BRYSON BROWN and GRAHAM PRIEST

CHUNK AND PERMEATE, A PARACONSISTENT INFERENCE STRATEGY. PART I: THE INFINITESIMAL CALCULUS

Received 29 January 2003; received in revised version 20 January 2004

ABSTRACT. In this paper we introduce a paraconsistent reasoning strategy, Chunk and Permeate. In this, information is broken up into chunks, and a limited amount of information is allowed to flow between chunks. We start by giving an abstract characterisation of the strategy. It is then applied to model the reasoning employed in the original infinitesimal calculus. The paper next establishes some results concerning the legitimacy of reasoning of this kind – specifically concerning the preservation of the consistency of each chunk – and concludes with some other possible applications and technical questions.

KEY WORDS: chunking, infinitesimal calculus, paraconsistent logic

1. INTRODUCTION

The infinitesimal calculus of Leibniz and Newton is well known to have operated on an inconsistent basis. In particular, at different points in the calculation of a derivative infinitesimals had to be assumed to be both zero and non-zero.[1] How was the trick turned? The fact that arbitrary conclusions were not drawn, even though reasoning was carried out on the basis of inconsistencies, means that the inference procedure involved must have been paraconsistent (where contradictions do not entail everything). But what was it? It does not appear to have been any standard paraconsistent logic.[2] In this paper we will suggest an answer to this question. The computation was broken into chunks, and a limited amount of information was allowed to flow between the chunks. That is, certain information permeated chunk boundaries. We therefore call this procedure 'Chunk and Permeate'. In the next section we will give an abstract characterisation of the strategy. Following that, we will show how the strategy can be applied in the particular case of the infinitesimal calculus. In the next section, we will make a few comments concerning the conditions under which the strategy can operate coherently. In particular, we will show the legitimacy of the reasoning in the case of the calculus. We close with some comments on various issues that the procedure raises.

The strategy of Chunk and Permeate has a much wider application than the calculus. In fact, it seems to be employed not uncommonly in

Journal of Philosophical Logic **33:** 379–388, 2004.

science where inconsistency is in play, and especially where parameters whose computations depend on inconsistent underlying assumptions are then combined in some further calculation. A striking and famous example of this is the Bohr theory of the atom. That particular example will be analysed in a second part of this paper.

2. Chunk and Permeate

A common procedure for handling inconsistent information is to break it up into consistent fragments, and then to operate within these. This is a strategy that is at the root of non-adjunctive paraconsistent logics, of which there are numerous kinds. The procedure we will describe here has that feature. It also has a feature that standard non-adjunctive logics do not possesses: a certain kind of interplay between the chunks. Specifically, information is allowed to flow between the chunks. Of course, if *all* the information in one chunk were allowed to flow to another, this would destroy the chunking procedure; so there has to be a limit. Hence, there must be a mechanism for allowing partial flow. We might, then, think of the chunks as separated by membranes which are semi-permeable: permeable to sentences of some kinds, but not others.

So much for the idea. Let us give it a precise rendering. Let L be some classical language (of first or higher order). Let $\vdash$ be the appropriate classical consequence relation. If Σ is a set of sentences in the language, let $\Sigma^{\vdash}$ be the closure of Σ under $\vdash$. A *covering* of Σ is a set $\{\Sigma_i; i \in I\}$, such that $\Sigma = \bigcup_{i \in I} \Sigma_i$, and for all $i \in I$, Σ_i is classically consistent. If $C = \{\Sigma_i; i \in I\}$ is a covering on Σ, call ρ a *permeability relation* on C if ρ is a map from $I \times I$ to subsets of the formulas of L (e.g., those of some particular syntactic form). If $i_0 \in I$, we will call any structure $\langle C, \rho, i_0 \rangle$ a *C&P structure* on Σ.

If $\mathfrak{P}$ is a C&P structure on Σ, define the C&P consequences of Σ with respect to $\mathfrak{P}$, $\Vdash_{\mathfrak{P}}$, as follows. First, for each $i \in I$, we define a set of sentences, Σ_i^n, by recursion on n:

$$\Sigma_i^0 = \Sigma_i^{\vdash}$$

$$\Sigma_i^{n+1} = \left(\Sigma_i^n \cup \bigcup_{j \in I} (\Sigma_j^n \cap \rho(j, i))\right)^{\vdash}$$

Thus, Σ_i^{n+1} comprises what can be inferred from Σ_i^n *together with* whatever flows into chunk i from the other chunks at level n. Next we collect up the results of all finite stages:

$$\Sigma_i^{\omega} = \bigcup_{n<\omega} \Sigma_i^n$$

Finally, $\Sigma \Vdash_{\mathfrak{P}} A$ iff $A \in \Sigma_{i_0}^{\omega}$. The C&P consequences of Σ are thus the sentences that can be inferred in the designated chunk, i_0, when all information of the appropriate form has been allowed to flow along the permeability relations.[3]

A particularly simple case of a C&P structure, which we will meet in what follows, is where there are just two chunks, the *source* chunk, Σ_S, and the *target* chunk, Σ_T; where the flow of information is from the former to the latter (only); and where the target chunk is the output chunk. Note that in this case, clearly, $\Sigma_T^1 = \Sigma_T^{\omega}$. We will call such a C&P structure a *binary structure*.

We end this section by noting the following. The permeability relation allows information to be aggregated, but only in a controlled way. In standard non-adjunctive logics, controlled aggregation is obtained by appealing to variation in the partitions. For example, in the approach of Scotch and Jennings[4] the level of a set, $l(\Sigma)$, is defined as the least cardinal, n, such that there is a covering of Σ of size n (or ∞ if there is no such n). Consequence, $[\!\vdash$, is then defined as follows: $\Sigma\ [\!\vdash \alpha$ iff $l(\Sigma) = \infty$ or $(l(\Sigma) = n$ and for every covering of size n, $\{\Sigma_i; i \in I\}$, there is a $j \in I$ such that $\Sigma_j \vdash \alpha)$. The permeability mechanism is much more powerful than anything that can be produced in this way. Indeed, the effect of the Scotch and Jennings mechanism can be obtained simply with a family of C&P structures each of which has the empty permeability relation.[5]

3. The Infinitesimal Calculus

We will now show how the C&P strategy can be applied in the particular example of the original infinitesimal calculus. First, a word of caution. This is not an historical paper. We do not claim that Leibniz, Newton, and the other mathematicians who worked with the infinitesimal calculus explicitly presented their work in this way. There is therefore no point examining texts to determine whether or not they did. The strategy we give is, rather, that of a rational reconstruction. What we do claim is that this strategy captures the practice that they engaged in – however, in fact, these mathematicians conceived of it. Doubtless, also, they told metaphysical stories about the nature of infinitesimals (different for Leibniz and for Newton), which attempted to justify, or at least motivate, their practice. Again, this is not our concern here.[6]

Let L be the language of the second-order theory of the real numbers. We suppose that the language includes the functional abstraction operator, λ. We single out for special attention one monadic function symbol, δ. Intuitively, $\delta(x)$ is an infinitesimal part of x. We will normally write $\delta(x)$ as δx. We will also suppose that there is a functional, D. That is, if f is a function, so is Df. Intuitively, Df is the derivative of f.

We will call the C&P structure in question *LN* (for Leibniz–Newton). It is a binary structure, $\langle\{\Sigma_S, \Sigma_T\}, \rho, T\rangle$. The source chunk, Σ_S, contains the second-order theory of the reals (or at least, an axiom system for an appropriately strong part of this), together with the usual axioms for λ. There are two further axioms:

S1 $Df = \lambda x((f(x+\delta x) - fx)/\delta x)$

S2 $\forall x\, \delta x \neq 0$

The first of these is, in effect, the definition of the derivative. The second states the relevant property of infinitesimals in this chunk. The target chunk, Σ_T, is the same as the first, except that it contains neither of these axioms, but contains instead:

T1 $\forall x\, \delta x = 0$

This states the relevant properties of infinitesimals in this chunk. Clearly, $\Sigma = \Sigma_1 \cup \Sigma_2$ is inconsistent. The permeability function, ρ, is such that $\rho(S, T)$ is the set of equations of the form $Df = g$, where:

- neither f nor g contains D (or I when we go on to add it in a moment),
- f is a λ-term containing no occurrences of δ,
- g is of the form $\lambda x(h + p)$, where h contains no occurrences of δ, and p is a polynomial of powers > 0 of δx.

(These constraints are not necessary to make the computation work; they are necessary for the legitimacy proof of Section 4.) Since this is a binary structure, $\rho(T, S) = \phi$. Thus, information about derivatives is allowed to flow from Σ_S to Σ_T, but nothing is allowed to flow back.

To illustrate how the setup works, let us show how it delivers the derivative of the function $\lambda x(x^2)$. First, working within Σ_S, we have:

$$\begin{aligned} D\lambda x(x^2) &= \lambda x((\lambda x(x^2)(x+\delta x) - \lambda x(x^2)x)/\delta x) \\ &= \lambda x((x+\delta x)^2 - x^2)/\delta x) \\ &= \lambda x((2x\delta x + (\delta x)^2)/\delta x) \\ &= \lambda x(2x + \delta x) \end{aligned}$$

The last line is *kosher* since $\delta x \neq 0$. The equation $D\lambda x(x^2) = \lambda x(2x+\delta x)$ is in Σ_S. Allow this to permeate into Σ_T. Now, working within Σ_T, since $\delta x = 0$, we have $D\lambda x(x^2) = \lambda x(2x)$. Hence $\Sigma \Vdash_{\mathfrak{P}} D\lambda x(x^2) = \lambda x(2x)$, as required.

Note that it is important that this equation is not allowed to flow back into Σ_S, or within Σ_S we would have $D\lambda x(x^2) = \lambda x(2x+\delta x) = \lambda x(2x)$. It would follow that, say, $\delta 0 = 0$. Hence, Σ_1 would be inconsistent. We could therefore infer, say, $D\lambda x(x^2) = \lambda x(7x)$. This would then permeate back into Σ_T, making it inconsistent too.[7]

Integration can also be handled in the same C&P structure. We now suppose that the language is augmented by a new three-place functional $I(r, s, f)$, where r and s are numbers, and f is a function. We will write this in the more familiar fashion $\int_r^s f$. Additionally, in the source chunk, I is satisfies the equation:

$$\text{S3} \quad \int_r^s f = \sum_{i=1}^{(s-r)/\delta r} \delta r.f(r + i.\delta r)$$

Intuitively, δr is an infinitesimal which divides $s-r$, and the area under the curve between $x = r$ and $x = s$ is broken up into rectangles of width δr, so that there are $(s-r)/\delta r$ such rectangles. (In particular, then, $\delta r \neq 0$). The ith such rectangle has area $\delta r.f(r + i.\delta r)$).

Equations of the form $\int_r^s f = g$ are now also allowed to permeate from source to target, where f and g are subject to the same restrictions as before, except that g is now a numerical term of the form $h + p$, where h contains no occurrences of δ, and p is a polynomial of powers > 0 of δr. Finally, in addition:

- r and s are numerical terms containing no occurrences of δ, D or I.

Let us illustrate how integration may be performed by computing the integral of the function $\lambda x(x)$ between $x = r$ and $x = s$. Let $n = (s-r)/\delta r$. We then have in Σ_S:

$$\begin{aligned}
\int_r^s \lambda x(x) &= \sum_{i=1}^{n} \delta r.\lambda x(x)(r + i.\delta r) \\
&= \sum_{i=1}^{n} \delta r(r + i.\delta r) \\
&= \sum_{i=1}^{n} (r.\delta r + i(\delta r)^2)
\end{aligned}$$

$$
\begin{aligned}
&= nr.\delta r + \sum_{i=1}^{n} i(\delta r)^2 \\
&= nr.\delta r + (\delta r)^2 \sum_{i=1}^{n} i \\
&= nr.\delta r + (\delta r)^2 n(1+n)/2 \quad \text{summing the series} \\
&= nr.\delta r + n(\delta r)^2/2 + n^2(\delta r)^2/2 \\
&= r(s-r) + (s-r).\delta r/2 + (s-r)^2/2 \\
&\quad \text{since } n.\delta r = (s-r) \\
&= s^2/2 - r^2/2 + \frac{(s-r)}{2}.\delta r
\end{aligned}
$$

The final equation is then permeated into the target chunk, where $\delta r = 0$. In that chunk we then have $\int_r^s \lambda x(x) = s^2/2 - r^2/2$.

The computations we have just looked at are, of course, only examples of the way that reasoning in the calculus is handled in the C&P structure. They suffice to illustrate how it works, however.[8]

4. The Preservation of Consistency Under Permeability

In this section, we make a few observations and establish a few results about the legitimacy of C&P reasoning in binary structures. Let us suppose that Σ_T is a consistent chunk, and that the information α is allowed to flow into it. If $\Sigma_T \nvdash \neg\alpha$, then $\Sigma_T \cup \{\alpha\}$ is consistent. In particular, if Σ_T is incomplete with respect to α then $\Sigma_T \cup \{\alpha\}$ is consistent. Once more than one piece of information is allowed to flow into Σ_T, however, independence is no longer sufficient. For example, α and β may both be independent of Σ_T. Yet it may be the case that $\Sigma_T \vdash \alpha \leftrightarrow \neg\beta$. Hence, $\Sigma_T \cup \{\alpha, \beta\}$ is inconsistent.

Suppose, however, that the only information that is allowed to flow is in the form of identities (equations) of a certain kind. This is, in fact, a very common way for C&P to work. The values of certain quantities are computed in one chunk, and then passed on to the other. (The binary structure *LN* works in this way.) Suppose, specifically, that the permeable sentences are of the form $f(a_1 \ldots a_n) = c$, where $a_1, \ldots, a_n, c$ contain no occurrences of f, and where Σ_T is an axiomatic theory with no axioms mentioning f. If $\mathfrak{M}$ is any interpretation and t is an term, let us write the denotation of t in $\mathfrak{M}$ as $t^{\mathfrak{M}}$. Then it is not difficult to see the following.

LEMMA. *Let $\mathfrak{T}$ be a model of Σ_T and suppose that for all equations $f(a_1 \dots a_n) = c$, $f(b_1 \dots b_n) = d$ in $\Sigma_S \cap \rho(S, T)$:*

(∗) *if $a_i^{\mathfrak{T}} = b_i^{\mathfrak{T}}$, for $1 \leq i \leq n$, then $c^{\mathfrak{T}} = d^{\mathfrak{T}}$*

Then Σ_T augmented by the permeated information has a model (and so is consistent).

Proof. Since there are no axioms governing f, in Σ_T we can alter the interpretation of f in $\mathfrak{T}$ in any way we like, and still have a model of Σ_T. To obtain a model of Σ_T together with all the permeated sentences, $f(a_1 \dots a_n) = c$, we simply take the denotation of f to be the function that assigns $c^{\mathfrak{T}}$ to each $\langle a_1^{\mathfrak{T}}, \dots, a_n^{\mathfrak{T}} \rangle$. We can do this because the condition (∗) holds. □

One way in which we may be able to prove a condition of the form (∗) is as follows. Suppose that Σ_S is consistent, and so has a model, $\mathfrak{S}$. Let us suppose that the domains of $\mathfrak{T}$ and $\mathfrak{S}$ are the same, and that the interpretations of the vocabulary occurring in all the terms in the permeated equations are the same in $\mathfrak{T}$ and $\mathfrak{S}$. This situation will arise if, for example, the terms are purely numerical, and the numbers work in exactly the same way in $\mathfrak{T}$, and $\mathfrak{S}$. Now, if $a_i^{\mathfrak{T}} = b_i^{\mathfrak{T}}$ for all $1 \leq i \leq n$, we have $a_i^{\mathfrak{S}} = b_i^{\mathfrak{S}}$. But in $\mathfrak{S}$, $f(a_1 \dots a_n) = c$, and $f(b_1 \dots b_n) = d$. It follows that $c^{\mathfrak{S}} = d^{\mathfrak{S}}$, so that $c^{\mathfrak{T}} = d^{\mathfrak{T}}$, as required.

This argument will not work with the binary structure *LN*, since the terms c and d in question may contain infinitesimals, and the interpretation of these must be different in any source and target models. However, a more complex argument works.

LEMMA. *In the binary structure LN, we can find an interpretation $\mathfrak{T}$ satisfying condition* (∗).

Proof. Let $\mathfrak{S}$ be a non-standard model of analysis, of the kind familiar from non-standard analysis. In $\mathfrak{S}$ let δx be a number infinitesimally close to x (chosen, e.g., by some choice function). S1 and S3 can simply be taken as a definition of D and I in $\mathfrak{S}$, and the model clearly satisfies S2. Let $\mathfrak{T}$ be the standard model of analysis, and in this model, let δ be the zero function. This clearly makes $\mathfrak{T}$ a model of Σ_T. Now suppose that $a_i^{\mathfrak{T}} = b_i^{\mathfrak{T}}$, for all relevant i. (If f is D, $i = 1$; if f is I, $i = 1, 2, 3$.) The as and bs are numerical terms or λ-terms, containing no occurrences of δ, and so expressible in the standard language of real numbers. Hence, since $\mathfrak{S}$ and $\mathfrak{T}$ are elementarily equivalent as far as this language goes, $a_i^{\mathfrak{S}} = b_i^{\mathfrak{S}}$, for all the relevant i. Hence $c^{\mathfrak{S}} = d^{\mathfrak{S}}$. If c and d are numerical terms, they are of the form $h_c + p_c$ and $h_d + p_d$ where the hs contain no occurrences of

δ, and the ps are polynomials of powers > 0 of δ-terms. h_c and h_d denote reals (since they do not contain δ) and the ps denotes infinitesimals. Thus, since $(h_c + p_c)^{\mathfrak{S}} = (h_d + p_d)^{\mathfrak{S}}$, $h_c^{\mathfrak{S}} = h_d^{\mathfrak{S}}$, and so $h_c^{\mathfrak{T}} = h_d^{\mathfrak{T}}$. Since each p denotes 0 in $\mathfrak{T}$, $(h_c + p_c)^{\mathfrak{T}} = (h_d + p_d)^{\mathfrak{T}}$, i.e., $c^{\mathfrak{T}} = d^{\mathfrak{T}}$, as required. Similarly, if c and d are λ-terms, they are of the form $\lambda x(h_c + p_c)$ and $\lambda x(h_d + p_d)$ where the hs contain no occurrences of δ, and the ps are polynomials of powers > 0 of δx. If x denotes any real, h_c and h_d denote reals (since they do not contain δ) and the ps denotes infinitesimals. Thus, since $(h_c + p_c)^{\mathfrak{S}} = (h_d + p_d)^{\mathfrak{S}}$, $h_c^{\mathfrak{S}} = h_d^{\mathfrak{S}}$, and so $h_c^{\mathfrak{T}} = h_d^{\mathfrak{T}}$. Since each p denotes 0 in $\mathfrak{T}$, $\lambda x(h_c + p_c) = \lambda x(h_d + p_d)$ in $\mathfrak{T}$, i.e., $c^{\mathfrak{T}} = d^{\mathfrak{T}}$, as required. □

COROLLARY. *In the binary structure LN, the flow of information into* Σ_T *preserves consistency.*

5. Other Applications and Generalisations

The foregoing serves to illustrate the basic ideas involved in C&P inference. As we have already observed, this kind of inference seems to have important uses in empirical science. One example of this, the Bohr theory of the atom, will be the topic of Part 2 of this paper. Other examples abound. It is not uncommon to reason about a situation or a piece of equipment on the basis of various parameters, each of which is determined by applying theories with quite different underlying assumptions.[9] It is clear, in principle, how this reasoning works on a C&P basis. Each parameter is computed in a different chunk; the parameters are then allowed to permeate into further chunks in which the reasoning continues. Renormalisation procedures of the kind currently employed in relativistic quantum theory may well also be conceptualisable as C&P. The method may have applications, too, in computational information processing where inconsistent information must be handled. Indeed, the modularisation of computation lends itself to this kind of structure. The passing of parameters from one routine to another is standard fare in this kind of architecture.

It should be noted that the method may be generalised beyond what has been presented here. There is no reason why each chunk must employ classical logic: the chunks may have any logic. Indeed, there is no reason as to why every chunk must have the same logic. It is quite natural to suppose that different logics may be applicable in different contexts (= chunks).

Finally, note that the construction gives rise to many interesting technical issues. For example: what can be said in general about the constraints on information flow between the chunks which are sufficient (or even necessary and sufficient) for the maintenance of consistency? What other constraints on information flow are natural? (If a chunk has an underlying paraconsistent logic, consistency is not, of itself, important; but non-triviality certainly is.) What can be said about the computational complexity of C&P reasoning? (If we know the computational complexity of each chunk and of the classes of sentences allowed to permeate between chunks, what can be said about the complexity of $\Vdash$?)

This is not the place to go into any of these issues here. What is clear is that C&P reasoning seems ripe for a wide range of applications and corresponding theoretical investigations. Though the basic idea is simple, then, the method would appear to be one of significant power and utility.[10]

NOTES

[1] See Priest and Routley (1989), pp. 374–6.

[2] Simply allowing infinitesimals to have inconsistent properties leads to completely unacceptable results. Thus, suppose that $\delta x = 0$ and $\delta x \neq 0$. Since $0 = 2.0$, $\delta x = 2\delta x$. And since $\delta x \neq 0$, we can divide to obtain $1 = 2$.

[3] More generally, the output chunk may be context-dependent. Thus, for example, there might be two output chunks, i_0 and i_1, and we might look to i_0 if we want the conclusions of the system concerning micro-objects (or of a certain syntactic form), and i_1 for the conclusions of the system concerning macro-objects (or of a different syntactic form). We ignore these complications here.

[4] See, e.g., P. Scotch and R. Jennings (1980, 1989).

[5] Specifically: Σ $[\vdash \alpha$ iff (if $l(\Sigma) = n$ and $C = \{\Sigma_i; i \in I\}$ is a covering of Σ of size n, then there is a $j \in I$ such that $\Sigma \Vdash_{\langle C,\phi,j\rangle} \alpha$).

Proof. If $l(\Sigma) = \infty$ then both sides are true (the right-hand side being vacuously so). So suppose that $l(\Sigma) = n$.

Suppose that Σ $[\vdash \alpha$. Let $C = \{\Sigma_i; i \in I\}$ be a covering of size n. Then there is a $j \in I$ such that $\Sigma_j \vdash \alpha$. It follows that $\Sigma \Vdash_{\langle C,\phi,j\rangle} \alpha$. Conversely, suppose that $C = \{\Sigma_i; i \in I\}$ is a covering of Σ of size n; then there is a $j \in I$ such that $\Sigma \Vdash_{\langle C,\phi,j\rangle} \alpha$, i.e., $\Sigma_j \vdash \alpha$. So Σ $[\vdash \alpha$.

[6] For a history of the matter, see Boyer (1949), Ch. 5 and Cajori (1991), pp. 191–219.

[7] Computing second and higher-order derivatives is straightforward. For example, let us write DDf as $D^2 f$. Then we show that $D^2\lambda x(x^2) = \lambda x(2)$ as follows. We have shown that $D\lambda x(x^2) = \lambda x(2x)$ in Σ_T; $D\lambda x(2x) = \lambda x(2)$ is proved a similar way. Now the first of these entails $D^2\lambda x(x^2) = D\lambda x(2x)$; and this together with the second entails the result.

[8] The whole of a computation in a binary system can be captured in a single modal logic with a normal modal operator, $\Box$. The axioms are all of the sentences in Σ_T, $\Box\sigma$, for every $\sigma \in \Sigma_S$, and $\Box\sigma \to \sigma$ for every $\sigma \in \rho(S,T)$. (Note that the system is *weaker* than K_ρ.) If for some $\sigma_1, \ldots, \sigma_n \in \Sigma_S$, $\sigma_1, \ldots, \sigma_n \vdash \alpha$, $\Box\sigma_1 \wedge \cdots \wedge \Box\sigma_n \vdash \Box\alpha$. Hence, if $\alpha \in \rho(S,T)$, we may infer α, and complete the computation invoking the axioms in Σ_T. One might think of the worlds picked out by the operator $\Box$ as possible worlds realising the "fictional truths" of Σ_S. In such worlds, for example, the object δx, which is, in fact,

zero, behaves in a non-zero fashion. Information about the fictional world is taken to apply to the actual world provided that it does not involve the fiction.

[9] For an example of the former, see Feyerabend (1978), pp. 62–63; for an example of the latter, see Hacking (1983), pp 268–9.

[10] A version of this paper was given at a joint Adelaide/Melbourne logic mini-conference, held in Melbourne, May 2002. It was also given at a conference on the foundations of mathematics held at the University of Nancy in October 2002. We are grateful to those present on these occasions for helpful comments and thoughts, and especially to Johan van Bentham, Bryn Humberstone and Lloyd Humberstone.

REFERENCES

Boyer, C. B. (1959) *The History of the Calculus and its Conceptual Development*, Dover Publications, New York.
Cajori, F. (1991) *A History of Mathematics*, 5th edn, Chelsea Publishing Co., New York.
Feyerabend, P. (1975) *Against Method*, Verso, London.
Hacking, I. (1983) *Representing and Intervening*, Cambridge University Press, Cambridge.
Priest, G. and Routley, R. (1989) Applications of paraconsistent logic, Ch. 13 of Priest, Routley and Norman (1989).
Priest, G., Routley, R. and Norman, J. (eds.) (1989) *Paraconsistent Logic: Essays on the Inconsistent*, Philosophia Verlag, München.
Scotch, P. K. and Jennings, R. (1980) Inference and necessity, *J. Philos. Logic* **9**, 327–40.
Scotch, P. K. and Jennings, R. (1989) On detonating, Ch. 11 of Priest, Routley and Norman (1989).

BRYSON BROWN
Department of Philosophy,
University of Lethbridge

GRAHAM PRIEST
Department of Philosophy,
University of Melbourne,
and Department of Logic and Metaphysics,
University of St Andrews