is an element of A. A machine is said to converge to the truth in the infinite limit for this problem if when an input x is given, it prints out a sequence of "yesses" and "noes" in response and from some point on prints only the correct answer, though there may be, in general, no effective way of determining when it has begun to print only the correct answer. Such an algorithm exists if and only if A is a Δ^0_2 set. A machine exists which eventually stabilizes on the correct answer for a set A for all x which in fact belong to A only if A is Σ^0_2; a machine exists which eventually stabilizes on the correct answer for all x which do not belong to A only if A is Π^0_2.[27] The results of this paper entail that there exists no algorithm which converges to the truth in the infinite limit for final derivability from a recursive set of propositional premises. Moreover, there can be no algorithm for converging to the truth in the limit for those sentences which in fact are finally derivable from a given recursive set of premises nor can there be such for those sentences which are not finally derivable from the given premise sets.

The adaptive logicians say that we must be willing to work with and in inconsistent theories, at least provisionally. And in doing so, we try to minimize and localize inconsistency. At first sight, minimizing the effects of inconsistency looks like a semantical operation. For it seems to amount to an instruction to concentrate on models which make true a "minimal" set of inconsistent statements. Adaptive Logicians claim as an advantage of their approach that they have *proof systems* that produce finite provisional proofs that *approximate* the final adaptive consequences of an inconsistent theory:

> What if no criterion enables one to conclude from a proof whether some formula is or is not finally derivable from the premise set? [...] Roughly, the answers go as follows. First, there is a characteristic semantics for derivability at a stage. Next, it can be shown that, as a dynamic proof proceeds, the insight in the premises provided by the proof never decrease and may increase. In other words, derivability at a stage provides an estimate for final derivability, and, as the proof proceeds, this estimate may become better, and never becomes worse. (Batens, 2001, p. 63) [28]

The claim that as an adaptive proof proceeds it gives an ever better estimate of final derivability, is taken to be supported by the fact that the class of models of a given adaptive proof is always a superset of the class of models of an extension of that proof.[29] But the results of this paper indicate that not too much should be made of this fact. This is seen as follows: the collection of adaptive *provisional* consequences (as the "estimate" set of formulae derivable at a particular stage) always forms a

[26] See Kelly (1996).

[27] See Putnam (1965).

[28] See also Batens (2004, p. 480).

[29] See Batens (2004, p. 480) and Batens (1995, p. 301).

recursively enumerable collection of sentences. But the *final* consequences may, as we have seen, form a set which is much more complex than any recursively enumerable set. In such cases as have been considered in this paper, provisional consequences (i.e. those derived at a a stage) form a very poor approximation of the final consequence set. As we have seen, provisional consequence sets in general do become worse as well as better: there are proofs of finally derivable propositions that *must* infinitely often change their mind about the derivability of those propositions. Thus "derivability at a stage" *provably* does not form a good, monotone, method of estimation.

Acknowledgements Thanks to Liza Verhoeven, Kristof De Clercq, Guido Vanackere, Diderik Batens, two anonymous referees and to the members of the Luxemburger Kreis for valuable comments on earlier versions of this paper. The research for this paper was supported by a British Academy International Collaborative grant; the second author wishes to further acknowledge support from an EPSRC Fellowship.

References

Batens, D. (1995). Blocks: The clue to dynamic aspects of logic. *Logique et Analyse, 150–152*, 285–328.

Batens, D. (1999). Inconsistency-adaptive logics. In: (Orlowska, 1999, p. 445–472).

Batens, D., Mortensen, Ch., Priest, G., & Van Bendegem J.-P. (Eds.). (2000). *Frontiers of paraconsistent logic*. Baldock, UK: Research Studies Press.

Batens, D. (2001). A general characterization of adaptive logics. *Logique et Analyse, 173–175*, 45–68.

Batens, D. (2004). The need for adaptive logics in epistemology. In: (Rahman, 2004, p. 459–485).

Batens, D. (2005). The theory of explanation generalized to include the inconsistent case. *Synthese, 143*, 63–88.

Batens, D. (2005). A procedural criterion for final derivability in inconsistency-adaptive logics. *Journal of Applied Logic, 3*, 221–250.

Burgess, J. (1986). The truth is never simple. *Journal of Symbolic Logic, 51*, 663–681.

Gärdenfors, P. (1988). *Knowledge in flux. Modeling the dynamics of epistemic states.* Cambridge, MA: MIT Press.

Kelly, K. (1996). *The logic of reliable inquiry.* Oxford: Oxford University Press.

Orlowska, E. (1999). *Logic at work. Essays dedicated to the memory of Helena Rasiowa.* Heidelberg, NY: Springer.

Putnam, H. (1965). Trial and error predicates and the solution to a problem of Mostowski. *Journal of Symbolic logic, 30*, 49–57.

Rahman, S. (Ed.) (2004). *Logic, epistemology, and the unity of science.* Dordrecht: Kluwer.

Rogers, H. (1968). *Theory of recursive functions and effective computability.* NY: McGraw-Hill.

Sinaceur, H. (2000). Address at the Princeton University Bicentennial Conference on Problems in Mathematics (December 17–19, 1946), by Alfred Tarski. *Bulletin of Symbolic Logic, 6*, 1–44.

Urquhart, A. (1984). The undecidability of entailment and relevant implication. *Journal of Symbolic Logic, 49*, 1059–1073.

Weber, S. (2000). Mathematical analysis of priest's nonmonotonic version of Kleene's SK3. In: (Batens et al., 2000, p. 299–314).