
Aerial Monofluoroacetate in New Zealand's Forests

An appraisal of the scientific evidence.

New Zealand

31 January 2007

ERMA New Zealand
P O Box 131
Wellington

Re: Submission on Application No. HRE05002, Reassessment of 1080

Dear Sir or Madam:

We have requested a hearing on our submission. If ERMA wishes independent scientific review of the assertions in our submission, we would welcome this. With regard to issues of good experimental design and statistical inference, we suggest contacting international authorities at say the Johns Hopkins School of Public Health, or other internationally respected authorities. With respect to questions regarding the principles of ecology and the management of ecosystems, we would encourage you to contact the departments of ecology at Harvard, Stanford, or Cambridge Universities or at the University of California, Davis. These are the finest academic institutions in the world. It is our view that the safety of New Zealand's forests and our national reputation as an environmentally conscientious country deserves no less.

If ERMA provides the applicant with the opportunity to respond to and refute the material in our submission, we wish to be able to respond to their written comments in writing and to be present at any hearing that may pertain.

An electronic version of our paper with hypertext links to the full text of the cited literature will follow under another cover. This will make it possible for ERMA reviewers to have instant access to the original papers that we have cited.

Finally, we are requesting the raw data from the most important of the DoC/AHB studies so that they can be reanalyzed by us. Thus, far we have not obtained any of these data sets, but they may be available prior to the time of our oral presentation. If so we would like to include those results in our presentation.

We appreciate your careful consideration of our submission and look forward to receiving ERMA's decision regarding this application.

Respectfully,



Enclosure

Table of Contents

Table of Contents	1
Introduction	3
Issues not addressed	3
Limitations of this paper	4
Science, politics and the nature of this document	4
Monofluoroacetate facts	5
Experimental design and statistical inference	7
Controls	7
Blinded observation	9
P-values	9
Confidence intervals and statistical power	10
The role of random sampling	11
More than one study by more than one group of independent investigators	11
The absolute need for researcher independence: the human factor	11
The absolute need for financial researcher independence	12
Does it matter?	12
Summary and implications	13
The question of net ecosystem benefit (or harm)	14
Burden of proof	14
Unprecedented practice	15
New Zealand is unique	16
Ecosystem level studies showing net effect	17
The problem with Control Level 0 "research": who to believe?	17
An example study design	19
Summary	20
The species level evidence	20
Bird populations and mortality	20
The Miller et al study.	23
The Pierce et al study: decreased numbers of tomtits	24
Powlesland et al: The egregious case of robins and tomtits.	25
Other research cited in the AHB/DoC submission regarding native birds	29
Summary and conclusion	29
Invertebrates	31
Meads	31
DoC's reply to the Meads paper: Spurr 1994	32
The politics of 1080 and how the Meads' paper came not to be officially published	35
Conclusion	36
Other invertebrate studies.	37
Mammals: bats	40
Sub-lethal effects on birds and the ecosystem	40
Evidence that possums are importantly deleterious to native forests	42

Table of Contents

Four species studied: aerial 1080 associated with worsening of the forest canopy	42
The Payton et al (1988)	44
Bellingham et al	44
Mistletoe	45
Nugent et al (2002)	45
Sweetapple et al (2004)	48
Fuchsia study (2005)	49
The Green Report (2003)	50
Other studies	51
"Pests" other than possums	51
Conclusion	52
The Department of Conservation: guardian of the environment or typical bureaucracy?	52
Aerial 1080 and the control of bovine tuberculosis (TB)	53
What the evidence should be	54
Possums are a vector of bovine TB	54
Pasture margins are the place to work	54
How big a pasture margin buffer?	55
Summary	56
Alternatives to Aerial 1080	56
Ground-based trapping	56
Trap only the margins	57
Other approaches	57
By its own numbers DoC's programme has failed	58
Multi-pronged strategy	59
Conclusion:	60
Does the structure of research funding guarantee biased results?	60
Does the DoC/AHB submission misrepresent the truth about aerial 1080?	61
About the authors	71
Quinn Whiting-OKeefe, MD, Ma	71
Patricia Whiting-OKeefe, PhD	71
Summary and Conclusion: "First do no harm"	72
Personal comment	73
References	75

Introduction

Annually, the New Zealand Department of Conservation (DoC) and the New Zealand Animal Health Board (AHB) routinely drop from the air food laced with enough of an "extremely hazardous" (16) poison (sodium monofluoroacetate, also called compound 1080) into New Zealand's unique forest ecosystems to kill every person in New Zealand 8 times over. DoC rationalizes its policy by saying that it is necessary to control feral "pests". They claim that it benefits native species and forests and does not do significant harm. AHB believes that it is necessary to control bovine tuberculosis (TB). DoC and AHB have jointly applied (1) to ERMA to continue and extend authorization for this practice. The purpose of this document is to examine the scientific evidence supporting the contention that aerial monofluoroacetate (aerial 1080) is benign and beneficial to our forest ecosystems and the contention that it is essential to the control of bovine TB. Since these objectives and the evidence needed to support them are quite different, they will be dealt with separately.

Issues not addressed

One-shot use of aerial 1080 on true islands. Aerial 1080 has been used on true islands to eradicate feral mammals. The important feature of this is that it usually requires only one poisoning, or at most two.

Other uses of 1080. Monofluoroacetate itself is not the issue that we have investigated. It is rather the aerial application of food laced with 1080 into our forest ecosystems that is the subject of this paper. We suspect that any other broad spectrum poison would have similar effects, e.g., cyanide. The use of such poisons in traps that limit access to all but targeted species may be necessary and even desirable, but in any case is not the subject of this investigation.

Risks to Humans. We have not attempted to assess scientifically the risk of aerial 1080 to humans. Compound 1080 is a highly toxic chemical that will certainly kill humans if they are exposed to even small amounts, but this is true of many substances. It is fairly clear from the literature the aerial 1080 in the concentrations in which it is usually applied does not constitute a major risk to humans from water contamination, providing it is used and applied as it is supposed to be. So the risk comes down to that from accidents, errors and malice*.

Over the last two or three decades, there have been numerous reports of accidents and near accidents, of accidental animal poisoning and the like. As the use of 1080 becomes more widespread, its handling would be expected to become increasingly "routine", which means it is probably just a matter of time until something really serious happens. The other grave possibility is that of a child walking into a recently poisoned forest and eating some bait. Because DoC routinely drops aerial 1080 into forests that are near human habitation and that are frequented by humans, this risk would appear to be substantial, and indeed at least one child was almost killed (2)†.

Risks to domestic animals. We do not assess the risk to domestic animals in this document. It may be substantial, but we have not looked into the issue.

* As a weapon, 1080 would certainly qualify as one of mass destruction. A few kilograms put into the water supply in the right place could result in the death of hundreds or thousands of individuals.

† Personally, as a physician, I would not wish to be the one who signed off on this practice.

Limitations of this paper

We have not attempted to be exhaustive in our coverage of the research literature. We have selected papers which seemed to have the best methodology, that were frequently referenced by other authors or the DoC/AHB submission, or that contain important results. In this we are confident that we have not missed major studies on the central issues, but we have not reviewed every scientific paper in New Zealand that has anything to do with aerial 1080, nor do we think it would be useful to do so.

There were one or two apparently minor papers copies of which we were unable to obtain. There was one large retrospective management report (82) a copy of which we have not yet received from DoC despite a request and verbal assurance that it would be sent.

We would have liked to reanalyze the data from several critical studies, and have requested copies of the data. In the case of the Spurr invertebrate study (60), reanalysis may have been particularly helpful in reconciling the author's results with a previous study. However, we are informed by the Director of Landcare Research that Dr. Spurr was unable to locate the data and he would not be able "to search" for it until 30 January when this report is due.

The material covered in this document is vast in scope and it has been necessary to do our research in a relatively short time frame. Thus, it is possible, even likely, that we have made some errors of detail. For this we apologize if it turns out to be true. However, we are convinced that the bulk of the evidence is as we have represented it, and thus that the conclusions are substantially as we have stated.

Science, politics and the nature of this document

At the outset it was our intention to confine ourselves to the scientific evidence supporting the use of aerial 1080. However, it quickly became apparent that, although the scientific evidence is far from adequate to justify such an extraordinary national policy of indiscriminately spreading poisoned food throughout whole forest ecosystems, the scientific evidence is not the whole story.

We will show that the manner in which DoC has been interpreting the scientific evidence is as much a problem as is the evidence itself^{*}. There is a pattern of misrepresentation, omission, and distortion in DoC's writings and pronouncements so obvious and so flagrant that the scientific evidence could not be explained in the absence of documenting this aspect as well[†]. It is further clear that DoC-sponsored scientists are under considerable pressure from DoC management to consistently support the agency's statements regarding the real effects of 1080. Consequently, we have in several places documented and explained the political context in which scientific investigations have taken place.

However, we have attempted to make it clear from context when we are presenting a scientific, technical judgment as opposed to explaining the political, bureaucratic, and human context.

In writing this document, we have been candid, sometimes brutally so. We have not attempted to euphemize in order to protect the sensibilities DoC/Landcare Research

^{*} Relative to DoC, we have found much less evidence of misinterpretation on the part of AHB.

[†] The nearly universal presence of misrepresentation by omission is exemplified by a quotation of the purpose for the application from ERMA's web site: "... the reassessment of 1080 and substance containing 1080 (a vertebrate toxin) ...". In fact, compound 1080 is toxic to all organisms that burn carbohydrates to produce energy, i.e., all animals, not just vertebrates.

researchers or to spare DoC's management from embarrassment. We feel that when something as important as New Zealand's rich native forests and national reputation as an environmentally conscientious nation are at stake, it is too crucial a juncture for equivocation or pusillanimity. Rather, we believe that it will require the full force of plain language to effect a change. Nonetheless, it is not our intention to offend gratuitously and we apologize to the extent that we may appear to have done so.

Monofluoroacetate facts

Monofluoroacetate (1080) was originally developed and marketed as an insecticide (67).

It functions primarily by interfering with the citrate step in the Krebs cycle (25). The Krebs cycle is the major and an essential mechanism by which all air breathing creatures utilize food to produce energy. This means that it is toxic to all animals, essentially everything living except perhaps plants and some microorganisms.

Of course some species are more susceptible on a weight basis (Table 1). Remarkably, given that New Zealand uses 80 to 90% of the world's production in our forests (3,4), the susceptibility of most of New Zealand's native species have not been studied, as DoC unashamedly admits (5).

Table 1 Relative Toxicity of Monofluoroacetate

Species	LD50	Relative tolerance
	(mg/kg body weight to kill 50% of a population)	(LD50 for Species / LD50 for possums)
Dog	0.06	0.1
Pademelon	0.13	0.2
Bennett's Wallaby <0.2	0.20	0.3
Cat	0.40	0.7
Rabbit	0.40	0.7
Cattle, sheep, deer 0.2-0.6	0.40	0.7
Red-browed firetail	0.60	1.0
Possum	0.80	1.3
2 AU bird species 0.6-0.99	0.80	1.3
Rat	1.00	1.7
Wombat	1.50	2.5
Man	2.00	3.3
Finches	2.70	4.5
House sparrow	3.00	5.0
Chukar	3.50	5.8
Golden Eagle	3.50	5.8
Sulphur-crested cockatoo	3.50	5.8
Eastern quail	3.70	6.2
Parrots 8 species	4.00	6.7
Tasmanian devil	4.20	7.0
California quail	4.60	7.7
27 AU bird species 1.0-9.9	5.50	9.2
AU Insectivorous birds 3.4-18	7.30	12.2
Mallard	9.10	15.2
Birds 3-19	11.00	18.3
Mouse	13.00	21.7
Great Horned Owl	20.00	33.3
11 AU bird species 20.0-49.9	35.00	58.3

Monofluoroacetate and cyanide

Monofluoroacetate is very similar to sodium and potassium cyanide in its profile as a poison. Both are universally lethal to animals. Both have no antidote. Both are rapid acting, though cyanide is more so. Time to death after monofluoroacetate poisoning is quite consistent among species (Atzert, 1971, 6). Both have low environmental persistence when wet, though monofluoroacetate is more persistent than cyanide. The exact degree of persistence of monofluoroacetate is a matter of dispute in New Zealand. It depends dramatically on circumstances and varies widely but an average is about 50% loss in 24 days in baits (52). Weaver (7) concludes that there is evidence that, since degradation rates vary dramatically with temperature, in some circumstances it may persist for a very long time. This has not been adequately investigated.

Secondary poisoning is possible, and perhaps even frequent, with monofluoroacetate, but essentially impossible with cyanide. Cyanide is cheaper than monofluoroacetate.

The risk to humans is substantial according to the WHO, which classifies both as "1A extremely hazardous" (8). In discussing their relative merits, DoC and AHB listed cyanide as having the disadvantage of "risk to humans if ingested", but surprisingly did not do so for monofluoroacetate despite the fact that as little as 30 mg can be fatal to humans (9). In a definitive review done independent of DoC, Eisler noted (52):

Compound 1080 is highly poisonous to all tested mammals and to humans. There is no known antidote to 1080, and it has been impossible to resuscitate any animal or human during the final stages of 1080 poisoning

• **Table 2. Monofluoroacetate versus Cyanide (25)**

Factor	Comparison
Range of action	Both poison all animals
Human risk	Similar
Antidote	Cyanide advantage
Cost	Similar
Environmental persistence	Cyanide less
Secondary poisoning risk	Cyanide: non-existent
Speed of action	Cyanide faster (10 minutes vs. 1-24 hours), which may lead to higher probability of bait aversion*

Given Table 2, one might wonder what the big attraction is to monofluoroacetate over cyanide. Although we have no direct evidence, the answer seems to be in the politics. Because monofluoroacetate is relatively unknown, especially outside of New

* While this has been asserted by DoC and AHB as a major advantage of monofluoroacetate over cyanide, we can find no published study that would give scientific credibility to that claim, particularly as regards 1080 administered aerially at infrequent intervals.

Zealand, it is politically acceptable to indiscriminately drop food laced with monofluoroacetate into forests, whereas doing the same with cyanide would generate both a national and international outcry that would bring the multimillion dollar practice of dropping tonnes of a universal poison into our forest ecosystems to an immediate halt. In view of this, it is instructive to note how DoC and AHB represent monofluoroacetate versus cyanide (1, 9). They make several insignificant distinctions: they describe both as having "low" environmental persistence, but then fail (as noted above) to mention the human risk for monofluoroacetate.

A brief tutorial on experimental design and statistical inference

DoC and AHB, mostly through Landcare Research, are essentially the only sources of scientific investigation on the question of the effect of aerial 1080 on ecosystems. This is because no other country in the world is doing anything remotely comparable. This means that one cannot challenge the validity of DoC-sponsored research with independent studies done domestically or abroad. There is none. Thus, we must evaluate the quality of DoC research. To do this it is necessary to use accepted standards for experimental design and statistical inference as a benchmark against which to judge the quality of DoC's investigative work. This section reviews these principles.

Our intention is to provide a basic knowledge of the principles of experimental design and statistical inference for people who are not well versed in such arcane matter, so that they can read and understand the information presented elsewhere in this document that presupposes an understanding and appreciation of those principles. We provide here a few references for the principles described below. These span a wide range of detail and sophistication (10, 11, 12), but there are literally thousands of books and textbooks on the subject.

Controls

Virtually all scientific hypotheses have implied controls embedded in them. If one says, "Our forests got worse." The immediate question arises of relative to what have they worsened: relative to Hawaii's forests, relative to what they would be if we did not saturate them with 1080, relative to what they would have been if the possum had never been introduced, relative to what they would have been had Europeans not been introduced, etc.? The statement, "Our forests got worse" is entirely meaningless without the relevant comparison. When formalized into experiments, the comparison entity (or entities) becomes the "control", giving us an anchor from which to judge observed change.

In many respects controls are the key to good research in complex systems. The quality of the control(s) predetermines the quality of the scientific investigation, and to a substantial degree, the quality of the control group determines the validity of the results and the strength of the conclusion.

Controls can be categorized into a hierarchy.

Level 0: No control group at all

This is the category into which fits DoC's statements on the overall effect of aerial 1080 on our forests, namely uncontrolled observation (often by biased individuals). For example, in DoC's premier brochure advocating aerial 1080 (13) we find this statement regarding "mainland islands":

"Using 1080 in these forests has been successful in helping restore birdsong that was diminished before 1080 was first used."

ignoring the fact that "mainland islands" are more comparable to real islands than the forests usually poisoned by DoC^{*}, this assertion is based on nothing more than opinion, i.e., uncontrolled "observation". It is not based on science. It is an anecdote and as such is more likely to represent the prejudice of the writer than truth.

Level 1: Historical controls

In this case, the experimental group is compared to a previous state of the system under investigation. Many DoC studies fall into this class. Such controls have two major problems. First, historical circumstance is often not comparable to the current ones and, second, it is impossible to determine the cause of any observed difference (or lack of difference) between the control and experimental observations. In addition, historical controls are often accompanied by retrospective observations, which are notoriously unreliable. The literature is filled with examples of historically controlled research that turn out to be false when examined with simultaneous controls.

Level 2: Simultaneous controls

A few of DoC's studies have simultaneous controls. There are perhaps a dozen that bear directly on the question of the effect of aerial 1080 on our forests. The problem with simultaneous, but not randomized, controls is that one never knows whether the controls are really comparable to the experimental group. The chances that the controls are inherently different from the experimental group can be reduced by two techniques, and definitively eliminated by one[†].

Level 3: Simultaneous, matched controls

One can carefully examine what are thought to be relevant factors to assure comparability and attempt to prove this comparability statistically. Of course, one can never be sure that s/he has gotten all the relevant factors or that the factors examined are the correct ones. DoC sponsored studies almost never do this kind of comparability checking. In fact often (as will be seen) they simply ignore clear evidence of incomparability.

Level 4: Simultaneous, matched controls with diversity and multiplicity

One can have multiple and varied control and experimental areas that truly represent the range of conditions to which the study will be applied. None of the research that DoC cites to support its use of aerial 1080 reaches this level of control quality, and indeed, any study that did would have been most likely to go on to Level 5, randomized controls.

Level 5: Simultaneous, randomized controls

This is the highest standard of control quality. Really it should be Level 10 since none of the others approaches its ability to insure reliability of results. The concept of

^{*} ... because repeated applications of 1080 are not usually necessary on islands.

[†] Strictly speaking randomization does not eliminate the possibility that control baseline characteristics account for an observed difference. Rather randomization allows a researcher accurately to calculate the probability of that possibility, and thus consciously to decide how much chance of making an error he is willing to take.

randomization in research design was originally developed by R A Fisher* in the 1920's to support agricultural and genetics research. Randomized design is now the gold standard for experimental research in complex systems, for example, in clinical medical research and in biological systems. Though little known to the general public, it is among the most important discoveries of all time. The reasons for its power are subtle and deep, and beyond the scope of this brief discussion. It will suffice here to describe its effect in experimental inference. It removes the influence of most forms of bias, it validates the assumptions underlying the statistical tests, and it is the only way to prove causation in multivariate systems with substantial variation among analyzed parts. In the particular case being addressed in this submission, it is that aerial 1080 causes benefit or harm to our forest ecosystems. In short, it is the only path to the unvarnished truth.

We can find no DoC sponsored study in which the selection of control and experimental units was randomized—none, let alone one that bears on the issue of the effect of poisoning our forests with aerial 1080. The existence of one such study addressing the relevant questions would trump all the other "research", opinion, tradition and propaganda put together. Despite decades of dropping 1080 into our forests and despite hundreds of millions of dollars having been spent, that one essential study has not been done.

Blinded observation

Blinded observation in study design is the use of observers and assessors of experimental results who are unaware of the control status of the observations they are making. It is vital when the variables being observed are subject to judgment, which is virtually always true in biological field studies such as the ones we have reviewed in this paper. It prevents observer bias from influencing the outcome of a study. Most observer bias is not conscious or malicious. It is simply a function of being human. For example, people examining aerial photos of a forest to determine the degree of deforestation from possums is very subjective. None of DoC's studies that we have reviewed have blinded observers.

P-values

A P-value is the probability that a particular statistical result could have happened by chance. The lower the P-value the less likely that an observed difference (between treated and control area) was due to chance. By convention, scientific results are generally not considered to be "statistically significantly different" unless the P-value is less than 0.05 which mean there is a 5% percent chance that the observed difference was just an accident.

It is important to understand when reading scientific papers that the term "significant" usually means "statistically significant", and it bears no relationship to the concept of scientific significance. Thus, a difference might be statistically significantly different but not scientifically important, or it might be scientifically important, but not statistically different.

For example, a 1% drop in robin population numbers from aerial 1080 might be statistically significant (and thus real and reproducible) if the number of observations was great enough, but few would argue that it was ecologically or scientifically important. On the other, hand a 50% drop would certainly be ecologically important, but if the P-value were too large (>0.05) then it should be ignored, except, of course,

* Sir Ronald Aymer Fisher, (17 February 1890 – 29 July 1962) was a British statistician, evolutionary biologist, and geneticist. He was described by Anders Hald as "a genius who almost single-handedly created the foundations for modern statistical science" and Richard Dawkins described him as "the greatest of Darwin's successors", high, but highly deserved, praise.

as a guide to future research. This is not just mathematical sophistry. The consequence of disregarding these principles is that one will end up drawing a lot of false conclusions (and in the case at hand, might end up doing vast damage to our forest ecosystems).

As will be seen below, P-values were not calculated for many of the results on which DoC bases its claims of benignity and benefit of aerial 1080. In at least one case they were calculated selectively, which allowed the DoC sponsored researchers to claim a benefit to robin populations that did not exist and that was not reproduced later in the same study.

Confidence intervals and statistical power

Confidence intervals (CI), when appropriate, give some of the same information as formal power calculations (see below) and are much easier to understand. Most confidence intervals are calculated for a 95% confidence or a 67% confidence^{*}. Roughly, a 95% confidence interval tells one the range over which 95% of results would occur if the same experiment were done repeatedly.

Perhaps an example will help. Let us suppose that we did two experiments: one in which 4 of 10 robins died of aerial 1080 and a second in which 40 of 100 robins died. In both cases 40% died. This is the way DoC typically reports its results. However, common sense tells one that these are very different results. One would have much more "confidence" in 40/100 than 4/10. Confidence intervals quantify that "confidence" and express it in a standard form so that it is easily understood. The 95% CI for 4/10 is (19% to 74%), for 40/100 is (31% to 50%), and for 400/1000 is (37% to 43%). If 4/10 is the result, we know with 95% confidence that the true value for robin deaths is between 19% and 74%, which usually would not be close enough to make a decision about aerial 1080 in our forests, whereas 31% to 50% probably would be enough, and 37% to 43% would be overkill and a waste of scarce research resources. The point is that confidence intervals tell us how accurately a particular result is known and thus how much "confidence" we should put in them. Without them statistics are not interpretable and useless, or worse, misleading.

"Statistical power" is a more difficult concept, but it is vital when one is attempting to show that there is not an "important" difference between experimental and control values (e.g., robin populations after aerial 1080). Confidence intervals provide something of the same information as statistical power once the study is complete, but statistical power calculations done before a study is started allow one to design the study to have a predetermine probability of detecting a certain difference between treated population and controls. It allows the researchers to set their chance of drawing a falsely negative conclusion (that there is no important difference between control and treated populations). In what follows in this paper, we will see example after example of DoC sponsored researchers concluding that there was no difference between 1080 poisoned native species and those not poisoned when they had merely failed to detect a difference because the statistical power of their research was insufficient. It is not an exaggeration to say that this statistical error is the basis of most of DoC's claim that poisoning with 1080 is benign to native species.

For at least 30 years, since the age of computers, power, P-value and confidence interval calculations have been trivial to do. There is no excuse for not including them in published reports. Put bluntly, any researcher that publishes summary statistics without P-values and either power calculations or confidence intervals is either deliberately deceptive or incompetent. There are no other choices. The manner in

^{*} A 67% confidence interval is conventionally called the "Standard Error".

which DoC researchers have used P-values and power calculations (and more often not used them) will be seen below.

To some this may seem daunting and difficult to understand. However these people need not despair of being able to judge for themselves the quality of quantitative research. Just follow this rule: if a percentage or average is not accompanied by a P-value or confidence interval, it is worthless, or almost so, and should be disregarded. Think of the 40% example above.

The role of random sampling

Strictly speaking one can only generalize results to populations that are randomly sampled. However, true random sampling is rarely done*. Instead scientists rely on including in their study populations, multiple and varied representatives from the population to which the results will be generalized. The importance of this depends on how varied the subjects are known to be. Most scientists would agree that a breast cancer victim in New Zealand is quite similar to those say in the United States. Thus, results of studies done in the United States are assumed to "generalize" to New Zealand women. However, that is certainly not the case for forest ecosystems. So if we wish to generalize results to all of New Zealand's forests†, we MUST study a representative (if not random) sample.

All of the controlled studies regarding the effects of aerial 1080 on New Zealand forests involve a very few sites, usually less than three, and always relatively close to each other. Thus, the generalizability of all the claims is suspect.

More than one study by more than one group of independent investigators

The essential element that distinguishes an experiment from other kinds of organized observation is reproducibility. Before any assertion based on experiment can be considered a scientific fact, it must be reproduced by others who are socially, academically, and financially independent.

DoC sponsored research has often not been reproduced, and none of it is independent of DoC influence and therefore its agenda.

The absolute need for researcher independence: the human factor

Another reality driving the need for diversity and independence in research might be called the human factor. Scientific research is a struggle engaged in by people who are often, if not usually, passionately committed to their efforts. Their reputation, status and financial well-being frequently depend on being correct and getting positive results. Anyone who is honest with himself and has been there can tell you of the pressure and the tendencies that are consequent. One does not lie or actively

* (Warning: only read this footnote if you are really into experimental design. Otherwise it is not essential, or even desirable for that matter.) Random sampling refers to taking a sample randomly from the population to which one intends to generalize his results. Randomization of control and experimental groups is somewhat different and is almost always done in good research when it is possible. It accomplishes most of the benefit of a random sample from a population, but means that the observer is left to judge whether the set of study subjects faithfully represent the population to which one wishes to apply the result. For example, if DoC wished truly to discover the effect of aerial 1080 on our forests, it would first randomly sample plots from our forests and then randomize those plots to determine which were to be "treated" with aerial 1080 and which were to be "treated" with nothing or ground control or whatever. However, the first step might not be possible because not all forests were equally available. A reasonable substitute would be to select a "representative" set of plots and then randomize them as to "treatments".

† ... which we do since DoC is actively "treating" them with the same "therapy", or at least intending to do so.

misrepresent. He does not need to. It is easy enough to convince oneself of the "good reasons" why this result was flawed and should not be published, or why the statistical tests should be done this way or that way. The situation is worse when experimental conditions are difficult to control, as is usually the case in clinical medicine and environmental research. The net result is that many, perhaps as many as 60%, of positive results turn out not to be reproducible. Rigorous and prospective study design and strict adherence to protocols help, but the only real antidote to this very human problem is to insist that results are independently reproduced by others.

The absolute need for financial researcher independence

There is another kind of independence that is needed: financial independence. Any experienced scientist will testify that even for the simplest experiment there are a thousand ways to influence the way the results appear when finally published: choice of controls, how exceptions are dealt with, choice of statistical tests, the choice of what tests to report, where in the paper a fact is placed, conclusions, etc. The list is almost endless. Again randomization, blinding, formal protocols, and multiple researchers can largely obviate the inadvertent influence that bias and self-interest will introduce. Recognizing this, the Federal Drug Administration that authorizes all drugs and medical devices in the United States requires that pharmaceutical companies pay for multiple studies, usually randomized and double blind at multiple sites. As has been repeatedly pointed out, DoC sponsored researchers are not financially independent of DoC. This flaw in the execution of the aerial 1080 research should shed real doubt on its validity, especially when coupled with all the evidence of bias in the published reports themselves. In addition, none of the DoC and AHB research has been done with study designs that tend to immunize against influence and bias, and consequently virtually all have the taint that financially sponsored research inevitably engenders.

Does it matter?

Many people reading this document will be asking the question, "Does it really matter? Is it not good enough to do Level 1 or 2 research; after all, one will get it right most of the time." As the head of the Northern Coromandel Biosecurity Subcommittee, Douglas Wright commented to me in a startlingly unselfconscious communication defending the lack of good science supporting aerial 1080: "management trials" (read: Level 0 or 1, controls) are what has been and should be used.

Scientists have a saying, "If you wish to know a researcher's prejudices, read the results of his last uncontrolled study". One does not get half the truth with a half-good research design. One gets a result that will reflect the bias of the researcher, which may or may not be the truth.

It is not possible to prove this assertion, but it is possible to illustrate it with a particular case: human clinical research. In the last 60 years, clinical research has gradually evolved from what may be called organized anecdote* (which did little more than perpetuate rumors) into an experimental standard for clinical truth that can be summarized as randomized, double-blind controls with full statistical disclosure. This transformation has revolutionized clinical healthcare throughout the world because it means that clinical knowledge is no longer dependent on anecdote, opinion, or individual experience. Clinical knowledge can no longer be held captive to the prejudice of well-meaning advocates or of self-serving profiteers. Vanity and political power have taken back seats. Individual whim, academic position and self-aggrandizement no longer dictate clinical truth.

* Very much like the DoC's claims for the benefits and benignity of aerial 1080.