ACERA Project
0703

Title
Combining disparate data sources to demonstrate pest/disease status

Author(s) / Address (es)
Tony Martin / Department of Agriculture and Food, PO Box 1231, Bunbury WA 6231

Material Type and Status (Internal draft, Final Technical or Project report, Manuscript, Manual, Software)
Final Project report

Summary

The project aims to adapt existing methodology for evaluation of varied surveillance activities for pest- or disease-free status, so that it can fill the needs of both animal and plant domains. The major pre-existing methodology was developed for animal health needs, and this project will deliver two case studies (analyses) of plant and nuisance pest surveillance programs. In so doing the project will review the methodology and its current implementation (software), and will clarify and define issues for further research which pose obstacles to its adaptation, its use, and its general acceptance by end-users and trading partners.

Project schedule

The approved, revised schedule of project activities is as follows (progress towards each milestone since the last report is presented in *italics*):

Date	Task
29-31 May 2007	Initial project workshop, Canberra Successfully completed
9 – 11 October 2007	Case study workshop, Melbourne The workshop to develop further the two case studies (on karnal bunt of stored grain and the urban surveillance program for red imported fire ants (RIFA)) was held successfully in October. Nichole Hammond, Stephen Pratt, Tony Martin, Paul Pheloung, Angus Cameron and Mark Burgman attended. Nichole was close to completing the karnal bunt study, while Stephen used the workshop to solidify his plans for the RIFA study.
14 December 2007 *June 2009*	Case studies completed; draft reports delivered to project leader *Both case study reports have now been received, and are incorporated in this report.*
28 February 2008 *June 2009*	Review of methodology completed A meeting of the review team (Ray Chambers; Liwan Liyanage; Samsul Huda) with Tony Martin is planned for late January. Actual timing is dependent on the delivery of the case study reports, which are essential reading for the review team. *The review was, of necessity, completed before delivery of the fire ant case study report. It therefore does not make reference to the fire ant case study.*
28 February 2008	Additions and revisions to software completed Under way. The software enhancements for this project do not involve work on the

June 2009	analytical programme (written in R), but rather on the user interface (written in PHP). *The software adaptations for the project have been completed, and are described briefly in this report. The software has undergone a major re-write since the draft Final Report was submitted for this project.*
31 March 2008	Information workshop for potential end-users *This workshop took place, with support/funding from the CRC for National Plant Biosecurity and the Australian Biosecurity CRC. An audience of 35 from DAFF (BA, BRS, OCCPO, OCVO); QDPI&F; NSW DPI; ACT; PHA; PIRSA and CSIRO was entertained by the project team during the morning, with presentations on each of the case studies, and more generally on methodology. In the afternoon group discussions canvassed participants' views on the project's achievements and on requirements of future developments.*
30 April July 2008 *June 2009*	Final report delivered to ACERA Draft Final Report (without fire ant case study report) delivered July 2008 *The final Final Report is appended.*

Project management

All the contracted services were have now been completed and payments have all been made. The project was completed late, but on budget.

Combining disparate data sources to demonstrate pest/disease status

ACERA Project No. 0703

Tony Martin
Department of Agriculture and Food Western Australia

Final report; June 30th 2008

Acknowledgements

Disclaimer

This report has been prepared by consultants for the Australian Centre of Excellence for Risk Analysis (ACERA) and the views expressed do not necessarily reflect those of ACERA. ACERA cannot guarantee the accuracy of the report, and does not accept liability for any loss or damage incurred as a result of relying on its accuracy.

Table of contents

List of Tables

List of Figures

1. Executive Summary

Background and Aims

In line with the Sanitary and Phytosanitary (SPS) Agreement of the World Trade Organization (WTO), assertions of pest- or disease-free status made to trading partners must be scientifically justified. It is not possible to determine pest-/disease-free status with absolute certainty, so both international standards and trading partners adopt probabilistic assessments of status against agreed standards. Classically, quantitative evidence for pest-/disease-free status is derived from a stratified random survey of the population at risk, but in recent years there has been much interest in quantifying the results of non-random and general surveillance to reduce dependence on random surveys, as well as to provide more convincing, ongoing assessment of the probability that a population is free from a pest or disease. A methodology for achieving these objectives in animal disease surveillance has recently been developed; this project aimed to investigate application of the methodology to plant pests/diseases and to invasive pest species, through the following specific activities:

- Two case studies; one for a plant disease and the other for an invasive pest.
- Review and modification of web-mounted software
- Review of the methodology and its application to plant pests/diseases
- Workshop to communicate the results of the project to potential end-users, and gather ideas for further research and development in this area

Methods and Achievements

Stochastic scenario tree modelling of surveillance system components (SSCs) was applied in the two case studies:

- Karnal bunt surveillance in Western Australia: Karnal bunt is a fungal infection of cereal grains (wheat; durum wheat; triticale) causing contamination of stored grain and unsuitability of the grain for baking, etc. Two SSCs were evaluated: testing of delivery parcel samples and general siding samples

- Inspection of high-risk sites for Red Imported Fire Ants (RIFA) in the national urban surveillance program.

The end result of analysis of a single SSC is the sensitivity of the SSC, or the probability that it would detect the pest or disease if it were present at an agreed level (the design prevalence). We also used these results to estimate the probability that the relevant areas (WA for Karnal bunt; individual cities for RIFA) were free from the pest. These analyses highlighted issues associated with the different domains of application of the methodology, in particular in relation to design prevalence, defining the surveillance unit for analytical purposes, incorporation of continuous rather than categorical inputs, and media (diagrams and software) for representation and implementation of the models. All were constructively resolved and are addressed in this report.

This project brought together collaborators from government (State and Commonwealth), university and private consultancy; from plant pathology, invasive pest surveillance and animal health; and from the disciplines of mathematics/statistics, biology, epidemiology and modelling. We held 3 meetings of the project team to address general and specific methodological and terminological issues, and drew up a list of terms appropriate for application in plant pest/disease and invasive pest surveillance. The web-mounted software

previously developed for animal health analyses was then adapted to incorporate terminology relevant to these domains.

We held a one-day workshop to present our results to potential end-users and discuss with them their views on, and expectations for, the methodology. This was attended by over 30 and the main needs identified were for

- general acceptability of the results to trading partners;

- accessible means of analysis;

- ability to accommodate and use spatial data in determining risk of infestation and probability of detection for individual surveillance units.

Conclusion and Recommendation

This project has achieved its aim of evaluating the suitability of this methodology for plant pest and invasive pest applications, and has demonstrated such suitability. The wide range of participants ensured thorough evaluation from numerous perspectives. The strong interest from potential end-users confirms a demand for this methodology, and for its further development.

We believe it would be appropriate for a further project to develop suitable methods, techniques and software for facilitating analysis of surveillance activities in which continuous, spatially-defined variables are important determinants of probability of infection/infestation and/or probability of detection.

2. Introduction

The Sanitary and Phytosanitary Agreement of the World Trade Organization stipulates that member countries may take measures to mitigate sanitary and phytosanitary risks associated with trade in animals, animal products, plants and plant products, provided such measures are supported either by international standards or, in the absence of such standards, by scientific justification. In any reference to international standards or any import risk analysis the starting point is a determination of the pest or disease status (*i.e.* presence or absence of the pest or disease) of the importing and the exporting countries with regard to the pest(s)/disease(s) under consideration. It follows that determination of such status must be science-based. It is not possible to determine pest-/disease-free status with absolute certainty, so both international standards and trading partners adopt probabilistic assessments of status against agreed standards. Classically, quantitative evidence for pest-/disease-free status is derived from a stratified random survey of the population at risk, and other evidence supporting free status is assessed purely qualitatively, and is only used to support the conclusion drawn from the random survey. This approach has various drawbacks, principally the high cost of conducting random surveys to provide sufficient evidence, the ephemeral nature of the results of, and conclusions drawn from, such surveys, and the fact that most of the evidence for freedom is ignored. This other evidence is derived from general, passive and non-random surveillance activities, and in recent years there has been much interest in utilising the results of non-random and general surveillance to reduce dependence on random surveys, as well as to provide more convincing, ongoing assessment of the probability that a population is free from a pest or disease. To do this requires quantitative evaluation of the surveillance activities involved, and quantitative integration of the evidence for freedom derived from multiple, varied sources, including random surveys. A methodology for achieving these objectives in animal disease surveillance was described by Martin et al (2007). This project aimed to investigate application of the methodology to plant pests/diseases and to invasive pest species, through the following specific activities:

- Two case studies; one for a plant disease and the other for an invasive pest.

- Review and modification of web-mounted software implementing the methodology for animal disease applications in order to make it applicable to plant pest/disease applications.

- Review of the methodology and its application to plant pests/diseases

- Workshop to communicate the results of the project to potential end-users, and gather ideas for further research and development in this area

3. Methodology

The methodology forming the basis of this project provides a framework for estimation of the probability that each component of the surveillance system (each SSC) will give a positive outcome, given that the population is infested at a predefined level (the design prevalence), using a probabilistic model of each SSC. The surveillance system, and the SSCs of which it is comprised, are considered analogous to diagnostic tests applied to the population. Scenario (or event) trees (Hueston & Yoe, 2000; Martin *et al* 2007a) are drawn for each SSC to define the steps necessary for a surveillance unit in the population to give a positive outcome in the surveillance process. All factors affecting the probabilities that a unit will be infested or that an infested unit will be detected are included in the scenario tree model of the process. The probability of a positive surveillance outcome (the sensitivity of the SSC) may then be combined with those of other SSCs to give an estimate of the sensitivity of the whole surveillance system; this in turn may be used to estimate the probability that the population is free from infestation (at the design prevalence) given that the surveillance system had a negative outcome, a probability analogous to the negative predictive value of the diagnostic test on the population. Ongoing evidence for population freedom (*i.e.* negative surveillance outcomes) may be accumulated over multiple surveillance time periods using appropriate adjustments for the probability of introduction of infection / infestation into the population.

A key assumption in the application of this methodology is that in a pest/disease free situation the specificity of the surveillance system, and of each of its components, is perfect; in other words, there are no false positive results. The reason for this is that positive identifications of the pest/disease are not compatible with freedom from the pest/disease, so each SSC includes all the follow-up testing necessary to resolve any positive test result into a true positive or true negative.

Specific issues addressed in this project included

- *Terminology*

    Much of the terminology used in evaluation of surveillance is generally used and understood, being applicable in consideration of one or more of the following

    - agricultural and environmental biosecurity
    - surveillance in the context of biosecurity
    - mechanistic modelling of biological processes
    - probabilistic modelling and analysis
    - epidemiology

    Many terms necessary for the analysis of disparate data in assessing pest/disease free status have been defined by Martin et al (2007a) in the animal health domain; most of these are equally applicable in the other domains under consideration in this project, i.e. plant health and invasive pests. Some essential and generally applicable terms are presented in Table 1.

Table 1   Terms applicable to plant, animal and environmental surveillance analysis

Term	Meaning
Surveillance system	All the surveillance which is conducted for a particular pest/disease
Surveillance system component (SSC)	A single surveillance activity contributing to the surveillance system, which may comprises multiple SSCs.
Surveillance time period	The period of time over which results from an SSC are accumulated for analysis as a single block of information (often a year, or a month for a rapidly moving animal disease)
Surveillance unit	The basic unit processed in the SSC – often the individual animal or plant (see below)
Sensitivity	Probability of a positive outcome (in a test, SSC, surveillance system, etc.) given that the pest/disease is present; equivalent to the diagnostic sensitivity of a laboratory test
Risk of infection/infestation	Proportion of surveillance units or groups of units which become infested/infected when exposed to a risk (predisposing) factor
Relative risk of infection/infestation	Risk or infection/infestation in one group relative to another (i.e. ratio of risk in one group to that in the other).
Probability of freedom	The probability that the pest/disease is either not present in the reference population, or is present at less than the design prevalence
Design prevalence	Hypothetical level of infection/infestation in the reference population, set for purposes of estimating the sensitivity of a testing/surveillance program applied to the population.

- *Surveillance unit*

  The surveillance unit is the basic unit which is tested/observed/evaluated in the surveillance process under consideration. In many animal health applications, this is the individual animal, which may be seen to be sick, or may have a sample taken from it on which a laboratory test is performed. In the context of surveillance for plant pests/diseases and invasive pests, the surveillance unit, or that which is assessed by the surveillance process, is most commonly an area of land. In surveillance for pests of a harvested commodity, the surveillance unit will be a standard measure of the commodity, *e.g.* a tonne or a kg or a carton. See Table 2 for a wider range of possible applications.

- *Group or cluster of surveillance units*

  In animal health applications it is commonly assumed that contagious (and other) disease clusters in management groups of animals (commonly herds or flocks). This is critically important in analysis of negative surveillance findings for disease freedom, since individual units are clearly not independent of each other in terms of their probability of being infected, and multiple negative unit-level results from a single group do not carry the same value as negative unit-level results from multiple groups. Analogous clustering / grouping levels in plant pest and invasive species surveillance are not immediately obvious, if indeed they exist at all. Table 2 shows some potential examples of such groupings, but we emphasise that each application, with its own combination of pest/disease and host/environment, must be considered individually.

- *Design prevalence*

  This is the standard, hypothetical, infection/infestation level for which a surveillance system is evaluated. In the context of animal health the meaning is clear: it is the

proportion of animals (or groups) infected in the reference population. For pests/diseases of plants, design prevalence is best conceived as a concentration of infested plants in an area, and will most commonly be expressed as one infested plant per surveillance unit, where the surveillance unit is an area of land of fixed size; *e.g.* one infested plant per hectare, or one per 100m^2. Its units are therefore area^{-1}. For invasive pests design prevalence will take a value of one "pest unit" per surveillance unit, *e.g.* one ant nest per hectare. In the context of diseases of harvested agricultural products (*e.g.* wheat) the design prevalence will relate to a standard quantity of the product, *e.g.* a carton or a kilogram.

Table 2   Examples of appropriate terms for surveillance units and groups in different contexts

Example Context	Unit[*]	Group[**]
Terrestrial animals	Animal	Herd
Aquatic animals	Animal	Cage, pond, net, lake
Bulk handled product • Karnal bunt of grain	50 gram sample	Lot
'Continuous' host or habitat • RIFA (urban surv.) • Wheat aphid	Square metre	Hectare Square kilometre
Trapping • RIFA (drift net) • Coddling moth • Fruit fly • Urban surv.	Trap catchment area • Linear metres • Square metres (to probability threshold)	Orchard, Block Trapping area 'Area of interest' 'Quarantine approved premises' 5 kilometre grid square
Inspection of fruit • Fruit fly (apples)	Apple	Orchard Block (Possibly tree as second grouping level)
Inspection of plants • Banana • Tomatoes	Tree Plant	Property
Vegetable crop • Onion smut	10 metres of row	Field
Public reporting • RIFA	Property	Suburb
Weeds	Square metre	Weed management area Square kilometre Hectare

[*]**Unit**:   That to which the conclusion of presence or absence (of pest or disease) applies; that which is 'tested' or observed.

[**]**Group**: Potential clustering level

3.1.   Case studies

Two case studies were decided upon to evaluate the application of the methodology across different potential scenarios, covering a fungal plant disease, an invasive insect pest, a targetted (non-random) SSC, an SSC comprising representative sampling, and general ("passive") surveillance activities. Our initial intention was to include the combination of multiple SSCs in these studies, but time constraints resulted in only one SSC being analysed for each case study, and neither of these involved general/passive surveillance. The case studies and SSCs selected were:

1.   Surveillance for karnal bunt (*Tilletia indica* infestation) of stored grain in Western Australia. This study was conducted by Nichole Hammond, a plant pathologist and PhD student at Murdoch University, as part of her PhD. Nichole's selected SSC was a targetted program of sampling and testing of stored grain, conducted over three years.

2.   Surveillance in urban areas for red imported fire ants (*Solenopsis invicta*; RIFA). This invasive pest is established in the Brisbane area, and is one of the target species in the national urban surveillance program. A range of surveillance activities are undertaken at the State level aimed at "pre-emptive" detection of RIFA in and around unaffected cities. Stephen Pratt and Greg Hood of the Bureau of Rural Sciences in the Department of Agriculture, Fisheries and Forestry carried out this study in collaboration with Paul Pheloung of the Office of the Chief Plant Protection Officer, which administers the urban surveillance program. The selected SSC comprised targetted inspections of high risk premises, which is an ongoing program carried out in all States.

A brief outline of each case study follows, highlighting issues of significance to the project's objectives. Full reports on each case study are appended to this report; karnal bunt in Appendix 9-1, and RIFA in **Error! Reference source not found.**.

3.1.1.   Karnal bunt

3.1.1.1.   Background

Karnal bunt is a fungal disease of wheat, durum wheat and triticale caused by *Tilletia indica*. Karnal bunt causes portions of the seed to be replaced by masses of teliospores known as sori and can be spread in infected or contaminated grain, and on contaminated machinery and agricultural products. Karnal bunt is an important pathogen in international trade, with many countries having phytosanitary restrictions in place to mitigate the risk of introduction, and would cause considerable damage to Australia's economy through loss of domestic and international markets if it were to become established. Maintaining pest free status for Karnal bunt is important in maintaining Australia's grain export markets.

Many countries require a declaration of freedom from Karnal bunt for imports of grain. Under the Agreement on the Application of Sanitary and Phytosanitary Measures decisions on restrictions relating to the trade of plant products must be "technically justified" and based on scientific evidence (World Trade Organisation 1995). Therefore, there is a need for methods to evaluate both specific and general surveillance data to justify claims of pest free areas for Karnal bunt.

In Western Australia Co-operative Bulk Handling (CBH) is the main company that handles and stores grain prior to export. Grain is delivered to CBH by the grower, at a number of sites located throughout the Western Australian wheatbelt known as receival sites. Each delivery of grain by a grower is considered a delivery parcel. Delivery parcels are sampled at delivery to determine the grade of grain being delivered, and to check for abnormal seed and contaminants. A portion of these samples can also be collected for other tests, such as testing for Karnal bunt. General siding samples (GS samples) are bulked samples of a portion of the grain sampled from each of the delivery parcels. Each of these samples represents the bulked grain received at a particular site of a particular grade of grain. These samples are collected for further laboratory testing for purity, moisture and protein analysis and portions of these samples may also be collected for other tests.

A surveillance program for Karnal bunt in grain samples was initiated in 1997 by the Department of Agriculture and Food Western Australia. This program ran for three years and included samples from the 1997/98, 1998/99 and 1999/2000 harvests. Grain samples were collected from delivery parcels and from General Siding samples, and tested for teliospores of *Tilletia indica* using the sieve wash method and microscopic examination for teliospores in the collected particulate matter. Assessment of the efficacy (sensitivity) of this surveillance activity forms the subject of this case study, and a full report is appended as Appendix 7-A. Other surveillance system components are in place for Karnal bunt in WA, and these are also being assessed by Nichole as part of her PhD, but were not included in this case study.

3.1.1.2. *Scope of analysis*

The surveillance activities under consideration were modelled as two separate SSCs:

- Delivery parcel SSC;
- General Siding (GS) sample SSC.

The reference population for this analysis is the grain of susceptible species grown in WA.

The surveillance time period for the analysis was a growing season.

Surveillance unit In both of these SSCs the units of surveillance are not convenient, discrete entities, but rather a "continuum of opportunity" for the occurrence and detection of the fungus. The surveillance unit in each SSC is given by the SSC's name; the delivery parcel sample and the GS sample, the sizes of which were therefore dependent on the sampling process. In both SSCs the units (or potential units) were considered to be grouped in *lots*; a discussion of the meaning of this term is given in Appendix 9-1 paragraph 2.3.3.

3.1.1.3. *Scenario tree model*

The same scenario tree (Figure 1) was used for both SSCs, the differences between them being the nature of the lot and the sampling process. Separate stochastic, quantitative models of the sampling process were used to model the probability that individual test samples were infected at the design prevalence, based on within-lot design prevalence, variations in the size and constitution of lots, and sampling variability. Risk factors for infection were the climatic region of the state and the host species, both of which operate at the lot level, since lots are confined to individual regions and to one type of grain. They are represented in the tree by the nodes REGION and HOST. The sensitivity of the diagnostic test varies with spore concentration, so the tree also contains the detection category node LEVEL OF SPORES, which

categorises units (grain samples) by the concentration of *Tilletia* spores they contain. There is a single detection node, representing the sieve wash test.

Two levels of design prevalence were used in this analysis; the proportion of lots infected ($P^*DeliveryParcel$), and the proportion of grains infected within an infected lot. Values used for the former (1%; 0.5%; 0.1%) are based on appropriate levels for early detection of an outbreak, and on prevalence levels observed in countries where the disease occurs. Values for P^*Grain, the within-lot design prevalence, (3%; 1%; 0.0004%; 0.00004%; 0.000004%) were based on the range of observed levels in outbreaks in other countries, and on the "acceptable" level for general use in flour production from affected grain.

3.1.1.4. Implementation

The models were implemented in Microsoft Excel, with the stochastic modelling add-in PopTools (Hood, 2008).

3.1.1.5. Results

The sensitivity of the surveillance (*CSe_DeliveryParcel*) conducted by testing samples from delivery parcels during the 1997/98 harvest was 0.85, 0.62 & 0.18 for *P*DeliveryParcel* = 1%, 0.5% & 0.1%, across all values for *P*Grain* from 0.00004% to 3%. *CSe_DeliveryParcel* only decreased slightly at the lowest value used for *P*Grain* (0.000004%).

Figure 1 Simplified scenario tree structure for Karnal bunt SSCs modelling samples from both trucks and receival bins.
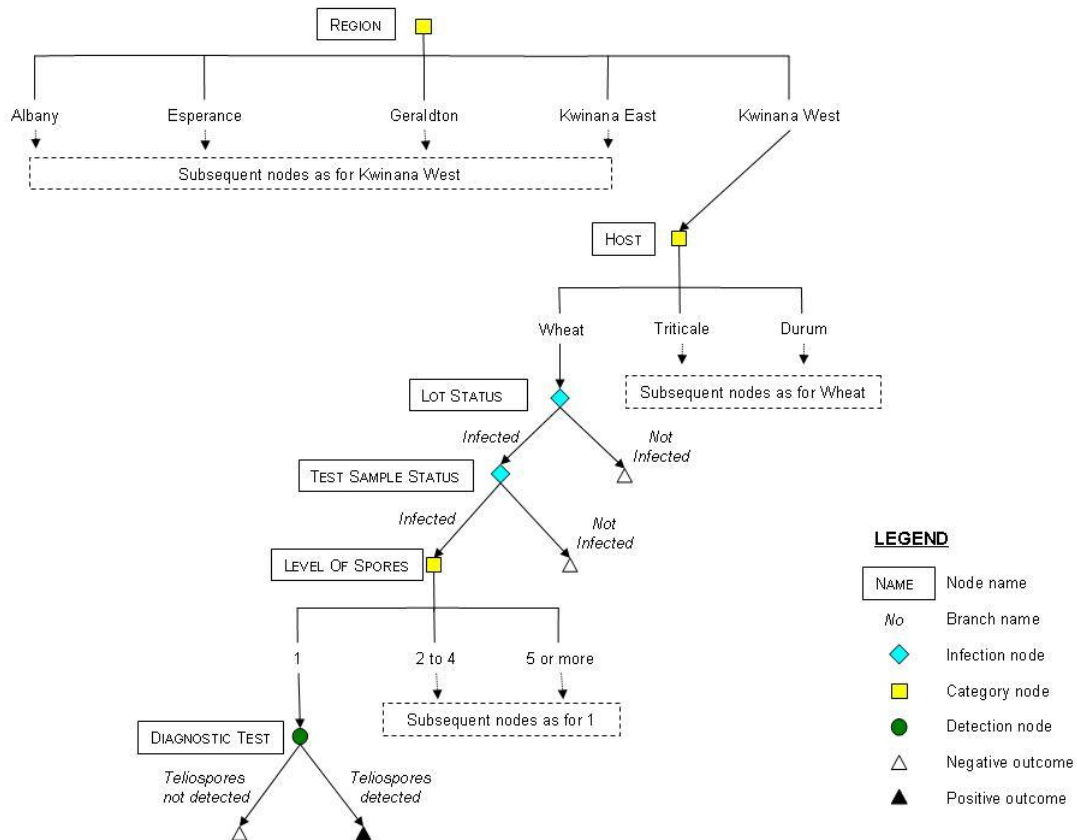
Annual sensitivity of the GS sample SSC, *CSe_GS*, was close to 1 for all combinations of design prevalences, except the lowest P^*Grain of 0.000004% (one infected grain per tonne) for which it was 0, and the combination of P^*Grain = 0.00004% and $P^*DeliveryParcel$ = 0.1%. See Appendix 9-1, tables 14 and 15 for detailed results.

These two SSCs both involve performing the sieve wash test on samples of harvested grain delivered to the bulk grain handling system in Western Australia. They have different component sensitivities, and the GS sampling program was more likely to detect Karnal bunt in WA grain during a year than was the Delivery Parcel sampling. Since the two SSCs are effectively testing the same grain using the same test, they are not independent of each other, so do not have any "additive" effect in terms of surveillance system sensitivity. For these reasons, *CSe_GS* was used to estimate the probability that WA was free from Karnal bunt (at the design prevalence). This estimation uses the process of Bayesian revision of an initial, prior, estimate of the probability of population freedom, based on the additional evidence for freedom given by the negative surveillance findings from a SSC with sensitivity as calculated. For the prior estimate of the probability that the population was infected with Karnal bunt a figure of 0.5 was used (equivalent to "we have no idea whether it is present or not").

Surveillance evidence from each year of GS sampling was accumulated (see Appendix 9-1 section 4 for details), and for all but the lowest combinations of design prevalences, the probability of population freedom at the end of the third year of surveillance ($PostPFree_{1999/2000}$) was effectively 1.

In sensitivity analysis conducted at the single design prevalence combination of P^*Grain = 0.00004% and $P^*DeliveryParcel$ = 0.5%, each input was evaluated over the full range of its input distribution, and none had any significant effect on $PostPFree_{1999/2000}$, which remained very close to 1.

3.1.1.6. Discussion

This case study has shown that this method can be applied to a plant pest and is potentially suitable for the evaluation of plant pest surveillance. The method provides additional transparency in calculation of the sensitivity of surveillance programs compared to traditional methods because the components of the surveillance system are described separately and not treated as a single test. This methodology also allows for the comparison between programs by comparing the component sensitivity (*CSe*) of the surveillance system components being analysed. It also provides a greater level of transparency for demonstrating freedom as each sample can be attributed to an area, targeting of high-risk subgroups in the population is clearly identified and the increased confidence from targeting high-risk subgroups is accounted for.

During this case study some issues become evident in sourcing appropriate data to include in the model for; 1) setting design prevalences, 2) determining relative risks, and 3) the sensitivity of diagnostic tests. These issues could all be resolved through further research specific to these plant biosecurity problems. Further research to resolve some of these issues would likely include; 1) diagnostic validation of the tests being used in the surveillance programs including calculation of diagnostic sensitivity and specificity, 2) research into the relative risks of the pest being present under different climatic conditions, and 3) observation of the prevalence of disease in established populations and when new introductions occur.

Another difference noted between the applications of this method to animal health surveillance and plant pest surveillance is that with plant pests we are often looking for very small numbers of propagules in large volumes, for example as in the Karnal bunt case study. Also, plant pests are also commonly not associated with a particular host and can be polyphagous, and this adds to the complexity of the models.

The major issue noted with the application of this method in the Karnal bunt example was that prevalence is often considered in terms of infected grains for seed borne diseases but the diagnostic test applied tests for individual propagules of the pathogen. This required additional modelling to convert the design prevalence from a grain volume as a unit to a measure of the number of teliospores per unit grain to allow for the sensitivity of the diagnostic test to be incorporated in to the scenario tree.

This particular example also raised the issue of the complex sampling procedures used when surveying for plant pests, where numerous small lots are often collected from a very large volume of grain, combined, and then divided until an appropriate sample size for laboratory testing is obtained. The incorporation of this type of sampling also raises the issue of clustering of the pest within the grain lot and what effects this may have on the probability that a laboratory sample may contain the pest. This issue has been addressed in sampling for other factors such as genetically modified organisms (GMO), where it has been found that a negative binomial distribution provides a better representation of the distribution of GMO seed in the lot than the Poisson distribution (Paoletti *et al.* 2003). Further study in this area is required to determine what effect clustering may have on the detection of plant pests in grain lots and the demonstration of pest freedom.

This is not the only surveillance activity that could provide data for demonstrating freedom from Karnal bunt in Western Australia. Other activities such as seed health testing and reports from the growers and other members of the public involved in grain handling would provide additional confidence and coverage for grain growing areas of Western Australia not covered by the grain surveillance, for example seed grown for on-farm use. The additional confidence provided by these surveillance activities may allow the number of samples tested to be reduced, potentially decreasing the cost of active grain surveillance.

This case study demonstrates that the scenario tree based methodology published by Martin *et al* (2007a) is suitable for the evaluation of plant pest surveillance and for determining the confidence in pest freedom. This methodology provides another tool, which can be applied for demonstration of pest free areas for trade in plant products. The method is transparent with each step in the surveillance system and influencing risk factors clearly described by the scenario tree. Probabilistic scenario tree analysis, the basis of this method, is used in many fields and is an accepted method for quantitative import risk analysis in field of plant biosecurity.

3.1.2.   Red imported fire ant

3.1.2.1.   Background

RIFA (*Solenopsis invicta*) is one of over 280 members of the widespread genus *Solenopsis*. Although the red imported fire ant is native to South America, it is best known in the United States, Australia, Taiwan, Philippines, and the southern Chinese province of Guangdong. In January 2005, several ant-hills belonging to fire ants were found in northern Hong Kong. The first detection in Australia of RIFA occurred in Brisbane in 2001 and triggered a national cost-shared eradication program led by the Queensland Department of Primary Industries. More information about the ant and its management in Queensland is available at http://www.dpi.qld.gov.au/fireants.

Under suitable conditions, RIFA forms 'super colonies' with multiple queens that spread rapidly and develop extensive colonies. They are omnivorous, opportunistic feeders, preying on invertebrates, vertebrates and plants. They destroy seeds, harvest honeydew from specialised invertebrates and also scavenge. Their foraging can affect whole ecosystems by reducing plant populations and through competition with native herbivores and insects for food.

The RIFA case study arose through a surveillance program administered by Office of the Chief Plant Protection Officer (OCPPO) [1]. Through the 2004 *Securing the Future* budget initiative, the Australian Government granted funds to OCPPO to assist state and territory government agencies implement pre-emptive surveillance for plant pests [2] at high-risk sites in urban areas (hereafter referred to as hazard site surveillance or urban surveillance). The rationale behind this program was that a small amount expended on surveillance for exotic pests that are not known to occur in the country will greatly benefit Australia if new incursions can be detected and controlled at early stages at relatively low cost. All states include RIFA as one of their targets for urban surveillance. A range of different surveillance activities are undertaken as part of the urban surveillance program for RIFA, including: inspection of high risk sites, provision of plant pest hotlines, targeted telephone surveys, trapping campaigns and other surveys.

3.1.2.2.   Scope of analysis

In this case study we analysed a single SSC; periodic inspection of high risk sites. The surveillance period was one year. and the surveillance unit was an inspection site. Data for each city were analysed separately. For each city, the "reference population" of units was the land comprising the city itself.

3.1.2.3.   Data

Site inspection data were obtained from the Northern Territory, Western Australia, New South Wales and Queensland, for the cities of Darwin, Perth, Sydney and Brisbane. Since RIFA is known to be present in Brisbane, "freedom" from this pest is not a possibility at present, so analysis of surveillance data to give a level of confidence in freedom is irrelevant.

---

[1] OCPPO is a branch of the Product Integrity Animal and Plant Health Division of the Australian Government Department of Agriculture, Fisheries and Forestry.

[2] In this context, 'plant pests' means 'pests and diseases of plants'.

However, it is still of interest to look at the sensitivity of the SSC as applied in Brisbane, in order to assess the efficacy of this surveillance activity in the city.

A questionnaire (see Appendix 9-2) was sent to each jurisdiction responsible for urban surveillance (governments of NT, WA, NSW, VIC, QLD, TAS and SA) aimed at acquiring data on inspections conducted, relevant details of sites inspected, and relevant information on the city (see Appendix 9-2 and the following description of the analytical model for details).

3.1.2.4. Model of SSC

Factors affecting the probability that a site would be infested are those affecting the probability of introduction of RIFA to the site (PATHWAY), and those affecting the likelihood of establishment on the site (HABITAT). These two factors were represented in the scenario tree by risk category nodes (Figure 2). In the absence of quantitative information defining these factors, each was divided into categories with a relative risk of infestation in the set: 1 (*very low*), 3 (*low*), 10 (*moderate*), 30 (*high*), or 100 (*very high*). A site in PATHWAY category *moderate* would therefore be 10 times as likely to be infested as a site in PATHWAY category *very low*. These relative risks of infestation were used to apply appropriate weighting to the value of inspections conducted at sites with different risks of being infested, while ensuring that the analysis was conducted for a given design prevalence.
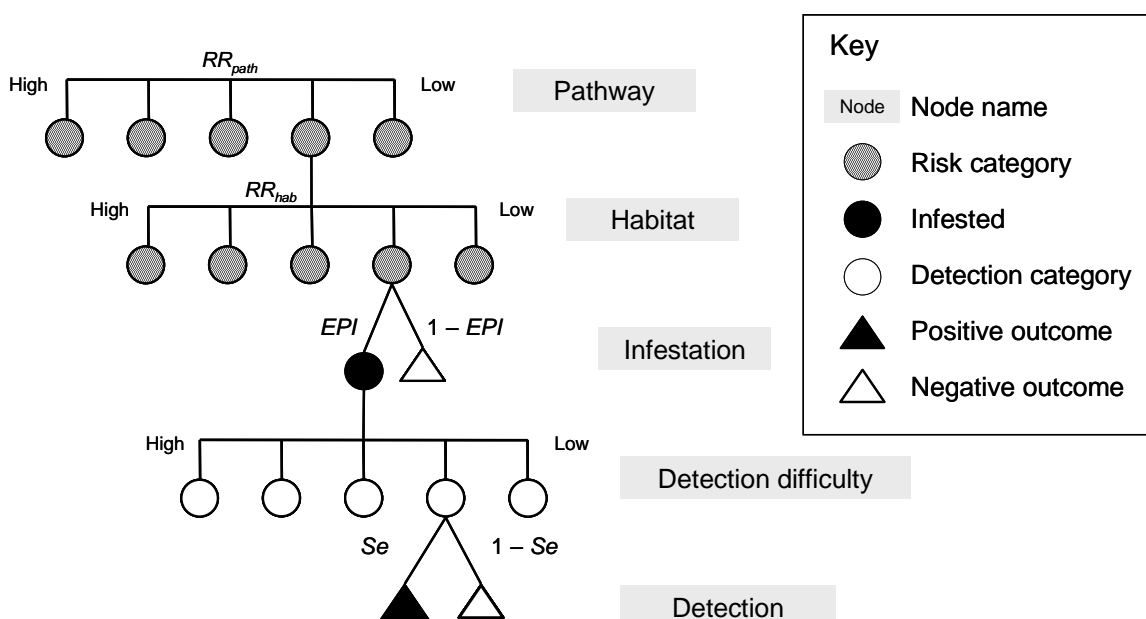
Figure 2   **Scenario tree diagram: urban surveillance for RIFA describing inspection of high risk sites in the program. Only one limb is completed; all branches of each category node within each city have the same structure.**

The probability of detecting RIFA when inspecting an infected site was modelled using a search function based on the *detection difficulty* of the site and the *effort* expended in searching the site. Detection difficulty was represented in the model by a detection category node with five branches/levels (1 to 5), and effort was modelled as a function of time spent on the inspection, the number of inspectors, and the area inspected. Then

$$SeU = 1 - e^{-\frac{t}{e^{\alpha + \beta D}}}$$

where SeU is the probability of detection; t is the effort expended at the site; D is the detection difficulty; α and β are constants estimated from the data.

Time spent on inspection, area of site, number of inspectors and detection difficulty for the site were all data provided by each jurisdiction.

3.1.2.5.   Model implementation

All calculations and simulations were performed using the R statistical language.

3.1.2.6.   Results

Records of inspections were received from NT, WA, NSW and QLD. Over a time period of one year, 120 sites were inspected in Perth, 151 in Darwin, 1,779 in Sydney, and 96 in Brisbane. With a design prevalence of one nest per thousand hectares, the probability of detecting one or more RIFA nests in each city given the surveillance actually carried out (Se_{act}) was estimated to be 0.615 in Perth, 0.412 in Darwin, 0.23 in Sydney and 0.211 in Brisbane (Table 3). This table also gives Se_{act} values for P^*_{Ha} of one nest per ten thousand Ha (*ie* 0.0001), together with estimated values for the hypothetical sensitivity derived from equivalent numbers of random (representative) inspections in each city (Se_{rep}) and values for the sensitivity ratio (*Se* ratio).

$$Se \text{ ratio} = Se_{act} / Se_{rep} \tag{1}$$

This sensitivity ratio is a measure of the relative efficacies of the two sampling designs: *targetted* (as actually carried out), and *representative* (random sampling).

Table 3   Surveillance component sensitivities
for urban surveillance site inspections for RIFA in Perth, Brisbane, Darwin and Sydney

City	K [a]	P^*_{Ha} [b]	Se_{act} (%) [c]	Se_{rep} (%) [d]	Ratio [e]
Brisbane	96	0.001	21.1	9.46	2.24
Brisbane	96	0.0001	2.34	0.985	2.37
Darwin	151	0.001	41.2	6.05	6.81
Darwin	151	0.0001	5.14	0.621	8.28
Perth	120	0.001	61.5	4.04	15.2
Perth	120	0.0001	9.02	0.412	21.9
Sydney	1779	0.001	23	28.7	0.803
Sydney	1779	0.0001	2.58	3.32	0.778

Notes: a: Number of inspections
b: Design prevalence expressed as the expected number of nests per hectare
c: Surveillance sensitivity—probability of detecting one or more times after K targeted inspections.
d: Surveillance sensitivity—probability of detecting one or more times after K representative inspections (10,000 simulations).
e: Surveillance sensitivity ratio, Se_{act} / Se_{rep} .

3.1.2.7. *Discussion*

In this study substantial issues that needed to be addressed at the start included

- Definition of the surveillance unit and the associated issue of defining the reference population

- determination of the meaning of design prevalence

- assessing the meaning, if any, of the concept of grouping or clustering of surveillance units

- quantification of inspection sensitivity, which was dependent not only on categorical variables, to which the scenario tree concept is well suited, but also on continuous variables including

    o area over which ants might be found
    o area of site inspected
    o number of people inspecting
    o time spent inspecting
    o density of vegetation and accessibility of terrain.

- representation of the model

    o logical framework for discussion and diagrammatic representation
    o appropriate environment for quantitative implementation

The principles involved in each of these are discussed in Section 6 below, while the approach taken in this case study is discussed briefly here.

Surveillance unit
Surveillance in this SSC is conducted on high risk sites. These vary considerably in their nature, in terms of land use, building coverage, size, vegetation, and general accessibility. The pest being sought establishes in the ground, so a volume or area of earth is the substance in which it will establish and be found. These considerations all point to the surveillance unit being an area of land, and we chose to define the unit in terms of the inspection site rather than an area of fixed size. This had a range of implications for determination of the sensitivity of detection, which are discussed under that heading below.

Design prevalence
It is clear that with the surveillance unit being an area, the design prevalence is the proportion of surveillance units infested, so nests per Ha was an obvious choice. We considered it highly unlikely that individual ants would ever be identified prior to the establishment of a nest, so in effect if ants are present a nest must be present. RIFA have the capacity to form large, complex nests with multiple queens as colonies develop, but in the spirit of the urban surveillance program, which aims to detect infestations as early as possible, we defined the design prevalence in terms of (single-queen) nests per Ha. The bigger the nest, the larger the area over which ants might be found, and we therefore kept to the standard of a small nest.

A further consideration was whether to define design prevalence in terms of nests per Ha or nests per site. We decided on nests per Ha because of the enormous variability in sites, and the need for the design prevalence to be applied to the whole of the reference population of units, ie the whole of the city. What is a "site" when it is not being inspected? Is it a suburb; a quarter-acre block; a street; a one-hectare area? How many sites are there in the reference

population? Treating the design prevalence as nests per Ha requires assigning the effective probability of infestation in terms of a standard one-hectare unit, and then applying a probability of detection based on the continuous, site-dependent variable *site area*.

Clustering of infestation in groups of surveillance units
We did not consider there to be any sensible grouping of surveillance units because of likely clustering of infestation. In animal health this is done because of the clear tendency of contagious disease to cluster in herds. RIFA nests will, naturally, tend to cluster around "parent" nests, and will spread radially, but in the context of early detection we considered it appropriate to assume independence of inspection sites in this regard. Having said that, clearly nests tend to establish in high risk areas (with regard to probability of introduction and/or probability of establishment), and these factors are accounted for in our model by using the risk nodes PATHWAY and HABITAT.

Quantification of inspection sensitivity
Given the variation in size of the surveillance unit over a continuous range, probability of detection is necessarily going to be determined by continuous, unit-specific variables. The continuous nature of these variables renders the categorical scenario tree inappropriate for representing each of the factors contributing to probability of detection, so this was represented in the tree by a single DETECTION node, whose branch probability was then calculated on a case-by-case basis.

Consistency of data and comparability of results
Information on proportions of each city falling into each of the branches/categories of the PATHWAY and HABITAT risk nodes was estimated by State government respondents to the questionnaire sent out at the start of the case study (**Error! Reference source not found.**.1). Understanding of the meanings of the 5 levels of each of these risk factors is subjective, and no attempt was made to ensure consistency across jurisdictions, so the quantitative estimates of the effective probability of infestation for each unit processed must be interpreted as being based on subjectively derived estimates, and not readily comparable among cities. Similarly, the *detectability* level of each site was estimated (subjectively) by each respondent (with input from site inspectors) so the resulting detection sensitivities for each city will not be directly comparable.

"Data" on all three of these factors (Pathway, Habitat and detectability) were necessarily derived retrospectively, in response to the questionnaire, since they were not recorded at the time of inspection. In some cases the inspectors who had actually conducted the inspections were not available to the questionnaire respondents at the time of the retrospective assessment of the three pieces of information.

4. Software

Software for implementation of stochastic scenario tree models of animal health surveillance system components was written in 2005 as part of a project of the Australian Biosecurity Cooperative Research Centre, and has been available since then over the web, at www.ausvet.com.au/freedom. This software has been rewritten to enhance its capability in line with advances in the methodology (as applied in animal health) concurrently with this ACERA project. This project has incorporated adaptation of the software to cater for applications appropriate to plant disease and invasive pest surveillance. At the outset of this project there were no known, substantial differences between animal health applications and plant health/invasive pest applications, so the adaptations envisaged were principally in terms of terminology and the relevance of the user interface to plant health and invasive pest users. These assumptions remain the case at the time of writing, since no clear additional or different needs for analyses in these "new" domains have arisen during the course of the project.
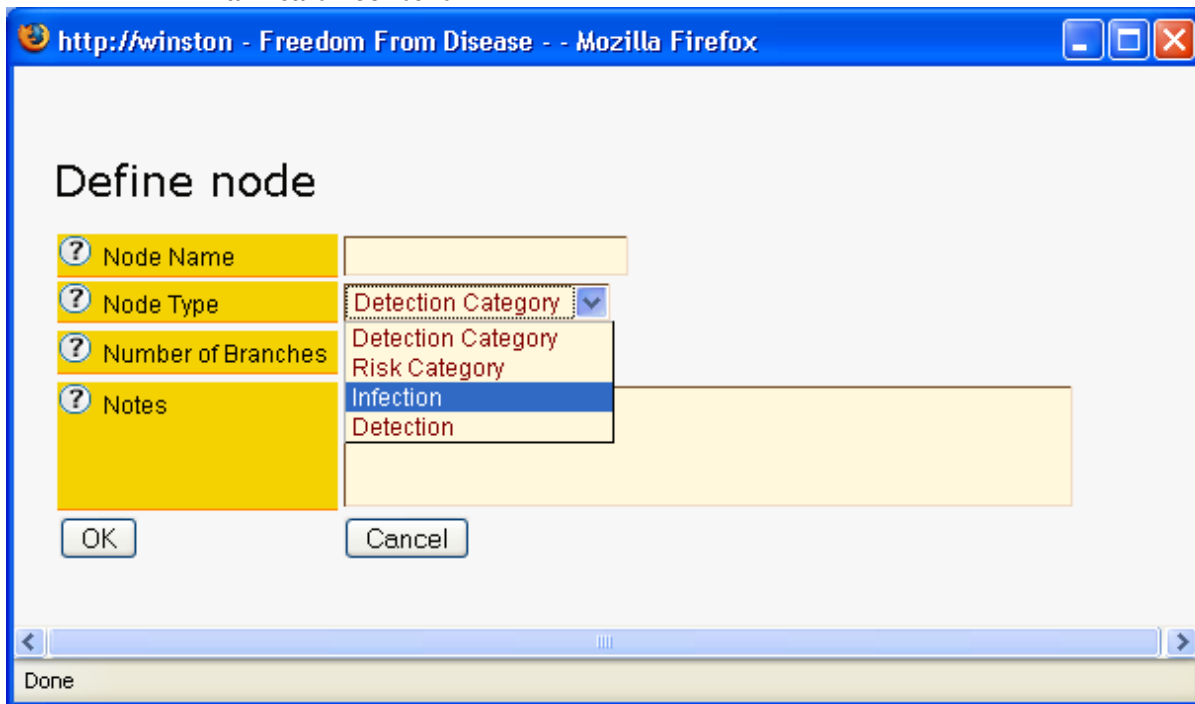
We have created a user interface which allows the user to specify the domain of the model to be created, and then applies appropriate terminology to subsequent steps in the model specification and analytical process. Sample screens from the new interface are shown here. In Figure 3 the user selects a "context" from the options *Animal health*, *Plant health*, and *Pests* on the opening screen for creating a new scenario tree. Having done this, subsequent

Figure 3 *Freedom* **software: context selection**
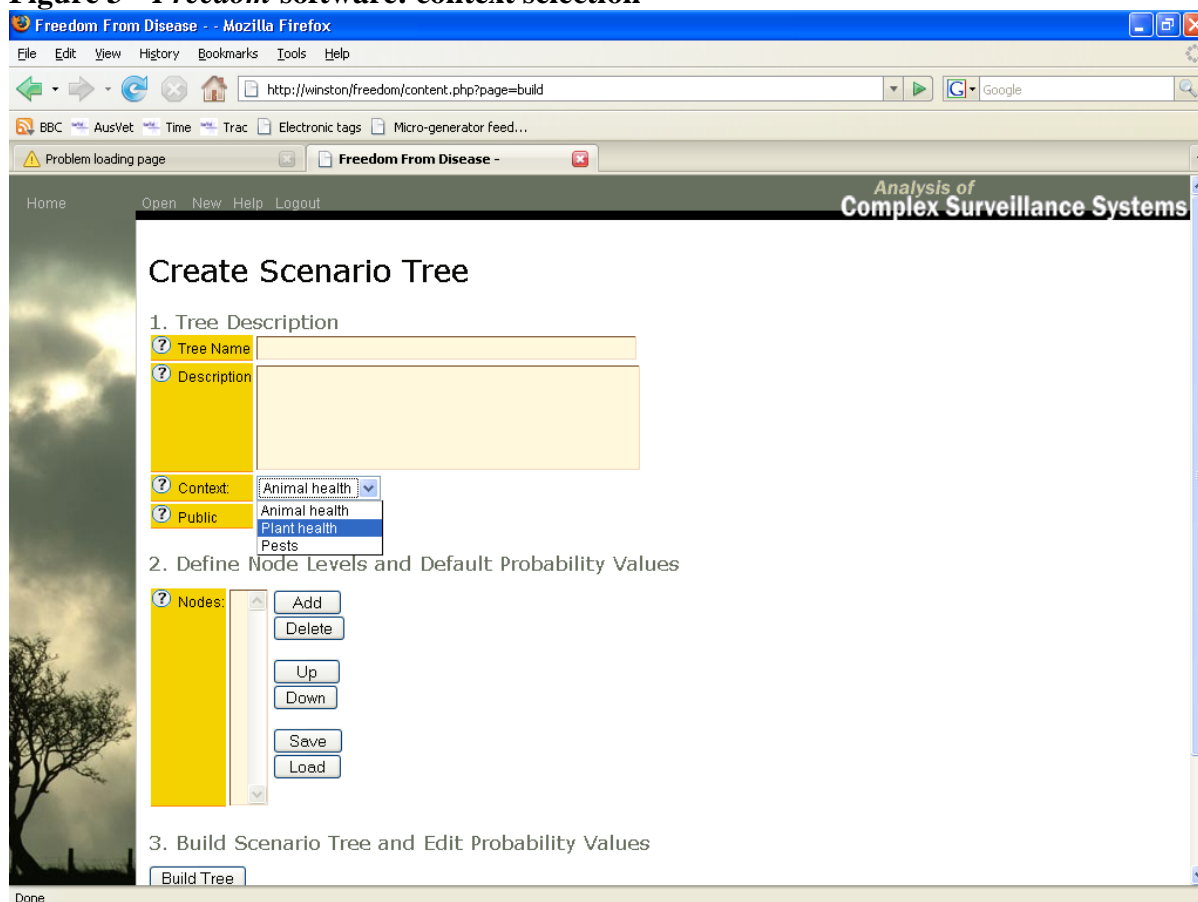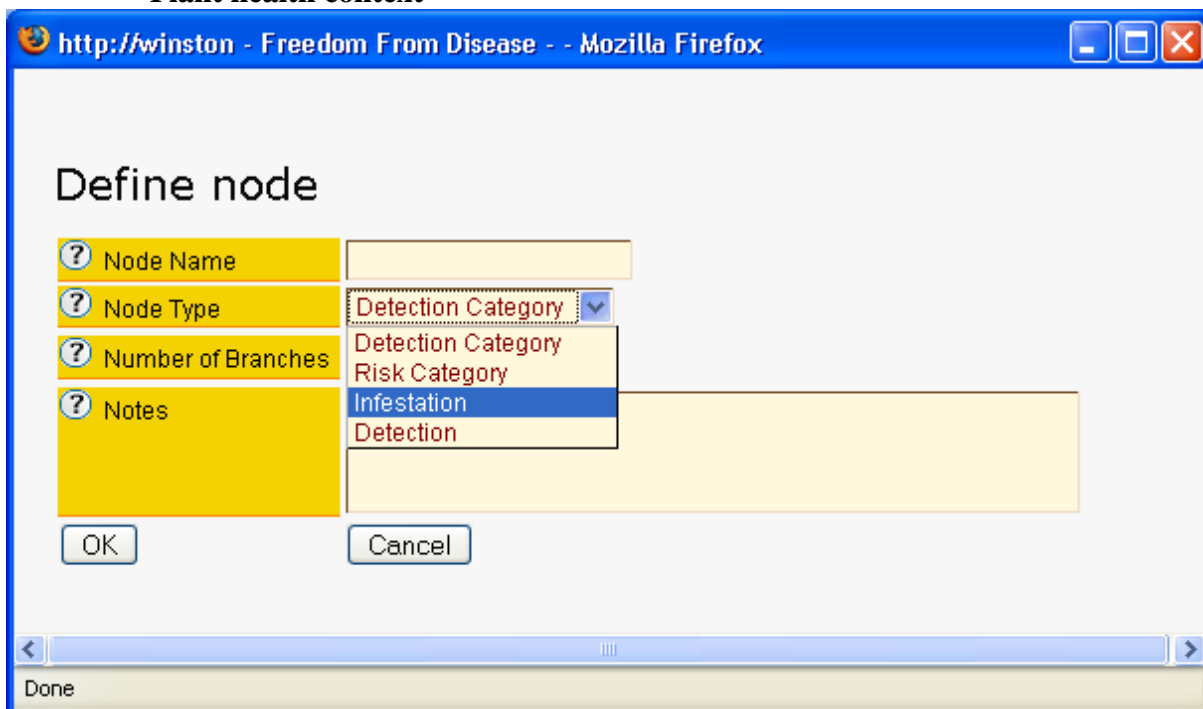
screens give context-relevant alternatives, such as those shown in Figure 4a&b.

Figure 4 *Freedom* **software: node type selection**
          Animal health context

          Plant health context

While this software adaptation is completed and ready for implementation, at the time of writing the major rewrite of the underlying software is still undergoing testing, so is not yet available over the web. It will be mounted on the AusVet website (URL as above), where a comprehensive *Guide to the methodology* is also located. This guide is animal health orientated, and is currently undergoing revision. The revision will incorporate terminological and procedural outcomes from this project.

5. Review of methodology

The methodology forming the backbone of this project (*i.e.* Stochastic scenario tree modelling as described by Martin *et al* (2007a)) was the subject of review within this project. Aspects of the methodology to be considered in the review were

- the methodology as developed for animal health applications
- applications to plant pest and invasive species as highlighted in the case studies of this project

The review team did not attend all project meetings concerned with developing the case studies, and depended on case study reports on which to base the review. Due to the late submission of the RIFA report, this could not be considered in the review, and its assessment of the RIFA case study is therefore limited to the material on that study that was presented at the end-user workshop in Canberra, 31 March 2008.. The review team's report is appended as Appendix 9-3.

The review team of Ray Chambers, Samsul Huda and Liwan Liyanage was led by Ray Chambers.

6. Issues

6.1. Application to "new" domains

Issues encountered in application of the methodology as developed for animal health applications to plant pest and invasive species applications have mostly been raised in other parts of this report, and are summarised or expanded here as follows.

- *Terminology*

  Much terminology is applicable across domains of application, and that which is not has been summarised briefly in section 3 above. Domain-specific terminology has been suggested and incorporated into this report and the web-mounted software at www.ausvet.com.au/freedom.

- *Design prevalence*

  This term is reasonably well established for "the amount of disease we would/will detect if it were present" in the context of both random surveys and general surveillance in the animal health domain. While its meaning is often not immediately apparent in the contexts of plant health and invasive pests, there is no other term which is better suited, so we conclude that we will stick to *design prevalence*. In this project we have encountered what is probably the norm for plant/invasive pests/diseases, but is not easily suited to the concept of a constant design prevalence; namely surveillance units whose magnitude varies on a continuous scale. This issue also arises in animal health surveillance, so it is not a "new" consideration.

  The approach taken in these case studies has been to set a design prevalence in terms of a standard quantity of whatever is the object of the surveillance (grain or terrain), *i.e.* effectively specifying a design *concentration* in the host/medium, and then to calculate unit-by-unit the effective probability of infestation based on the size of the unit as well as the design prevalence and applicable relative risks.

  Another issue in the application of design prevalence is how to deal with small groups of units where there is a grouping node in the scenario tree. In animal health applications this is most frequently encountered as the small herd; in the logical structure of the scenario tree, when a herd is infected we then need to estimate the probability that an individual unit (animal) within the herd is infected. In large herds this is the within-herd or unit-level design prevalence (P^*_U), but in small herds, particularly where $HerdSize \times P^*_U$ is less than one, this makes no sense. You cannot have less than one animal infected in an infected herd: if the herd is infected, then at least one animal must be infected. In small herds one must therefore use a value for P^*_U equivalent to there being one infected animal in an infected herd (an *integer design prevalence*), so the effective probability of an animal being infected within an infected herd is then (*applicable adjusted relative risks*) ÷ *HerdSize*.

  In the context of RIFA, there was no grouping level in the model structure, so the issue did not arise. For Karnal bunt the surveillance unit (the grain sample tested) consisted of 50g of grain, containing approximately 1250 individual grains. For this reason the design prevalence actually concerns the level of infestation present in the unit rather than the

probability that the unit is infested (as opposed to not being infested), and the problem of low design prevalence implying <1 unit infested was not an issue of logical concern. In addition, the assumption was made (and justified) in this study that the grain harvesting and handling process leads to a homogeneous mixing of teliospores throughout the lot. As a result these spores are no longer associated with individual grains, and the level of teliospores is uniform throughout the lot.

In theory, however, if a surveillance unit (*e.g.* a melon) in an infested group (perhaps a 15kg carton if we are dealing with a pest gaining access post-harvest) of not-very-many units is to be assigned a within-group effective probability of being infested, application of an *integer design prevalence* might well be appropriate.

- *Surveillance unit*

  Discussion of appropriate choice of surveillance unit occupied many hours of meetings of the project team. We have not improved on the approach used for animal health applications: "The unit of analysis is that unit for which results are generated by the SSC" (Martin *et al* 2007a; http://www.ausvet.com.au/freedom/pmwiki/pmwiki.php?n=Freedom.UnitOfAnalysis). In animal health there is an obvious tendency towards thinking of such units as being animals or clusters of animals (*e.g.* ponds); in other words that which becomes infected by the disease. However, it is this latter association which can lead to problems in wider application of the concept. RIFA infests a piece of land, and a nest occupies a piece of land of a readily definable size, but the SSC considered in this study involves inspection of sites of varying size, all of them quite different from the area occupied by a RIFA nest. In choice of surveillance unit it is therefore important to focus on that which is being assessed or tested by the SSC rather than that which is infested by the pest. When looking for an insect "hitchhiker" in a shipment of a harvested commodity, the distinction is clear, and it is the inspection unit of the commodity (perhaps the carton) for which results are generated by the SSC. But where the insect is a parasite of the fruit comprising the commodity, the individual fruit will probably be the unit for which inspection results are generated. Where a disease of an agricultural crop is the subject, a defined area of land on which the crop is grown may well be the unit, or alternatively, particularly with trees, it may be the individual tree rather than an area of land.

- *Grouping of surveillance units and clustering of infested units*

  Both case studies generated lively discussion on the issue of whether or not to include one of more "grouping levels" in the logical structure of the model of the SSC. Clearly the organisational structure of the SSC and the logical process of administration of the SSC may well involve levels of grouping of units. For example

  - animals are grouped in herds, which are grouped in administrative districts, which are grouped in larger (perhaps climatically or geographically distinct) regions;
  - fruit trees are grouped in orchards, which are grouped in growing areas, which are grouped in geographically distinct regions.
  - sites for urban surveillance may be grouped in localities and localities within cities
  - sites for urban surveillance may be grouped by habitat and by exposure to potential pest introductions.

Inclusion of these grouping levels in the tree structure and analytical process is generally desirable to enhance transparency and communication of the analysis. However, the only factors and levels that must be included in the analysis are those that affect the probability that a surveillance unit will be infected/infested, and those that affect the probability that an infected/infested unit will be detected. All others may be diagrammatically helpful, but will not affect the calculated sensitivity of the SSC, and when calculating SSC sensitivity these "superfluous" factors and levels (represented by nodes in the scenario tree structure) will simply add to the number and complexity of parameters to be estimated. For example if, in the first example above, we did not include administrative districts, we would need to estimate the proportion of herds processed in each geographic REGION. If, on the other hand, we did include an ADMINISTRATIVE DISTRICT node, we would need to estimate the proportion of herds processed in each administrative district within each region. Assuming we got all these proportions right, the results will be the same either way, but the inclusion of ADMINISTRATIVE DISTRICT may well mean much more, possibly unnecessary, work for the analyst.

The situation in which inclusion of a node representing a grouping level becomes important is where infection/infestation clusters within the groups. With a contagious disease of animals managed in herds, and thus physically separated from animals (units) in other herds (groups), the clustering effect is self-evident, and the probability of an individual unit being infected is different depending on whether its group is infected or not infected. In this situation a group-level infection node is necessary, and it is therefore necessary to have two levels of design prevalence, one representing the probability that a group is infected, and the second the probability that a unit is infected *given that it is in an infected group*. Sensitivity of detection is then estimated at the group level (*SeH*), for each group from which units are processed, and these values for *SeH* are then used as the basis for estimating the SSC sensitivity at the population level.

Where the pest/disease spreads radially from a focus of infestation, regardless of management groups such as orchards or farms (*i.e.* it can just as easily jump the fence as move within the farm), whether or not to include a farm infestation status node is less clear. There are situations in which it is appropriate (where there are risk factors for infestation acting at the farm level) and those where it is not. Development of general guidelines for inclusion of group-level infestation nodes is one area that is clearly in need of further work.

It is important for analysts to understand that specifying multiple clustering levels in a model requires specifying multiple levels of design prevalence, and this renders the whole concept of design prevalence less comprehensible, as well as requiring considerable care in their specification, since ultimately the effective unit-level design prevalence (the average probability that a unit will be infected given that the population is infected at the design prevalence(s)) is the product of all the design prevalences in the model. So, taking the case of the animal disease again, the effective unit-level design prevalence is $P^*_H \times P^*_U$ where there are two infection nodes, HERD STATUS and UNIT STATUS. In this example it may be conceptually and computationally helpful to think in terms of each of the 5 REGIONs of the country being a separate management grouping, and the among-herd, within-region design prevalence P^*_H only applying in an infected REGION. There is then a need for an among-region design prevalence P^*_R with a meaning of the probability that a region will be infected given that the country is infected. Clearly this will have to take a value of (at least) 0.2 in this situation (5 regions) so to maintain the effective unit-level

design prevalence and insert a REGION STATUS infection node we will need to adjust P^*_H and P^*_U such that the same effective unit-level design prevalence is now equal to $P^*_R \times P^*_H \times P^*_U$, perhaps by dividing P^*_H by 0.2.

In general it seems to us that 2 levels of design prevalence (*i.e.* two infection / infestation nodes) is quite enough.

- *incorporation of continuous input variables*

  As discussed above, where the pest or disease infests a host or medium that is most conveniently thought of as a continuum of opportunity for infestation, rather than as a discrete host entity, there will be a need to incorporate a range of continuous variables into estimation of unit-specific probabilities of infestation and detection. This does not present any problems for this methodology, and simply requires that calculation of SSC sensitivity should be based on determining unit-specific *EPI* and *SeU* for each unit processed.

6.2. Methodological issues not addressed in this project

"Combining disparate data sources to demonstrate pest/disease status" implies, among other things, the combination of multiple SSCs into a single estimate of the probability that the population is free from disease. While we intended to combine multiple SSCs, notably in the RIFA case study, we have not done so. This therefore remains untried and undemonstrated.

Much surveillance for exotic pests and diseases relies heavily on the public informing government of observations of pests/diseases. This information may be derived from the observer seeking assistance (including seeking a diagnosis or identification) through government channels (diagnostic laboratories; help lines; etc.) or a conscientious and well-informed member of the public informing "the authorities", or from informal contact between members of the public and members of the biosecurity services. These surveillance activities are grouped as *general surveillance*, and contribute a substantial part of our surveillance evidence for freedom from pests/diseases. Quantifying these non-random sources of evidence is an important and essential contribution to *combining disparate data sources to demonstrate pest/disease status*. In this project we looked only at targetted sampling programs. Associated analysis of general surveillance SSCs for Karnal bunt is being undertaken by Nichole Hammond, and will be completed in 2009. An animal health example has been published (Martin, 2008), and others are in press. There is no clear reason why the methodology successfully applied to general surveillance for animal disease should not be applied equally successfully to plant/invasive pests/diseases.

6.3. Unresolved issues across all potential domains

Issues awaiting attention across all potential domains of application for the methodology include the following.

- How best to incorporate economic considerations into evaluation and comparison of surveillance activities
- How best to make comparisons of surveillance activities. What should the comparison be based on? Marginal cost per 1% increase in *SSe*? Sensitivity ratio?

6.4. Potential methodological developments

Possible new developments expanding the methodology were flagged by the closing workshop of this project in Canberra on 31 March 2008 and are summarised as follows.

- incorporation of spatially-related information into analyses. Climatic, topographical, and habitat-related information are, by definition, spatially distributed, and form essential components of pest- and disease-spread modelling. These and other variables, which may be continuous or discrete in nature, are therefore important influences on both probability of infestation and probability of detection. Various analysts/researchers have approached this need in recent years and the development of accessible and appropriate methods in this area was the most strongly expressed demand at the Canberra forum.

- Analysts would like to be able to select from a series of available analytical templates for the appropriate tool to analyse their own surveillance system. For example:

  - analyses of insect trapping programs will probably require similar analyses across a range of insects;

  - in the urban surveillance program the same surveillance activities are being used to detect a range of different pests.

- There is a need to incorporate modelling of rare and extreme weather events into surveillance planning. How does this fit with the methodology for evaluating efficacy of pest/disease detection?

- Long term trends in the risk of pest/disease establishment are very important in surveillance planning, and it is difficult to obtain long-term funding when one continues to turn up negative results. How can they be incorporated into evaluation of surveillance efficacy?

6.5. Adoption

Issues facing potential adoption and utilisation of this methodology within Australia and in other countries have been canvassed more in the context of animal health than in other contexts, and include the following.

- This methodology is concerned primarily with analysis of varied surveillance data supporting freedom from pests and diseases. This poses two problems immediately:

  - it is attempting to use non-random data, which historically has not been seen as admissible in providing quantitative support for disease freedom, and is seen as "soft" information

  - use of methods and data which are not established and accepted in international trade discussions is not possible unilaterally – the methods and evidence need to be acceptable to trading partners and the wider international community before they are any use for "internal" analyses.

Having said this, all countries are faced with the same problem in demonstration of disease freedom; how to optimise surveillance effort and expenditure, and provide satisfactory evidence both internally and to trading partners which does not involve major expense on an ongoing basis. This is an issue for both exporting countries and importing countries which wish to establish and maintain SPS measures to protect national health status. Within the context of animal health, the general level of awareness of the need for appropriate techniques is clearly demonstrated by the ongoing level of interest in the methodology of this project, with over 300 surveillance analysts and managers from 34 countries having attended courses on the methodology held in 9 different countries, and the interest showing no sign of decreasing. A range of publications are now appearing presenting animal health applications of the methodology.

It should be noted that the reason why this particular methodology is attracting interest is not because it is right, but because a package of some sort has been put together which is accessible as such, and within which or alongside which a range of analytical approaches may be used – anything appropriate in fact. It is an umbrella, and is not exclusive.

- Convincing trading partners that one is free of a pest or disease, whatever means are used to quantify the probability of freedom, is based primarily on accumulation of evidence over time. the methodology used in this project has addressed the quantitative accumulation of evidence for pest/disease freedom, and its application to Karnal bunt surveillance in WA is presented in the attached study report (Appendix 9-1).

7. Recommendations

Feedback from the Canberra workshop and wider sources suggests that there is considerable demand for recognised and accessible procedures for quantifying the efficacy of different, varied surveillance activities for pest/disease detection. The associated estimation of "confidence" in pest-/disease-free status is also something that people want to be able to base on the whole range of surveillance activities undertaken. This demand has led to an ongoing program of communicating the principles of this methodology to animal health surveillance analysts.

Incorporation of spatially defined inputs into models of SSCs is important across all domains of application, but particularly for plant pest/disease and invasive species applications. The concepts of risk density and sensitivity density need to be developed. Interfaces between spatial data and relevant analytical software are needed, providing analysts with straightforward access to processes for linking standard spatial data sets to calculations of SSC sensitivity and probability of freedom. Both the analytical methods and the software need to be developed. A range of people are interested in, or even actively involved in, this work, and their coordination into a focussed project could provide the desired, usable outcomes.

8. References

Hood, G. M. (2008) PopTools version 3.0.5. Available on the internet. URL http://www.cse.csiro.au/poptools

Hueston,W.D., Yoe, C.D., 2000. Estimating the overall power of complex surveillance systems. In: Proceedings of the Ninth International Symposium on Veterinary Epidemiology and Economics, Breckenridge, USA.

Martin PAJ, Cameron AR, Greiner M. (2007a) Demonstrating freedom from disease using multiple complex data sources 1: A new methodology based on scenario trees. *Prev Vet Med* 79; 71-97.

Martin PAJ, Cameron AR, Barfod K, Sergeant ESG, Greiner M. (2007b) Demonstrating freedom from disease using multiple complex data sources 2: Case study—Classical swine fever in Denmark. *Prev Vet Med* 79; 98-115.

WTO, 1995. Agreement on the Application of Sanitary and Phytosanitary Measures. World Trade Organization, In: http://www.wto.org.

9. Appendices

Appendix 9-1   Karnal bunt case study report
click on image of cover page to open report in your pdf reader

CASE STUDY:
SCENARIO TREE ANALYSIS OF GRAIN SURVEILLANCE
FOR KARNAL BUNT IN WESTERN AUSTRALIA

ACERA Project 0703
Combining Disparate Data

Nichole E.B. Hammond,
CRC for National Plant Biosecurity
Murdoch University
Western Australia

Appendix 9-2   Red imported fire ant case study report

click on image of cover page to open report in your pdf reader

Australian Government
Bureau of Rural Sciences

Combining disparate data
to demonstrate pest status:
Pre-emptive surveillance
for fire ants in urban areas

Stephen J. Pratt & Greg M. Hood

April 2009

Appendix 9-3   Review of methodology

Scenario Tree Modelling of Disease Surveillance – Some Comments and Suggestions

Ray Chambers
Centre for Statistical and Survey Methodology
University of Wollongong

24 July 2008

Overview

It is clear that being able to convincingly establish freedom from both disease and pest infestation is fundamental to both Australia's bio-security as well as to its access to international markets for its agricultural outputs. In what follows we therefore use 'disease' as a generic term to describe the presence of biological pathogens as well as insect and plant pest infestation. We also note that a number of surveillance systems are currently operational within Australia with the aim of providing fast and reliable identification of disease presence (and consequently enabling efficient eradication if the disease is found) in a number of key biological and agricultural systems. These systems are typically quite complex, involving the integration of information about disease presence from a number of different sources, ranging from ongoing scientific inspection schemes through point in time surveys of the 'at risk' population to opportunistic data capture as a consequence of unrelated interactions with this population. The question then arises: If none of these different, but potentially overlapping, data collection activities record the presence of disease in the 'at risk' population, how sure can one be that the disease is actually absent? In particular, how can one then estimate the probability that this population is free from disease?

The aim of ACERA project 0703 is to demonstrate how a particular type of statistical modelling of data obtained from a surveillance system can be used to provide an estimate of such a probability. The model used is hierarchical and is referred to as a scenario tree, and the objective of the project is to demonstrate how models of this type can be used to integrate information from disparate components of a surveillance scheme with the aim of estimating the probability that the disease is present *given that the surveillance system has not found*

evidence for this presence. The scenario tree model does this by allowing estimation of the probability of the surveillance system detecting presence of disease in the population given that its prevalence is low (i.e. it is present but is rare). If the estimated detection probability in this situation is high enough, then one can infer that there is a very low actual probability that the disease is present given that no detection has actually occurred. For a discussion of the use of scenario trees in disease surveillance, see Martin, Cameron and Greiner (2007).

In practice, a scenario tree is a hierarchical model for the operation of a surveillance system within a particular environment. We can think of such a model as being made up a set of states, each one of which corresponds to carrying out a particular sequence of surveillance actions given a particular set of environmental conditions, and such that the set of all these states completely characterises the operation of the surveillance system within the environment of interest (i.e. any combination of environmental conditions and surveillance system outcome can be associated with one and only one of these states). For example, in the Karnal Bunt application that is part of this project (see later) one such state is defined by a positive diagnostic outcome for a test sample taken from a siding wheat sample obtained from the Kwinana West wheat growing region in a particular year. Here the surveillance actions correspond to the taking of samples and testing for Karnal Bunt and the environmental conditions correspond to the year of production (and associated climatic characteristics), place of origin of the wheat and its type of delivery/storage.

A scenario tree model for a surveillance system describes the relationship between three types of variables. The first is the underlying prevalence of the disease in the population subject to surveillance. In general, this is a non-negative variable. However, in most scenario tree models it is dichotomised to the variable D with states D_+ = 'disease present' and D_- = 'disease absent'. Second, there is the variable indicating the outcome of the surveillance. Again, this is often dichotomised to the variable T with states T_+ = 'surveillance shows disease present' and T_- = 'surveillance does not show disease present'. Note that both these dichotomisations are typically for convenience. There is nothing to stop D being multi-category (e.g. D_+ can correspond to different levels of disease prevalence) or presence/absence being defined relative to some acceptable 'background' level of disease prevalence. Similarly, T can have many levels (or even be continuously distributed) depending on the type of outcome of the surveillance process. However in most cases, the

primary interest is in disease presence/absence and so it is usually necessary to define T in terms of a range of outcomes that indicate disease presence, i.e. focus on the dichotomous variable defined above. The ultimate aim is then to estimate $\Pr(T_+|D_+)$ and $\Pr(T_-|D_-)$. Note that an effective surveillance system will have both these probabilities close to one, so that the false negative probability

$$\Pr(T_-|D_+) = 1 - \Pr(T_+|D_+)$$

and the false positive probability

$$\Pr(T_+|D_-) = 1 - \Pr(T_-|D_-)$$

are both close to zero. Note that the scenario tree modelling approach of Martin et al (2007) makes the assumption the surveillance system includes all testing necessary to ensure that there is zero probability of false positives.

Finally, there are the variables that characterise the structure of the scenario tree. These depend on the type of surveillance and the population at risk, with the only proviso being that the tree includes all possible combinations of environmental conditions and surveillance actions and that no combination is repeated anywhere in the tree. That is, the different states of the tree correspond to a partition of the sample space of the joint distribution of these variables. If we let $X_k; k = 1, 2, \mathrm{K}, n$ denote the set of possible states for the surveillance system, then a scenario tree model allows us to calculate $\Pr(T_+|D_+)$ by specifying values for $\Pr(T_+|D_+, X_k)$ and $\Pr(X_k|D_+)$. We can then integrate these components to get

$$\Pr(T_+|D_+) = \sum_{k=1}^{n} \Pr(T_+|D_+, X_k)\Pr(X_k|D_+).$$

Typically one would assume that $\Pr(X_k|D_+) = \Pr(X_k)$, but this does not have to be the case. For example, the mode of surveillance could adapt to the underlying prevalence of the disease. A key implicit assumption is then that

$$\sum_{k=1}^{n} \Pr(X_k|D_+) = 1$$

i.e. the surveillance is comprehensive for the 'at risk' part of the population. Otherwise we must condition analysis on the in-scope population for surveillance activity or assume that $\Pr(D_+) = 0$ for situations not in scope (e.g. export wheat not delivered to a CBH facility and hence not subject to surveillance).

From the above, it is clear that specification of the conditional probabilities $\Pr\left(T_+ \mid D_+, X_k\right)$ and $\Pr\left(X_k \mid D_+\right)$ are crucial to specification of a scenario tree. In this project, these probabilities have been specified stochastically, i.e. they have been defined by first specifying stochastic mechanisms for the different environmental conditions and surveillance actions that together constitute the tree and then simulating the operation of these mechanisms given an underlying disease prevalence in order to obtain these probabilities. Note that there is no requirement that these component mechanisms be independent of one another, although in practise this is often the case. Also, this type of model is rather suited to what might be called 'single-mode' surveillance, where a single surveillance action is carried out for each state in the tree. Many surveillance systems are 'multi-mode', i.e. they involve different types of surveillance activity at different levels in the tree. For example, a surveillance scheme could involve both a scientific sampling scheme that varies in intensity according to different levels of environmental factors and an ad-hoc observation scheme that operates across a combination of these factors (i.e. further 'up' the tree). Specification of the tree structure for this scheme then requires a model for the *joint* operation of the scientific sampling and the observation processes. This can be complicated due to possible interactions.

Comments on the Karnal Bunt study

This study was aimed at constructing and evaluating a scenario tree model for the Karnal Bunt surveillance system in operation in Western Australia. Karnal Bunt is a fungal disease of wheat (including durum wheat and triticale) that could significantly impact on Australian exports of this commodity if it became established. In Western Australia the main surveillance system for Karnal Bunt is based on sampling and testing of wheat deliveries to Co-operative Bulk Handling (CBH) depots in the state as well as wheat stored on these sites. This is a 'single-mode' surveillance system since detection is based entirely on evaluation of samples from these sites. Note that CBH is the main company that handles and stores grain produced in WA prior to export, so it is not unreasonable to assume that this surveillance system essentially covers all wheat exports from WA. However, it should be kept in mind that since the scenario tree model is restricted to CBH activities, it tells us nothing about risk of Karnal Bunt in WA export wheat that does not pass through CBH depots.

These caveats aside, the scenario tree modelling exercise for this application is carefully thought through, with sensible models used for the various stages in the scenario tree and operating parameters for these models justified either by reference to relevant scientific studies, expert opinion or from careful consideration of the sampling processes. Inevitably, there are parts of the scenario tree where subjective assessment is used to decide on the stochastic mechanism, but these are clearly identified and can be modified if necessary. The tree itself is defined in terms of three 'environment' variables (Region, Host and Sample Source) and two 'surveillance' variables (Lot and Test Sample), with disease presence measured in terms of Level of Spores (0, 1, 2-4, 5+) within a test sample, thus allowing for imperfect detection.

If the assumptions implicit in this scenario tree are valid, then it provides a sensible model for operation of the Karnal Bunt surveillance system in WA. Of course, if some of these assumptions do not hold, or are questionable, then the usefulness of this model is reduced. Such questions will no doubt be raised by plant production experts as well as by biologists more familiar with spread of Karnal Bunt. From a statistical perspective, however, there is an interesting technical issue with operation of this scenario tree model that could perhaps justify further investigation. This issue relates to the implicit assumption that mechanisms operating in distinct branches of the tree are independent of one another. Thus, for example, the random variable used to indicate the presence of Karnal Bunt in wheat deliveries to Kwinana West CBH depots is assumed to be independent of the corresponding variable for the CBH depots in the Kwinana East region. Similarly, the variable corresponding to detection of the disease in a sample taken from a truck delivery is assumed to be independent of the variable defined by detection in a sample taken from a siding, even though both detections take place in the same laboratory using the same diagnostic equipment. Such conditional independence assumptions (where the conditioning relates to the different 'path' traced through the tree) are commonly used because they allow easy simulation. However, this does not mean that they are necessarily reflective of reality. For example, given that different wheat growing regions are contiguous (or nearly contiguous) one would expect some association in the presence/absence of Karnal Bunt in different regions due to similar climatic conditions. Whether such associations are important in terms of assessment of risk is another matter, however.

Comments on the RIFA study

As their name implies, Red Imported Fire Ants (RIFA) are an imported pest to Australia. They are established in the Brisbane area and are the subject of surveillance in all parts of urban Australia. This surveillance is state-based and multi-mode, including inspection of high-risk urban areas, phone hotlines, household surveys and trapping campaigns. In theory, therefore, RIFA surveillance represents a good test of the capacity of scenario trees to model a complex surveillance system. Unfortunately the modelling that was carried out as part of the ACERA 0703 project did not attempt this. Instead, a scenario tree was constructed for the most straightforward component of RIFA surveillance, which is the inspection of high-risk urban sites. Since all sites are inspected, this tree is quite straightforward, being composed of three environment variables (City, Pathway, Habitat) and one surveillance variable (Inspection). In addition, just four cities (Darwin, Perth, Sydney and Brisbane) were actually modelled. Probabilistic mechanisms for RIFA spread were defined using information on Pathway and Habitat variation in these cities together with subjective values for relative risks of RIFA infestation. Interestingly, these risks appear to be based on an assumption of (log) additivity of risks for Pathway and Habitat, which seems rather strong. On the other hand detection probabilities are more reasonably defined in terms of the difficulty associated with inspection of the different sites and a measure of the effort expended, using data obtained from the relevant inspection authorities in each city.

Aside from the fact that this scenario tree does not model the full extent of RIFA surveillance, and so cannot be used to measure risks associated with this programme, its construction does throw up a number of statistical issues, which are discussed in the relevant report. One of these, that scenario trees do not allow for easy representation of continuous environmental and surveillance variables, is also worthy of further comment. Given that a scenario tree is essentially a decision tree representation of a system, this statement is true – the tree has to categorise continuous variables in order to define its 'branches', thereby imposing a default hierarchical structure on the surveillance process. However, this does not mean that computation of probabilities associated with a particular node of the tree necessarily has to be restricted just to outcomes leading up to that node. This is a manifestation the conditional independence assumption discussed earlier, and, though convenient, it is not strictly necessary. Unfortunately, there does not seem to be a simple way of representing situations in

a tree structure where different nodes 'interact' (as would be the case where they in fact arise from categorisation of an underlying continuum of risk). This again seems worthy of future research if scenario trees are going to be promoted as a general-purpose tool for modelling surveillance systems.

Can we use scenario trees more objectively?

As noted earlier, the main output from a scenario tree model is an estimate of $\Pr\left(T_+ | D_+\right)$ in terms of a posterior distribution for this conditional probability. However, the question of real interest to regulators is the inverse one - how sure are we about disease free status given that the surveillance system has *not* recorded disease presence? That is, what can we say about $\Pr\left(D_+ | T_-\right)$? From Bayes Theorem we see that

$$\Pr\left(D_+ | T_-\right) = \frac{\Pr\left(T_- | D_+\right)\Pr\left(D_+\right)}{\Pr\left(T_-\right)} = \frac{\Pr\left(T_- | D_+\right)\Pr\left(D_+\right)}{\Pr\left(T_- | D_+\right)\Pr\left(D_+\right) + \Pr\left(T_- | D_-\right)\Pr\left(D_-\right)}.$$

Given that D_+ is true, operation of the scenario tree model gives us the posterior distribution of $\Pr\left(T_- | D_+\right) = 1 - \Pr\left(T_+ | D_+\right)$, while $\Pr\left(T_- | D_-\right)$ is usually very close to 1 (and is exactly one for surveillance systems that include all necessary testing to resolve positive test results into true positives or false positives). Consequently, we only need to quantify $\Pr\left(D_+\right)$ in order generate to a posterior for $\Pr\left(D_+ | T_-\right)$. This is usually done by imposing a (subjectively defined) prior distribution on this marginal probability.

However, for some regulators there could be a preference for taking an alternative, more objective approach where we use a scenario tree model for the surveillance system to directly estimate a distribution for $\Pr\left(D_+ | T_-\right)$ that does not require assumptions about disease prevalence, in the sense of only using the empirical fact that no disease has been detected. Consequently, there seems to be an opportunity for extending the effectiveness of the scenario tree modelling approach by developing such an approach. In this context, a reasonable first stab would be to note that for the surveillance system to not find any evidence of disease presence, every part of it (i.e. every state X_k in the scenario tree) should record a negative

result. Consequently we could generate a distribution of consistent values for $\Pr\left(D_+ | T_-\right)$ from the identity

$$\Pr\left(D_+ | T_-\right) = \sum_{k=1}^{n} \Pr\left(D_+^k | T_-^k\right) \Pr\left(X_k\right)$$

where a superscript of k denotes restriction to state X_k of the scenario tree, by combining draws from the joint distribution of the $\Pr\left(X_k\right)$ - which can be singular if we are only interested in the actual population underpinning the scenario tree – with draws from a noninformative/objective joint distribution for the $\Pr\left(D_+^k | T_-^k\right)$. The key issue therefore is how to specify this 'noninformative/objective' joint distribution.

One possibility is to construct it as the product of uniform distributions on confidence intervals for each of the $\Pr\left(D_+^k | T_-^k\right)$. That is, suppose that for each value of k we can calculate a non-negative value $u_k^{1-\alpha}$ corresponding to the upper limit on a $100(1-\alpha)\%$ confidence interval for $\Pr\left(D_+^k | T_-^k\right)$. Then we can generate values of this conditional probability by making draws from a uniform distribution between zero and $u_k^{1-\alpha}$.

This hybrid approach depends on our ability to calculate the $u_k^{1-\alpha}$. In this context, we note 'The Rule of Three' (or R3): Given no successes are observed in n independent Bernoulli trials, each with common success probability π, then a good approximation to a 95% confidence interval for π is $[0, 3/n]$ (Jovanovic and Levy, 1997). Or, perhaps more appropriately, a better (Bayesian) rule: Given $\pi : Beta(1, b)$, the widest 95% posterior credibility interval for π is $\left[0, 3/(n+1)\right]$.

As stated, R3 holds for *independent and identically distributed* data. In the scenario tree context, our 'negative' observations are not identically distributed. They may not even be independent. There is also the issue whether the confidence coefficient α should be the same for each k or just take this value when state-specific intervals are combined across the whole scenario tree.

This suggests that further research could be carried out on whether we can generalise the R3 approach to determine values of $u_k^{1-\alpha}$ for α 'close' to zero within a scenario tree context. If such values can be determined, then by sampling within these limits and integrating the values thus generated with draws from the distribution of possible values of $\Pr(X_k)$ we could, for a specified value of α, generate an 'objective' distribution of values for the probability $\Pr(D_+|T_-)$ that is the real focus of interest.

Conclusion

Scenario trees represent a potentially useful tool for modelling the stochastic structure of surveillance systems. However, our capacity to use these trees to integrate information obtained from surveillance systems that operate across different 'modes' of surveillance remains untested, since both applications considered in the ACERA 0703 project corresponded to 'single mode' surveillance, where it is clear that the trees can be constructed and used to obtain detection probabilities given specified levels of prevalence of the disease.

Are scenario trees suited to all types of surveillance? Here some doubts have arisen, mainly because these tree-based models implicitly impose a conditional independence structure on detection 'events' in different parts of the tree. There may be some situations where this conditional independence is inappropriate. However, without further research it is impossible to say whether this imposition of conditional independence has the capacity to seriously compromise the performance of the tree in terms of estimating the surveillance system's detection probabilities.

Finally, there is the issue of the objectivity of a scenario tree model. These models, by construction, do not take account of the fact that in the situations of most interest, no evidence for the disease has been found. This 'null' information can be used to determine 'objective' prevalence estimates that can then be translated, via the tree structure, into estimates of detection probabilities that may be more acceptable from a scientific perspective. Again, further research is needed to see if this is feasible.

References

Jovanovic, B.D. and Levy, P.S. (1997). A look at the Rule of Three. *The American Statistician*, 51(2): 137-139.

Martin, P.A.J., Cameron, A.R. and Greiner, M. (2007). Demonstrating freedom from disease using multiple complex data sources: 1: A new methodology based on scenario trees. *Preventive Veterinary Medicine*, 79(2-4): 71-97.