Australian Government
Department of Agriculture and Water Resources

cebra
Centre of Excellence for Biosecurity Risk Analysis

Ministry for Primary Industries
Manatū Ahu Matua

CEBRA Report Cover Page			
Title, ID, & Output #	1502E, Exposure Pathway Model for Forest Surveillance: Stage 2, Output ID 6		
Project Type	Spatial Analysis	CEBRA Project Leader	Tom Kompas
Sponsoring Org.		Project Sponsor	Steve Hathaway, Director Biosecurity Science, Food Science & Risk Assessment, Regulation & Assurance Branch, Ministry for Primary Industries
Project Leader/s	Paul Stevens, Investigation and Diagnostics Centres and Response Directorate, Operations Branch, Ministry for Primary Industries	Collaborator/s	Bayesian Intelligence, Monash
Project Objectives	This project is based on the output of CEBRA project 1402B, and a prototype model generated from the MPI and the NZ Forest Owners Association (NZFOA) and NZ research collaborators. This project aims to quantify levels of pest exposure along pathways into the country in order to facilitate a robust nationally comprehensive surveillance programme for invasive pests that allocates resources on the basis of risk and is defensible internationally. The models (sub-models for each pest and each pathway) and risk maps will enable quantification of relative risk of entry (exposure) and establishment to suggest the highest risk areas for pest invasion and establishment. This will likely allow more effective and efficient allocation of surveillance resources and therefore a higher likelihood of detection in time to eradicate successfully. The completed Bayesian Network model (combining all the sub-models) will also allow investigation of the effects of allocating different amounts of intervention effort in different parts of the pathways investigated, allowing some examination of the trade-off between extra effort prior to the border and extra resourcing for border interventions.		
Outputs	Final Report. Outputs include Bayesian Networks exposure pathway models and biosecurity risk maps, recommendations on highest exposure areas for targeting surveillance.		

CEBRA Workplan Budget	**Year 20XX-20XX**			

Project Changes	N/A
Research Outcomes	The New Zealand forest industry are full steam ahead with using the model to plan their annual surveillance programme. Following a pilot project getting underway this month they are planning rollout of the full system in August of this year. MPI will also be using the model to update risk mapping for its High Risk Site Surveillance programme for the 2017-2018 season. Season planning starts in May with full rollout of fieldwork in September. From this you can see that MPI and the NZ forest industry are very happy with the output of the project and we look forward to any future collaborations.”
Recommendations	
Related Documents	
Report Complete	

Report Tracking:	SRC Reviews:	Date:
	DAWR/MPI Policy / Operational Reviews:	Date:
	BRSC Endorsement ( ) Yes ( ) No	Date:
	Report published to CEBRA Website ( ) Yes ( ) No	Date:

Exposure Pathway Model for Forest Surveillance: Stage 2

Final Report

Version: 3

August 11, 2017

Developed for CEBRA and MPI New Zealand

Prepared by Bayesian Intelligence Pty Ltd
Consultants: Steven Mascaro and Dhananjay Thiruvady
Email: steven.mascaro@bayesian-intelligence.com
dhananjay.thiruvady@bayesian-intelligence.com

ABN: 33 128 355 386
P.O. Box 8025
Monash University LPO
Clayton, Victoria 3168
Email: bi@bayesian-intelligence.com

BAYESIAN INTELLIGENCE

Contents

1 Executive Summary

The Stage 1 prototype exposure model was designed to investigate the value of a Bayesian network (BN) based approach to estimating the probability of pest entry and exposures across New Zealand. It proved that the approach was practical and that it was able to capture our understanding of how pests enter the country. The Stage 2 model builds on the successful earlier model by expanding and improving it in a variety of ways. In addition, custom software (called Spear) has been created to run and manage the model and its associated results.

The Stage 1 model focussed on three main types of pest – Asian gypsy moth (AGM), wood-boring and bark beetles (WBB) and Phytophthora – along with a host of associated items (such as containers, vehicles, sea vessels and passengers). We added two new pest classes – pine shoot moth (PSM) and Fusarium – which is less than the six new pests originally envisioned. After substantial consultation, we concluded that these five pests provide excellent coverage of the *features* of pests that may enter, and hence additional pest models were not needed. This gave us more time for other work, including elicitation, sensitivity analysis and Spear.

Stage 2 includes a Generic pest submodel that can be used to represent any pest at a low level of detail. The Generic pest can be used to handle a range of possibilities including: unknown threats; poorly understood pests and pest classes; new pests that may not fit any of the existing models; or simplified versions of existing pest classes that benefit from a common structure. In addition to the Generic pest submodel, the database and software used to drive the model was re-architected to allow for a clean separation of pests and items, allowing any pest to be combined with any item by merely specifying the appropriate parameters.

We explored the value of incorporating monthly factors into the exposure model. This work, which focussed on the AGM, showed that taking into account monthly factors does add value to the models, and has been extended to other pest submodels where those pests are influenced in important ways by seasonal factors. It also suggests there may be benefits in running the exposure model at a monthly resolution by default for certain pests.

Many improvements have been made to the parameters used in Stage 1. Most of the changes have been driven by elicitation workshops with experts and other individuals that have good familiarity with the domain. A modified version of the Delphi technique was used to elicit values for parameters such as infestation rates, detection rates and a range of pest- and item-specific parameters. While such parameters are always subject to large amounts of uncertainty, most of the parameters in the model are now well supported by data, literature or elicited estimates.

A substantial start has been made on performing a sensitivity analysis of the exposure models. Sensitivity analysis is critical to understanding the robustness of a model and to assessing how uncertainty impacts the conclusions we might draw. While time did not allow for a comprehensive sensitivity analysis of the full system, we have conducted analyses of several of the system's key submodels. The results were insightful and influenced model development. They also highlighted key factors for which good estimates were required (such as the time spent by items and their pests at various locations), and other factors that were less important.

At a broad level, the resulting exposure maps appear reasonable, and match well with the maps produced by the original prototype exposure model. The results for AGM on the used vehicles pathway match very well with the results produced in Stage 1, and this is as expected given the robustness of that model. The maps for Phytophthora fare less well, but this was also expected due to the lack of knowledge surrounding Phytophthora and the changes made to its model.

An important output of the project is the software created to run and manage the models. Many parts of the exposure model will need to change over time. For example, the number of items entering the country will grow or contract from year to year, land use and population densities will shift across the country, and the location of transitional facilities, shops and tourist sites will not always remain the same. For any given year, the model needs to be updated to take account of these changes, and exposure maps will need to be regenerated. The software automates the running of the exposure model, along with all of the tasks needed to generate exposure maps. This allows the model to remain relevant well into the future, without requiring any core changes. Importantly, the software also allows for exploratory runs, with methods for run management built in, and provides tools for aggregating pathways in a way that allows for a rudimentary overall risk analysis.

The following is a summary of the key outcomes of this project:

- The exposure model provides an explicit, transparent, flexible and robust representation of the nature of an entry pathway and its constituent pathway points.
- The model provides good coverage of the key classes of pests that may affect forest health.
- The model allows unknown or poorly understood pests to be modelled with a Generic pest.
- The model can take account of seasonal factors.
- The model and software contains a clean separation of items and pests, and there is now greater flexibility in how pests and items can be combined.
- Model parameters and data can be updated independently to create new exposure maps.
- The model and its software can be deployed as part of a surveillance workflow (through its user interface) and can also be embedded in other software directly.
- The Spear software has been created and is flexible enough to be used for both this and future biosecurity projects.

The following future work is recommended:

- Further improvements to the software, in particular to improve the automation of data updates, as well as to improve the tools available for analysing the model results.
- Extension of the model to handle pests that arrive by wind.
- Further sensitivity analysis work, including a full system sensitivity analysis if feasible.
- Further refinement of the (expert-based) parameters through future elicitation workshops.
- One or more published papers on the exposure model and its software.
- Software integration with the survey optimisation work.
- Extension of the full system (i.e., including exposure, establishment and survey optimisation) into a full risk model, that allows for a complete and proper risk assessment.
- Formal validation of the model, including comparisons against existing systems, such as MPI's High-Risk Surveillance Sites (HRSS) system.

2 Introduction

The New Zealand Forest Owners Association (NZFOA), in partnership with NZ's Ministry for Primary Industries (MPI), is upgrading their current Forest Health Surveillance system (FOA, 2015) to provide better early warning of invasions from a variety of pests that could do harm to New Zealand's plantation forests and, in turn, economy. To maximise the value and effectiveness of the Forest Health Surveillance system (FHS), surveillance components (such as traps, walkthroughs and aerial surveys) need to be deployed where they have the highest chance of identifying invasions early enough to maximize the response options available. In 2014, MPI engaged Bayesian Intelligence to create a prototype exposure model to explore the plausibility of a Bayesian network (BN) based approach to forest surveillance. After a successful demonstration of the technique across several important import pathways and pests, MPI requested that this prototype (now known as the Stage 1 model) be extended to cover a larger, more general range of pests, items and pathways.

An '*item*' in this context refers specifically to anything that a pest can be attached to and that can enter into the country. This includes things like used vehicles and containers, but also less discrete things like wind and soil. For this and other key terms, we strongly recommend the reader consult the brief glossary in Section 8.

The Stage 1 model focused on 3 pests – Asian gypsy moth (AGM), wood-boring and bark beetles (WBB) and Phytophthora. These were combined with 6 items (used vehicles, sea containers, sea vessels, wood packaging, wooden furniture and passengers) to create 7 unique pathways. Excellent data was available for the items, while much less certain information was available for the pests, particularly for infestation rates and certain aspects of pest phenology. The quality of information for detection and treatment rates was moderate to low – these were treated as mostly similar across pests based on rough estimates from experience. The results of the model were encouraging, aligning with expectations on how exposures would be distributed across New Zealand and, in general, this proved the basic approach to be sound.

After the Stage 1 prototype, it was expected that Stage 2 would largely consist of adding further pests, items and pathways to the system, along with creating customised models that catered to the unique features of particular pest-item combinations. However, it became apparent early on that many of the new pests and items to be modelled had features that were strongly similar to pests and items that had already been modelled. Furthermore, it was recognised that focusing too much on the unique traits of a single pest would restrict the generality and flexibility of the model.

The direction of the project was changed to accommodate this new understanding. Instead of the six new pests that were originally slated for inclusion, this was scaled back to just two new pests – Fusarium and Pine Shoot Moth (PSM). As before, modelled pests are intended to cover major classes of pests (e.g. eggs, live pests, moths, beetles, soil-based pests, etc.) rather than specific pests, hence the need to model only a few pests. Furthermore, a Generic pest model has been created. This can act as a model for almost any pest, but only at a very high level of abstraction. In addition, the underlying model (in particular, the database schema for parameters and all the related code) was refactored to create a clean separation between pests and items, allowing arbitrary combinations of existing pests and items, with changes needed in only a few places in the parameters file.[1]

The reduction in the number of pests modelled allowed for enhancements in other areas. More emphasis was placed on the software – now called Spear – taking it from the original idea of a basic tool for running the models, to incorporating run and data management along with some light data analysis. An initial sensitivity to parameters analysis was performed on each of the key submodels, and an initial capability for performing this sensitivity analysis was incorporated into the software. In addition, both in-person and online elicitation sessions were run for several key model parameters, which allowed placeholders and idealised parameter values to be replaced with much better estimates. The elicitation sessions also led to major revisions and improvements to the Phytophthora model, which was (and still is) known to be a highly speculative and uncertain model. While a more formal validation of the model was considered outside the scope of this project, the model structure and relationships were examined and accepted by experts (from MPI, Scion and the forestry industry) at a one-day workshop in Rotorua in November of 2015 and MPI is currently assessing the model against its existing systems.

This report deals with the changes that have been introduced in the Stage 2 model, and presents some of the exposure maps and other results produced by the model. Note that the model is constantly evolving and improving, so the presentation here is very much a snapshot of how the model and its results look at the time of writing. In Section 4, we present an overview of the elicitation sessions that took place during the project, with the aim of providing transparency on this approach to determining parameter values. We also present the results of an early sensitivity analysis on some of the submodels within the system, highlighting some of the parameters that affect the results the most, as well as identifying some of the parameter ranges for which the model can still be considered sound.

2.1 Related Documents

For an explanation of the general approach along with a detailed description of the Stage 1 model, please see the "*Forest Health Surveillance Exposure Mapping Prototype*" technical report (Bayesian Intelligence, 2015). For more information about the architecture of the model and the software, see "*Spear: Technical documentation*". For information on how to use the software created for this project, see the user guide "*Spear: Exposure modelling software for invasive forest pests*". (Both the technical documentation and user guide are included in the Spear software package.)

[1] A windborne model was to be created as well, but proved to be more challenging based on the state of other available software. The windborne model was therefore omitted from the software, but a proposal has been developed for including it in future. See Section 3.4.3 for more details.

3 Approach

Here we look at the approach taken to creating and extending the models in Stage 2. We first provide a brief overview of the model in Section 3.1 (for full details, please consult the Stage 1 technical report), and, in subsequent sections, focus on changes that have been introduced in Stage 2. Note that an updated data dictionary for all the variables across all models (old and new) has been provided in Section 9, which also highlights the changes made in Stage 2.

The model changes were motivated by MPI's requests (introducing month as a factor and including pests which are seen as threats), experience from the previous stages of modelling (e.g. splitting items and pests), and considerations of generality (e.g. being able to handle potential future threats). The changes essentially make the models more robust, generic, and allow future changes to be incorporated more easily. The parameters were obtained in several ways including from data, literature and through elicitation. In some cases (such as used vehicles numbers), data and literature provide reliable parameter values. However, there are a number of parameters for which accurate values were not obtained easily and for these cases, elicitations were carried out. Since this was a significant task in itself, the elicitation is described separately in Section 4.

3.1 Generic Model Structure

The prototype makes exclusive use of BNs for the modelling of the system.[2] There is no single BN or set of BNs that describes the whole system. Rather, BN models are generated from a script coupled with a configuration spreadsheet describing how the pieces fit together for each exposure pathway. The configuration spreadsheet also contains parameters that are to be entered into the generated networks, or in cases where the number of parameters required is very large (in particular, for spatial data), pointers to other files that contain the parameters. This approach is useful because the system being modelled is large and there are many shared components across the system that can be described using common, generalised submodels. The approach is also highly scalable, and adding or changing pathways does not require any changes to the design of the system.

While models are generated dynamically for each exposure pathway, they nonetheless share a common overall structure, shown in Figure 1. Each of the boxes represents a submodel, and submodels that share a colour also share the same internal structure. Every model has a flow that matches the structure of the pathway. This begins with Pest Arrival, which models how many pests will arrive at the first point on the pathway (typically, the sea or air port). Once the distribution over this quantity is determined, we model the Pest Activity at the entry point (i.e. how many pests will die or escape and how they will develop while at the entry point). We also model the effectiveness of any Treatments that may be applied to items at the entry point, which may include intentional and unintentional impacts on the pests. After this, we carry over the distribution for the remaining quantity of pests to the next point in the pathway for the item,

[2] However, as designed, it would be quite easy to "plugin" other kinds of model as sub-components, which may be convenient when working with existing models.

modelling Pest Activity and Treatment again in the same way, and we continue on in this way until we reach the end point of the pathway.[3]

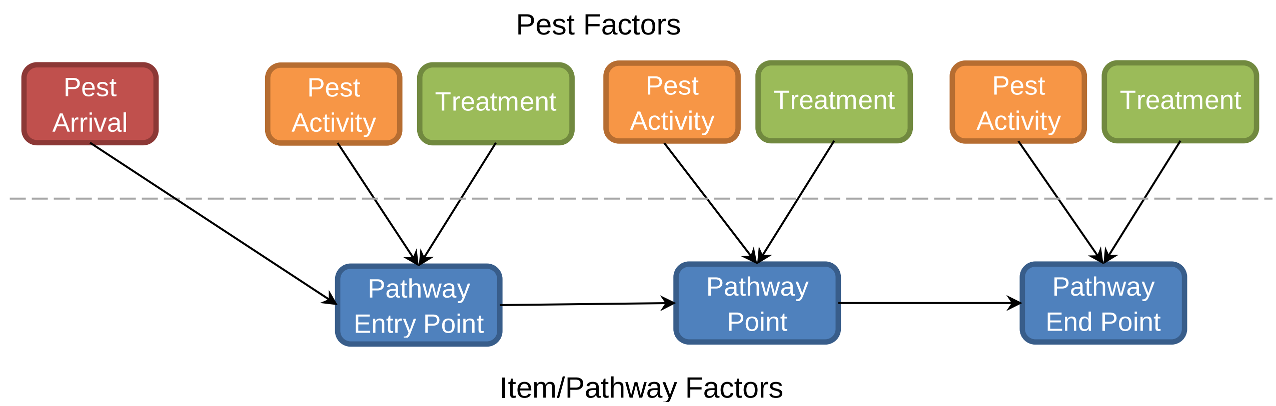

Figure 1: Common structure for all exposure pathways

As indicated in the figure, there are two key sets of factors in the exposure models: a) pest factors and b) pathway factors. Pest factors refer to the properties of the pest itself, such as its mortality or reproduction rate, as well as how it is affected by its environment and human interventions. Pathway factors refer to how the pests move from point to point, which is often determined exclusively by the movement of the item on which the pests are travelling.

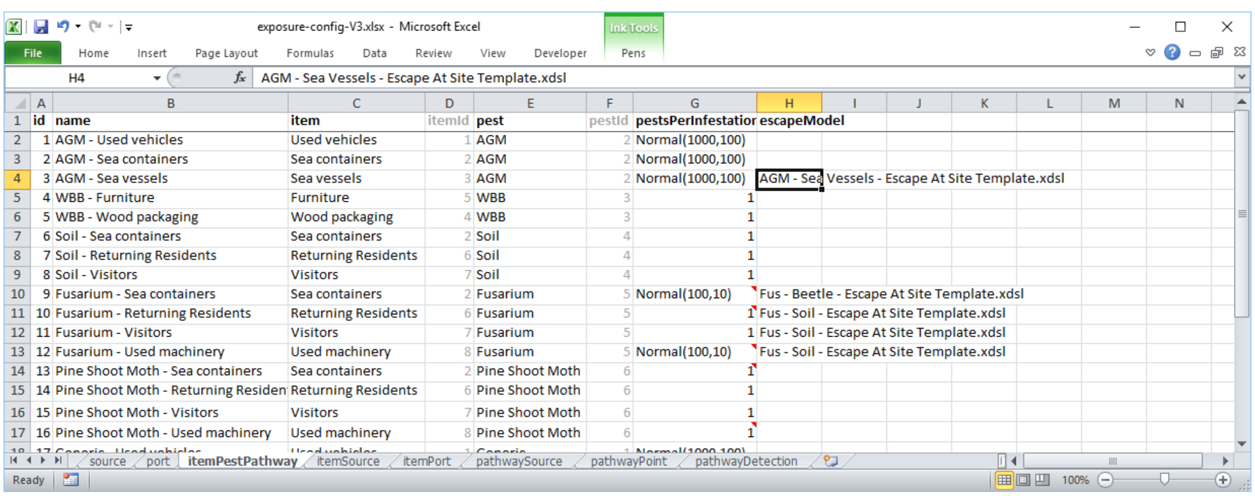

id	name	item	itemId	pest	pestId	pestsPerInfestatior	escapeModel
1	AGM - Used vehicles	Used vehicles	1	AGM	2	Normal(1000,100)	
2	AGM - Sea containers	Sea containers	2	AGM	2	Normal(1000,100)	
3	AGM - Sea vessels	Sea vessels	3	AGM	2	Normal(1000,100)	AGM - Sea Vessels - Escape At Site Template.xdsl
4	WBB - Furniture	Furniture	5	WBB	3	1	
5	WBB - Wood packaging	Wood packaging	4	WBB	3	1	
6	Soil - Sea containers	Sea containers	2	Soil	4	1	
7	Soil - Returning Residents	Returning Residents	6	Soil	4	1	
8	Soil - Visitors	Visitors	7	Soil	4	1	
9	Fusarium - Sea containers	Sea containers	2	Fusarium	5	Normal(100,10)	Fus - Beetle - Escape At Site Template.xdsl
10	Fusarium - Returning Residents	Returning Residents	6	Fusarium	5	1	Fus - Soil - Escape At Site Template.xdsl
11	Fusarium - Visitors	Visitors	7	Fusarium	5	1	Fus - Soil - Escape At Site Template.xdsl
12	Fusarium - Used machinery	Used machinery	8	Fusarium	5	Normal(100,10)	Fus - Soil - Escape At Site Template.xdsl
13	Pine Shoot Moth - Sea containers	Sea containers	2	Pine Shoot Moth	6	1	
14	Pine Shoot Moth - Returning Residen	Returning Residents	6	Pine Shoot Moth	6	1	
15	Pine Shoot Moth - Visitors	Visitors	7	Pine Shoot Moth	6	1	
16	Pine Shoot Moth - Used machinery	Used machinery	8	Pine Shoot Moth	6	1	

Figure 2: An example view of the configuration spreadsheet, showing the 'itemPestPathway' sheet. See the Spear user guide for more details

The model is, in fact, a collection of BN models, submodels and scenarios tied together by the Spear software via a single configuration spreadsheet (see Figure 2 for an example). The spreadsheet is organised in the style of a set of database tables, and allows users to specify

[3] Travel between points was not considered significant enough to include in the model, however the approach can easily accommodate this by adding points to the pathway that explicitly represent travel.

submodels, parameters and external data (including GIS layers and external tables) that are needed by the model. Each run of the model stores a snapshot of the configuration spreadsheet so that models can always be re-run. For more details on the configuration spreadsheet, please see the Spear user guide.

We now describe the main submodels.

3.1.1 The Pest Arrival Submodel

The main purpose of the Pest Arrival submodel is to describe (the distribution over) the quantity of pests that arrive at the entry point for a pathway. The structure of this model is shown in Figure 3, with the node of interest 'PestQuantity' at the bottom, which indicates how many pests ultimately arrive at the entry point. Since we are focused on what occurs after arrival, this submodel has been kept very simple, containing only those factors that are specifiable from easily available data.

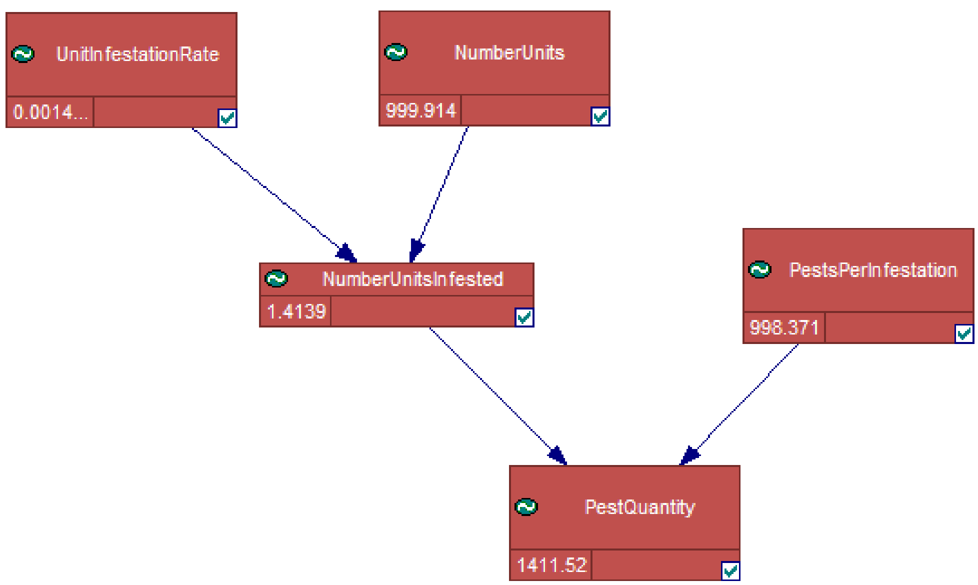

Figure 3: The Pest Arrival Submodel template. The numbers in the boxes are expected values. The formulas used for this model are described in Error! Reference source not found.**. Note that the parameters that produced these expected values are examples only (from the AGM on used vehicles pathway), and are set for each pathway from the configuration spreadsheet. (See Section 9 for variable definitions.)**

3.1.2 The Treatment Submodel

The Treatment submodel identifies the proportion of pests that are successfully treated (that is, successfully eradicated) at any point in the pathway (Figure 4). Treatment may be traditional treatment (such as fumigation) or it may also be an incidental human intervention that leads to the death of the pest (for example, washing the containers at the cleaning depots). These nodes operate with population proportions, rather than with individual pests. This is so as to keep the submodel units consistent with those for the Pathway Point submodels, allowing a simplified, inline treatment of this submodel. Indeed, this submodel appears in the Pathway Point submodel in the implemented system. The output of this submodel is the ProportionTreated node, which indicates the proportion of pests successfully treated.

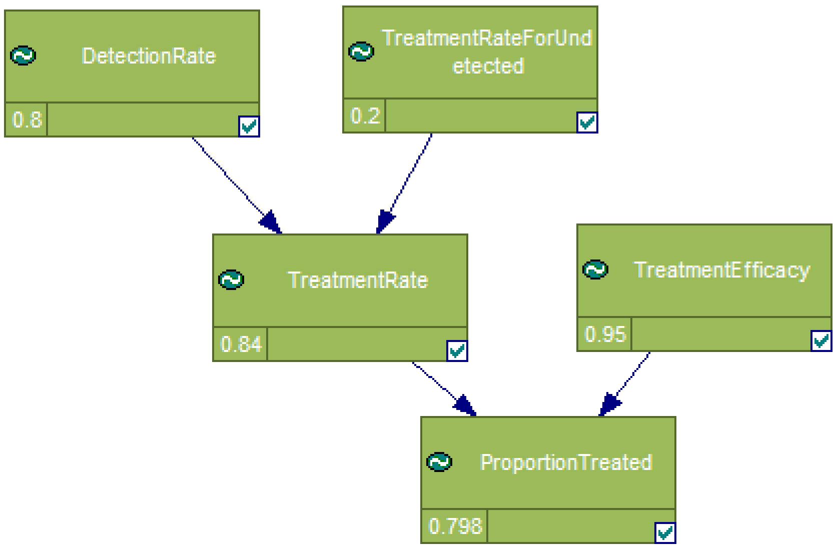

Figure 4: The Treatment submodel template. (See Section 9 for variable definitions.)

3.1.3 The Pest Activity Submodels

The Pest Activity submodels (sometimes called Escape submodels) describe how pests develop and the factors that affect whether or not they escape the item to which they're attached. In contrast to the other models described here, it can be modelled at the level of an individual pest *or* a pest population. Individual pest modelling allows for a much more natural causal description of the factors that affect a pest's development, behaviour and mortality in many cases. However, it has the implication that the nodes cannot simply be linked to other submodels, nor can node distributions be directly copied across (but they can be treated as parameters in the other submodels).

The Pest Activity submodels are highly dependent on the pest in question (and in some cases the pathway), and so each pest has its own submodel. However, every Pest Activity submodel has a set of inputs and outputs in common, which include TimeAtSite (input), EscapesAtSite and Survives (both outputs). (For population-based models, the two equivalent outputs are instead ProportionThatEscapeAtSite and ProportionSurvive.) TimeAtSite specifies how long the pests remain at a particular site in days; EscapesAtSite gives the probability that a pest escapes at that site; and Survives gives the probability that a pest survives to move on to the next pathway point.

3.1.4 The Pathway Point Submodel

The Pathway Point submodel is the heart of the system, tying together the outputs from all the other submodels. This submodel is designed to click together into arbitrarily long chains, similar to a dynamic Bayesian network (DBN).[4] Conceptually, this works as indicated previously in Figure 1. For example, in the case of AGM on used vehicles, we will have four copies of the Pathway Point submodel connected end-to-end, with the port-specific model first, then registration sites, car yards and finally the end point somewhere in the country. The structure remains the same in all cases, with only the parameters differing. Figure 5 shows the structure of the Pathway Point submodel, and **Error! Reference source not found.** shows the node definitions.

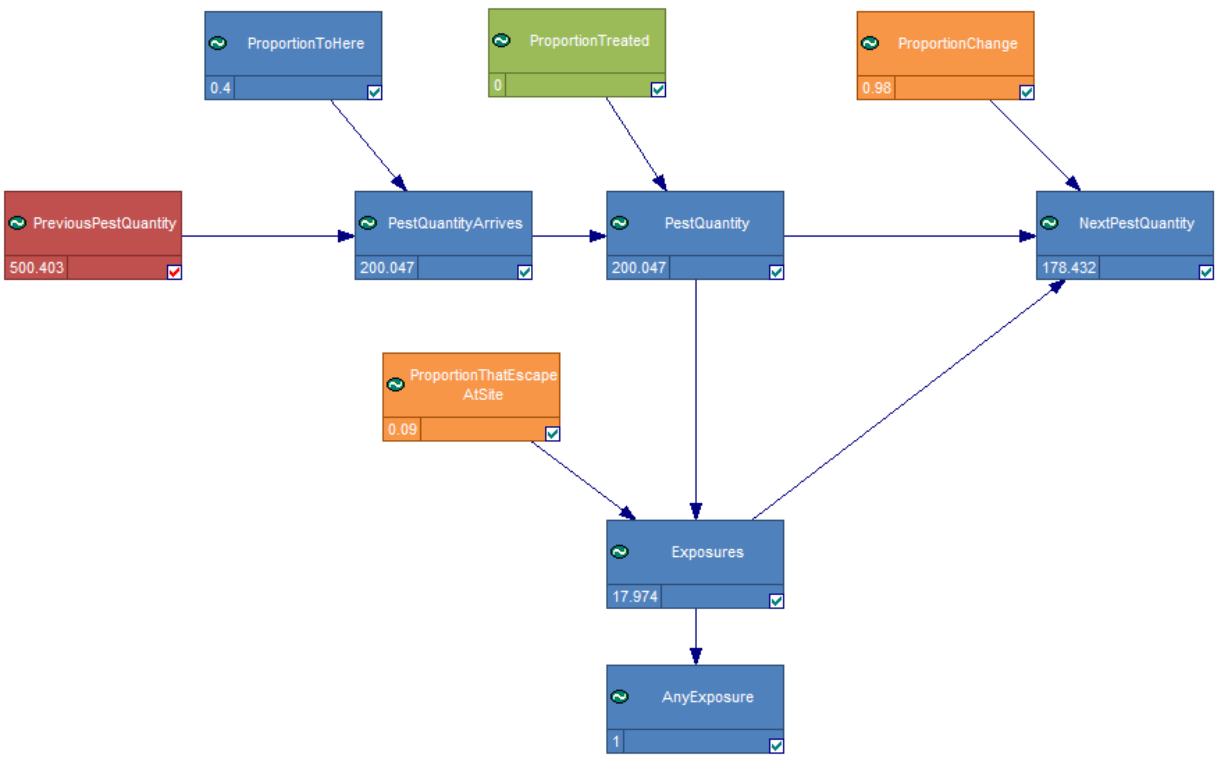

Figure 5: The Pathway Point Submodel template. (See Section 9 for variable definitions.)

3.1.5 Inputs

Inputs come from each of the submodels described above, and these are coloured in Figure 5 according to the submodel from which they derive.

The Pest Arrival submodel supplies the pest quantity distribution for the PreviousPestQuantity node when at the first point on the pathway; subsequent points simply copy over the value from the previous point's NextPestQuantity. The ProportionTreated submodel is invoked for every

[4] Note that it differs a little to a standard DBN, because the variables in each 'slice' refer not just to different values of a variable at different times, but also to values of a variable at different locations.

pathway point; when it's not applicable, the ProportionTreated input is simply set to 0 (as can be seen in the figure). The values for the ProportionThatEscapeAtSite and ProportionSurvive nodes are taken from the Pest Activity submodel.

There is one more input node, which is *not* derived from a submodel: ProportionToHere. Recall that the Pathway Point model describes what occurs across a set of locations (e.g., across all ports or across all car yards). As pests and items move further down the pathway, they are divided at each transition. For example, used vehicles coming into a single port will be sent to different registration sites, vehicles coming from a single registration site will be sent out to different car yards, and so on. The idea is illustrated in Figure 6 for AGM. Each path through the locations connected by arrows constitutes an instance of a single exposure pathway – clearly there are quite a few pathways here, and there will be many more in any real model. To make the network tractable, not all of these pathways are modelled individually. Instead the outputs for all the locations in one stage are summed together, and then redistributed to the next stage. Redistributing them uniformly doesn't work well in many cases (some end point areas, for instance, have a very high density, and should receive many more used vehicles than others). This is where ProportionToHere is used, allowing us to specify how many items (and thereby pests) arrive at a given location within a pathway point from the overall total coming from the previous point.

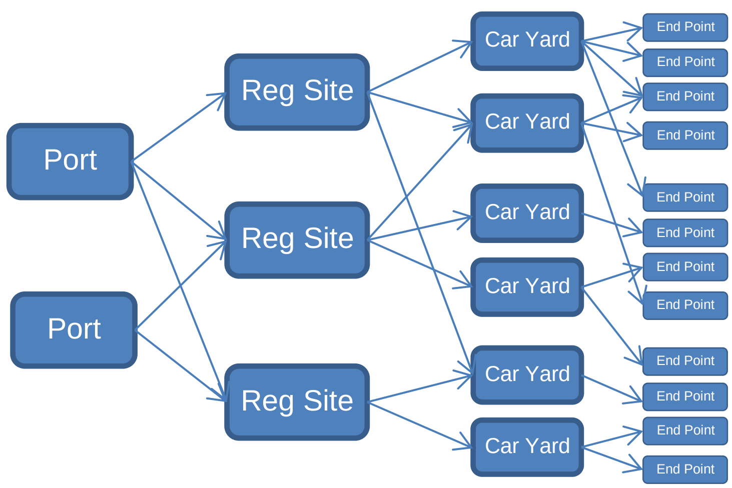

Figure 6: The fanning out of items/pests along many different pathways

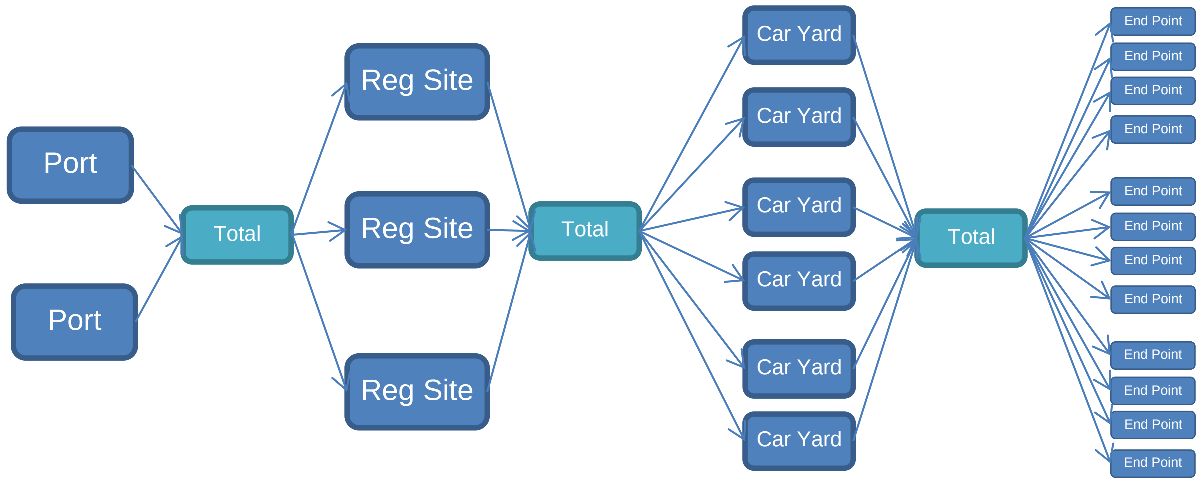

Figure 7: Simplifying the pest pathways

With this in mind, the model flows quite naturally. Along the centre, we have a sequence of pest quantities. Each subsequent quantity is some fraction of the previous quantity, based on removing or filtering on the proportion specified in the nodes above (as well as removal of any escaped pests from the nodes below). We then observe the number of exposures at some point along this internal path, namely after treatment, but before deaths and (of course) exposures are deducted from the quantity.

3.1.6 Outputs

The key output node from the full model is the 'Exposures' node, which gives the probability distribution over the number of exposures in a year at a given pathway point. Exposure here specifically refers to the point of escape (i.e. uncontrolled release) of the pest into the environment. An exposure event can occur in several ways:

- When a pest leaves the item on which it hitchhiked into the country (such as when a gypsy moth egg attached to a used vehicle hatches)
- When the item carrying a pest itself escapes (such as when soil containing phytophthora falls off into the environment)
- At the point that the import item ceases to be tracked
- When a pest arrives in the country in the absence of an obvious item to which it attaches (for example, Culicoides blown into Australia via cyclones)

The exposure distribution is determined by using the pest quantity (PestQuantity) at that point, coupled with the proportion of pests that escape (ProportionThatEscapeAtSite). This is modelled as a binomial distribution, with n=PestQuantity and p=ProportionThatEscapeAtSite.

‘AnyExposure’ simply gives the probability of *any* exposure (i.e., 1 or more) in a single year at that pathway point.

3.2 Model Changes in Stage 2

3.2.1 Dealing with the breadth of possible threats

Ideally, the full FHS model will assist the NZFOA and MPI in targeting their surveillance to achieve the greatest chance of detection for any given amount of surveillance effort.[5] This goal applies to *all* pests and items, known and unknown. Of course, there is an impracticably large number of pests and items, and the exposure model focuses on just a small number of these. However, this model is intended to be more general than what it may at first seem.

The modelled pests – AGM, WBB, Phytophthora, Fusarium and PSM – are intended to cover broad classes of possible pests (. This is due to an expectation (by MPI and its experts) that most forestry pests will share key traits with at least one of these.[6] Thus, AGM acts as a representation for any pest that enters the country solely in the form of eggs, PSM for those that enter as larvae, and WBB for those that enter as live adults. Both Phytophthora and Fusarium represent microbial pathogens that are embedded within some carrier (such as soil). While Phytophthora only has a single notable carrier, Fusarium represents pests that can arrive in the country via multiple carriers (namely, soil and live pests). MPI also considers windborne pests to be important to forestry health, and the model does not yet handle these; however, see Section 3.4.3 for a discussion of how these may be included in future.

In addition, the model includes a Generic pest submodel that is intended as an ultimate catch-all that can be easily tweaked to fit different pests well at a high level of abstraction. While this means accuracy (essentially, model depth) suffers, it ensures that the model is very general (i.e. it has good breadth). A description of the Generic pest model is given in Section 3.4.4.

Table 1: Pest escape models and their features

Escape Model	Pests	Degree-days	Arrival state	Example item
AGM	AGM-like	Yes	Egg	Used vehicles
Generic	(Any)	No	(Any)	(Any)
PSM	PSM-like	Yes	Larva	Live plants
Soil	Phytophthora, Fusarium	No	Live	Sea containers
WBB	WBB, Fusarium	No	Live	Wood packaging

[5] There is the caveat that this should be weighted by the value we attach to outcomes: for example, how much damage a pest is expected to cause or how much it costs to do different kinds of surveying effort. (The latter is the subject of a later stage of the full FHS model.) As per the Stage 1 model and report, we do not consider values here.

[6] By no means are we suggesting that the behaviour of these pests is exhaustive of the wide variety of possible pest behaviours. Rather, we are suggesting that these pest models (or some small additional subset) take account of most of the key influences that affect the chance of a pest escaping an item.

Table 1 provides a summary of the key escape models provided by the Stage 2 model (both existing and new), along with an overview of their key features.

In addition to a multitude of pests, we also have quite a few items and pathways. For several reasons, the number of possible items is not so much of a problem. For one, there are many fewer items than pests. Items entering the country also tend to be well tracked (other than for things like wind), so it is easier to add a thorough model of an item pathway. In addition, like pests, items share various features with each other. Possibly the key feature is that many (though certainly not all) items tend to move through high population areas. Note that we have not created a 'generic' item (the way we have created a generic pest), as there is too much variation between items and too few items to generalise over.

The above considerations make it clear that a model that only applies to specific pests on specific items can, in fact, apply quite generally to many pests and items, even ones not yet modelled. Just how true this is would depend upon the results of a thorough sensitivity analysis. We do not do that here, but we do perform a restricted sensitivity analysis on many of the models (see Section 5). We expect that one of the results of a thorough sensitivity analysis is that item factors would contribute much more to the final exposure map than do pest factors.

3.2.2 Decoupling item and pest parameters

In the Stage 1 model, items and pests were tightly coupled as it was expected items would largely be pest specific. A pathway could only be chosen where specific item-pest combinations were pre-defined. *Every* parameter had to be specified for a new item-pest combination, even if those parameters applied only to the item (and hence were the same for all pests) or only to the pest (and hence were the same for all items). If one item is mostly associated with one pest, this would not be such a problem. However, on further examination it was found that several pests shared the same items and that it was limiting to be restricted to specific item-pest pathways. Hence, we have decoupled items and pests and the current system, in principle, allows any combination of the available items and pests.

Despite decoupling pests and items, the current software requires that all pest-item combinations be specified ahead of time in the parameter file. This is because there are still several parameters that depend on the specific pest-item combination, and these would be difficult to present in the interface without a more elaborate way for managing them. If the user wishes to explore a combination that is not already contained in the parameter file, they can achieve this by modifying the parameters file as follows:

1. Add a new row to the 'itemPestPathway' tab that defines the pest-item combination
2. Add rows to the 'pathwaySource' tab that specify the infestation rates for various combinations of source port and pest-item.
3. Add rows to the 'pathwayDetection' tab that specify the detection and treatment rates for a pest on each pathway point for an item

It is anticipated that in a future iteration of Spear, these steps can be performed dynamically upon performing a model run, without having to change the parameters file in order to allow for simpler model running and exploration.

3.2.3 Correction for small quantities

The pest escape models are used to compute a probability of escape for a pest at any given point in an item pathway. For example, we can use them to compute the probability of an AGM escaping at a registration site or a car yard. The model uses a binomial distribution, taking the number of pests that arrive at a particular point (n) and the probability of escape (p_e) as parameters, to determine the distribution $B(n, p_e)$ over the number of pests that escape at that point. It makes use of the same type of binomial distribution to compute the results of treatment, mortality and so forth.

To keep the model simple, the number of pests arriving at a point is approximated (e.g. the model calculates the expected number of pests going to a given car yard from the sum of all registration sites), rather than computed exactly (e.g. combining the distributions for every registration site → car yard combination). If this number is high, rounding this number to the nearest whole number is unproblematic. However, for low numbers, this leads to some inaccuracy which was not dealt with in the Stage 1 model. A more significant problem, which went unnoticed in the Stage 1 model, is that GeNIe does not apply rounding or flooring to n, but rather applies a ceiling function.[7] As a consequence, for very low expected n, such as $n = 0.01$, GeNIe computes $B(\mathrm{ceil}(0.01), p_e) = B(1, p_e)$, which gives higher estimates for the number of pests that escape *and* that continue on the pathway (and that go untreated, etc.) than should have been the case.

This has been corrected in the new model by multiplying each such binomial by a simple correction factor, $n/\mathrm{ceil}(n)$.[8] At larger values of n, this value approaches 1, and therefore has a negligible impact on the output. At smaller values of n, particularly < 0.5, this adjusts the Binomial output so that it has the correct mean.

3.2.4 Updates to the Phytophthora model

From the very beginning of the project, it was known that the Phytophthora model would be subject to the most uncertainty of all the pest models. A relatively simple prototype Phytophthora escape model was created in Stage 1, which was pared back to be even simpler still (removing an account of sporing in the process). Most of the parameters in this model were taken to be placeholders. Fortunately, the item pathways (namely, air travellers and soil on sea containers) for Phytophthora were modelled with a great deal more certainty, so the model was expected to be at least partially informative.

[7] This is highly counterintuitive!
[8] Since GeNIe does not have an externally accessible ceiling function, ceil is implemented as $\mathrm{round}(n + 0.499999)$.

In Stage 2, we conducted an elicitation session with experts (described in Section 4) on many of these uncertain Phytophthora parameters. It was quickly determined that some of the basic assumptions that went into the Phytophthora model were either false or strained. While the basic structure of the simple escape model held up reasonably well, some of the assumptions underlying the parameters used for the model did not. In particular, the experts broadly agreed that almost any sized soil sample taken from a location of concern would contain Phytophthora. (Note that this is assumed to apply to any type of Phytophthora, importantly including those not currently present in New Zealand.)

Hence, it was decided that the focus of the model should change, and that soil itself be treated as the pest (i.e. a substance for which escape events are tracked). One small issue with this is that the question of dangerous levels of soil is left entirely open by this model. Strictly speaking, this is true of all the pest models (for example, perhaps the escape of one AGM or beetle is not so dire) – however, it is especially true for Phytophthora.

3.3 Monthly Models

MPI requested that an experimental monthly component be incorporated within the AGM model since several factors are affected by specific months. For example, temperatures are of course month specific and can have significant effects on the development of eggs or larvae. These aspects are, in some sense, lost in a yearly model and it is therefore likely to be less accurate.[9]

For Stage 2, we created a monthly model for AGM and several other pests. We also created changes to the item submodels, which can be applied to any pest-item pathway. Support for running monthly models is not yet included in the software interface, but can be run via the command line.

3.3.1 Factors influenced

Many factors vary according to the month. In some cases, obtaining monthly data for these factors is not straightforward and in other cases, the impact of month is known (or strongly believed) to be minimal. The following subset of factors were chosen as the most significantly affected by time of year and the most useful to provide higher time-resolution modelling for (see Section 9 for variable definitions):

- Average minimum and maximum temperatures (AGM and PSM)
- Egg age (AGM)
- Larva age (PSM)
- ProportionToHere in the PathwayPointTemplate model. Currently, due to data limitations, a monthly model is not meaningful. Hence, it has no influence, but a user may update the relationship between Month and ProportionToHere if they can specify an accurate influence.

[9] Note: an example of this difference is illustrated in Section 6.3.

- The generic model incorporates a Month node which initially influences no other factor. If monthly data is available for a pest being modelled, the model can be updated as needed.
- In the Pest Quantity Template model, NumberUnits and UnitInfestationRate are input parameters and are not affected by the month. However, it is recognised that these parameters could indeed be affected by a particular month or an export month and therefore these factors have been made available in the model. It is up to the user to modify these parameters appropriately when modelling specific pests.

In all cases, monthly variations have been encoded directly into the BN models (using equations), rather than stored in the parameter files. Data for most of the item parameters for the monthly components of models were obtained relatively easily and can be considered accurate, due to being reports directly from organisations like Statistics New Zealand. Pest parameters also often had good data backing them (due to a wealth of observations in the pest's original environment). Of particular interest, egg age and larva age were determined by considering months when pests arrive from the country of origin. Hence, there are many months where there would be no oviposited eggs or larvae.

If the monthly approach is to be integrated into the full model in future, very few changes would now be needed to the BN models. Most of the changes would involve the parameter and data files.

3.4 New Pest Models

In addition to the original pests (AGM, WBB and Phytophthora), two other pests were considered significant threats. These are pine root pathogens (covered by the Fusarium model) and the pine shoot moth (covered by the PSM model). Additionally, a "Generic" pest category has been introduced to model other or unknown pests at a more abstract level.

3.4.1 Pine root pathogen

The pine root pathogen (also referred to as Fusarium or pitch canker) can be transmitted via multiple vectors such as soil and beetles (i.e. WBB). Since these vectors had already been modelled explicitly as individual pests in Stage 1, an escape model was initially built by combining the Phytophthora and WBB models. However, analysis of the resulting model showed that combining the factors across these models leads to uncertain behaviour and inaccurate parameter values. Hence, the combined model was discarded.

Broadly speaking, the pine root pathogen can enter by two main vectors and therefore two separate pest models have been used. The first is a soil model which is very similar to the Phytophthora model and the factors, influences and parameters have not been modified.[10] The second model is the WBB model, since beetles are another possible vector for Fusarium. The particular item pathway being considered determines which pest model is used.

[10] Note however that the Phytophthora model itself has undergone significant changes, particularly in interpretation, as described in Section 3.2.4.

The items of interest for Fusarium are sea containers, returning residents, visitors and used machinery. Currently, for sea containers, the default model is the beetle model. For returning residents, visitors and used machinery the soil model is used as the default. These can be changed through the parameters file by changing the "escapeModel" in the "itemPestPathway" sheet if required.[11]

3.4.2 Pine shoot moth

The characteristics of the PSM are similar to the AGM. However, the chance of this pest entering New Zealand is significantly less likely. Nonetheless, many of the parameters used for the AGM are also applicable to the PSM.

The items of interest for the PSM are sea containers, returning residents, visitors and used machinery. As mentioned earlier (see Section 3.2.2), other items can be combined with the PSM as needed.

The escape model for the PSM is very similar to that of the AGM. The main difference is that the PSM would arrive in larval form and hence "egg age" in the AGM is replaced by "larva age" in the PSM. Of course, depending on the month, the egg age or larva age can be very different, with the data being obtained from the literature.

3.4.3 Windborne pests

There is a small chance of windborne pests arriving in New Zealand. Hence, MPI requested that such a pathway also be included in the model.

A tool to simulate windborne pest movements has already been created at CSIRO, called TAPPAS (https://tappas.csiro.au/login/auth). TAPPAS is able to model the long-distance, wind-based spread of pests from a source point to a range of destination points (called a forward model). TAPPAS produces a map, showing the concentration (or equivalently for our purposes, probability) of such pests landing (or reaching some other height) at locations around the point. It is also possible to run TAPPAS in reverse (called a backward model): i.e. to provide a *destination* point, and then to calculate the probabilities for a pest originating from other locations and then landing at that

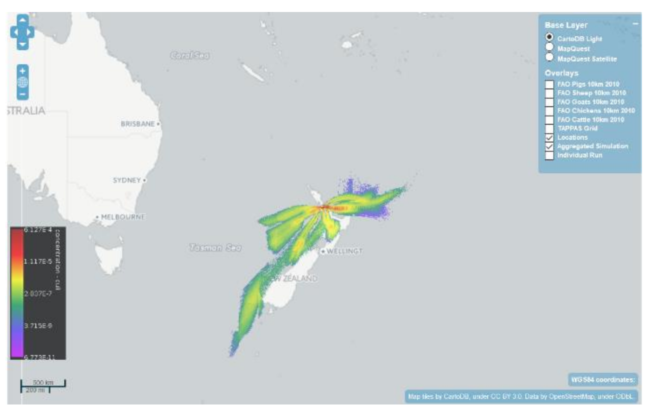

[11] The software does not currently explicitly support multiple escape models for a single item-pest combination. However, much the same effect can be achieved by duplicating the item (and its associated data and parameters) as needed.

point.[12] An example output map for a backward model is shown at right.

We explored this tool extensively and, while it appears to have good potential, it is under development and is not currently of production quality. Furthermore, the parameters required by TAPPAS are numerous and highly demanding (and the model slow-running relative to our needs), making a systematic exploration of the TAPPAS simulation a large modelling project all of its own.

Nonetheless, we believe that a simulation such as TAPPAS will be very useful, perhaps even essential, to modelling the probability of pest exposure via this pathway. Thus, windborne pests are currently not modelled but may be incorporated in the future.

It is useful to consider how we can model windborne pests entering New Zealand using TAPPAS (or a similar simulation) in the future. Assume the following:

- We have a specific pest of interest.
- We have some sense of the (probabilistic or numerical) distribution of that pest in its source countries.

Furthermore, our goal is to estimate the expected number[13] of pests of this type arriving at any given point (for any given year) across New Zealand.

One simple approach would be to sample the source locations according to the probabilistic distribution of the pest, and then to run TAPPAS forward from these sampled points. We can then simply sum and normalise the outputs from each of these runs to produce a probability estimate for where pests will land. If we have the numerical distribution of the pests in the source countries, we can weight each TAPPAS output map by the quantity of the pest in the source location for that map.[14] This would then give us the *expected number* of exposures (and not just the probability of an exposure) for areas across New Zealand. It would also give us an exposure map that is directly comparable to the maps produced by the current exposure model.

One issue with this approach is that it is likely to be highly inefficient and quite possibly impractical. If too small, our source sample may result in maps that completely miss the most important destination locations. A better, more practical, approach may be to choose a set of representative destination points within New Zealand.[15] In particular, such points could be concentrated in areas of particular concern. For each such point, we could then run TAPPAS in reverse, giving us an origin probability map for each point. Each such map tells us:

[12] In the near future, it may also be able to perform a model run from a single source point to single destination point (called a forward-backward model). This functionality is present in the software but, at the time of writing, does not work.
[13] Or, if we wish, distribution.
[14] This would require partitioning the source pest distributions, counting all pests nearest to the source location as if they had emanated from that specific source location. This can be done quite easily by constructing (say) Voronoi polygons, but the usefulness would depend on the quality of the data.
[15] We could use these points to partition New Zealand, again using something like Voronoi polygons.

$$P(O = o_i | D = d)$$

for a set of origin points $\{o_i\}$, where $i = 0..n$, and a single destination point, d. We could use Bayes' theorem (coupled with the pest distribution at the source) to calculate the reverse probabilities; i.e. $P(D = d|O = o_i)$, which essentially gives us n forward TAPPAS simulations that all happen to wind up at d. As above, we could then aggregate these and combine them with the numerical source pest distribution to produce an expected number of pests that appear at that point. This too would give us an exposure map that is directly comparable to the current exposure model's maps. However, it is much more likely that this approach will be computationally feasible.

3.4.4 Generic pests

As mentioned earlier, the Generic pest was requested by MPI so that rarer pests or ones that cannot be planned for can be modelled at a later stage. It also serves the purpose of being able to focus some of the model explorations on just the item pathway. While we have especially created a Generic pest model, we have not created generic versions of any other parts of the model (i.e. the pest pathway and point template models remain the same). This is largely because those other parts of the model are already well generalised.

The Generic escape model has been deliberately built to be simple (Figure 8). It contains 6 nodes: 2 input nodes (TimeAtSite and Month, the latter unused and for compatibility only); 2 output nodes (EscapesAtSite and Survives); and 2 “parameter” nodes (DailyEscapeRate and MortalityRate).

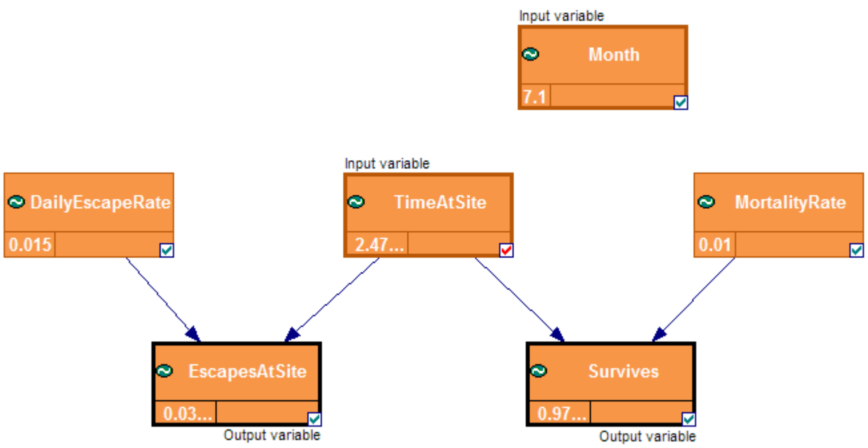

Figure 8: The Generic pest model

Using the other pest models as a guide, the main influencing factor is TimeAtSite. The first output node, EscapesAtSite at site, is a function of TimeAtSite and is modelled as an exponentially decreasing distribution. This is in order to capture the effect that the probability that the escape has occurred after any length of time increases, but tapers off as the length of time increases (since the pest is more and more likely to have already escaped). Furthermore, the proportion that escapes at site is influenced by a daily escape rate (DailyEscapeRate), i.e. a pre-determined constant which captures the proportion of pests that could escape on any particular day.

The second output node, Survives, is also affected by TimeAtSite via an exponentially decreasing distribution. However, an additional parameter, MortalityRate affects the population by further reducing its size over time.

A Month node is made available in the model, but only for compatibility purposes when running the model as a monthly model. This node does not otherwise influence any of the other variables.

The software allows the generic model to be run much like any other. At this stage, the software does not currently provide any specialised support, however it is expected that future versions would include such support. This would allow a user to select an item, couple it with a generic pest and merely specify *just* the DailyEscapeRate and MortalityRate, or perhaps (with monthly model support) also the relationship between Month and these two variables.

3.5 Parameters

The Stage 2 model contains an extremely large number of parameters. Where possible, parameters were first drawn directly from data collected by government, public or industry bodies. These are listed in Section 3.5.1 below. When data was not available, we then checked the literature for estimates of these parameters (Section 3.5.2). In the absence of strong estimates in the literature, we consulted directly with MPI and Scion's experts (Section 3.5.3). In some cases, parameters were either well-known (for example, logically required) or, conversely, too difficult to estimate by any of the above means. In such cases, we have used essentially a priori parameter values (assumed parameters).[16]

3.5.1 Data

The following is a list of parameters that was determined, either wholly or substantially, from data:

1. Transit times (for seaports and airports)
2. List of pathway points (such as ports, shops, yards, etc.)
3. Latitudes and longitudes (for pathway points that have a specific location)
4. Proportion from source (for each item-source port combination)
5. Average minimum, maximum and base temperatures
6. Time at site
7. Certain WBB pest-specific parameters (Activity, Average life time temperature, Lifespan – derived from data collected from Hylurgus and Anoplophora)

[16] Just a reminder that every model (whether a BN, spreadsheet, mental model or anything else) contains both explicit and implicit parameters. Implicit parameters include things like model choices and omitted parameters, which can be endless in number. We mostly stick to a discussion of explicit parameters here.

8. AGM infestation rates

Details on data sources can be found in the comments of the parameter spreadsheet files.

3.5.2 Literature

The scientific literature is often a good starting point to obtain parameter values, but are not always a good fit. This is often because the statistics stated in the literature have been expressed in a form that does not suit the needs of the model. Nonetheless, we can often adapt these data and statistics to the model by making some additional assumptions. The Stage 1 model made heavy use of statistics from the literature where raw data was not available. For example, the detection rate for the AGM on used vehicles was obtained from Wedde *et al.* (2005) which was a monitoring survey of used vehicles entering New Zealand. While we continue to use such data and statistics in Stage 2, we have shifted the focus much more to expert elicitation.

The following is a summary of the parameters that were determined from the literature:[17]

1. Yearly item units: for port-item combinations
2. Proportion to here: for port-item combinations
3. Infestation rates: for source (airport or seaport)-pest-item combinations
4. Egg, hatching and max age: AGM

Details on literature sources can be found in the comments of the parameter spreadsheet files.

3.5.3 Expert Elicitation

In general, several parameters may not be straightforward to determine via data or literature. Even amongst those parameters which can be determined, several can often be unreliable. Hence, an alternative approach is to use the combined knowledge of experts.

We make use of the IDEA protocol (Hanea et al., 2016), which is in turn a modified form of the Delphi elicitation technique (Brown, 1968; Linstone & Turoff, 1975). In particular, the protocol employs a specific question structure as a means of neutralising a range of cognitive biases beyond what the original Delphi technique is capable of alone.

When running an elicitation session, we select a set of suitable parameters, on the basis of which parameter values are least certain (had the least support in terms of data or background literature) or otherwise have the greatest impact on the outcome. We do not elicit values for all possible parameters as there are too many to handle well within an elicitation session. Instead, we choose a subset to focus on.

[17] There is some overlap with data, due to data that has been reported relatively unprocessed in the literature.

Participants are asked a specific question per parameter. The participants provide answers, which are collated and presented back to the group, question-by-question, in an open discussion. This is followed by a second round in which the same question is answered again.[18] The second round estimate (the average of the group's estimates) will better reflect the group's understanding of the what the parameter value should be. A detailed description of the implementation of this method and the ensuing results are presented in Section 4.

3.5.4 Assumed Parameters

Stage 1 contained a very large number of (explicitly) assumed parameters. In some cases, these parameters were intentionally specified with dummy or placeholder values – some value that allowed the model to run and nothing more. In other cases, the values were based on a informal expert opinion. For Stage 2, we have done a great deal of work to reduce the number of these parameters by making use of more extensive and systematic expert elicitation, as well as by making use of improved data sources. The model no longer contains any values that can be considered placeholders. However, there are still a few parameters that have yet to be specified with any rigour.

The Phytophthora model specifically contains a large number of extremely uncertain parameters. This is partly due to lack of a good understanding of Phytophthora's distribution and spread, which makes the model structure, and in turn parameters, for Phytophthora equally uncertain. In addition, parameters relating to distributions (other than means) have often been chosen based on what we as modellers (with some limited help from experts at MPI and Scion) have considered plausible given our discussions with experts, rather than on anything more solid. This includes parameters such as standard deviations, as well as the choice of distribution itself. (The distribution is typically assumed to be normal, where there is any variation at all.)

As with the Stage 1 model, the most significant assumed parameters lie on the pest-side of the model. Parameters associated with item pathways tend to be well covered by data. The main exceptions are the parameters associated with air traveller activity at the end point. While resident locations are expected to be well approximated by population density, visitor locations are likely spread across tourist sites and attractions. While this is taken into account in the model, the relative proportion of tourists that visit (say) golf courses, nature sites or camp sites is currently only crudely estimated.

In Stage 2, we have made an attempt at exploring the impact of these assumptions by way of a sensitivity analysis on various parts of the model (see Section 5). That analysis, however, is still very preliminary.

[18] In principle, several rounds can be conducted leading to improving estimates with increasing number of rounds.

4 Delphi Elicitation

A modified Delphi elicitation (Brown, 1968) was conducted to determine values for the less certain parameters as observed during the development of the models. The elicitation was conducted twice in different ways with (nearly) the same participants. The participants were encouraged to answer all the questions but were not required to do so. The aim was to get estimates from 11 participants with a minimum of 5-6 estimates for every question.

Firstly, an in-person elicitation was conducted at Scion Research in Rotorua, New Zealand on the 31st of March, 2016. The session consisted of 11 participants who were considered experts for different organisations around New Zealand. The session lasted one day where the same set of questions were presented in two rounds with a discussion in-between rounds as well as during the second round.

Secondly, an online elicitation (http://fhsdelphi.bayesian-intelligence.com/) was conducted with the same participants between the 9th and the 15th of June, 2016. It also consisted of two rounds and a discussion took place during the second round. While the online interface was an unfamiliar setting for the participants, it was expected that this would have no effect on the results or parameters elicited.

The participants were required to provide four estimates for each question in each round. The first three estimates were a low, best (middle) and high value within a range that they would assign to the question. The fourth estimate was a confidence (percent out of 100) level they associated with their answer. While only the average of the best estimates were used in the model, the ranges and confidences were used to inform the choice of parameters in the sensitivity analysis.

The range of values to be selected varied across the questions. Most questions required answers in the form of probabilities (typically rendered as percentages) and others required answers in units such as days or grams (see Appendix A.1 and A.2). In order to analyse the questions as a whole, the percentage difference from round one (R1) to round two (R2) has been calculated: $100 \times \frac{(R2-R1)}{R1}$. Individual question results are also analysed and discussed where interesting points were seen.

4.1 In-person elicitation

The questions varied across pests and items (see Appendix A.1) with a focus on the AGM and Phytophthora among pests and used vehicles among items.

Figure 9 shows the percentage difference from round two to round one for the high, medium and low estimates for each of the 15 questions averaged across the participants. A summary of the findings of the in-person elicitation are as follows:

1. The best estimates (PDRM) do not change substantially between rounds. However:

a. There is a slight increase in the estimate for the question about an egg mass ballooning to shore (Q2).
b. There is an increase for the questions about whether soil contains Phytophthora (Q6 a,b,c).

2. The high estimates (PDRH) also do not change substantially between rounds. However:
    a. There is an increase for two questions (Q1b and Q3): time that a used vehicle spends at a registration site and the amount of soil attached to a container when it arrives at port.
    b. There is a slight increase for the soil-based questions (Q6 a,b,c).
3. There are often large percentage differences for the low estimates (PDRL: Q2, Q4b, Q5) but on occasion, small differences as well (Q4, Q4c).
4. The estimates for the treatment of an identified egg-mass are consistently high (average-high: 0.99, average-low: 0.8). This is expected since once an egg-mass is found the treatment should be effective.

Regarding the questions concerning the presence of Phytophthora (Q6 a,b,c), there was a significant discussion between most of the participants. These questions and their answers were contentious as some experts were of the opinion that any quantity of soil will always contain Phytophthora. This opinion appears to be reflected in the results as all estimates (low, best and high) increase in round 2. The discussion led to changes in the Phytophthora model, with the focus being shifted to the movement of soil rather than of the Phytophthora itself (as discussed in Section 3.2.4).

Question 2 also stimulated a prolonged discussion. In particular, the chances that larvae would make it to shore were dependent on a number of factors (from larvae not being detected to larvae surviving on the surface of a ship in the sun) of which each participant only considered a subset.

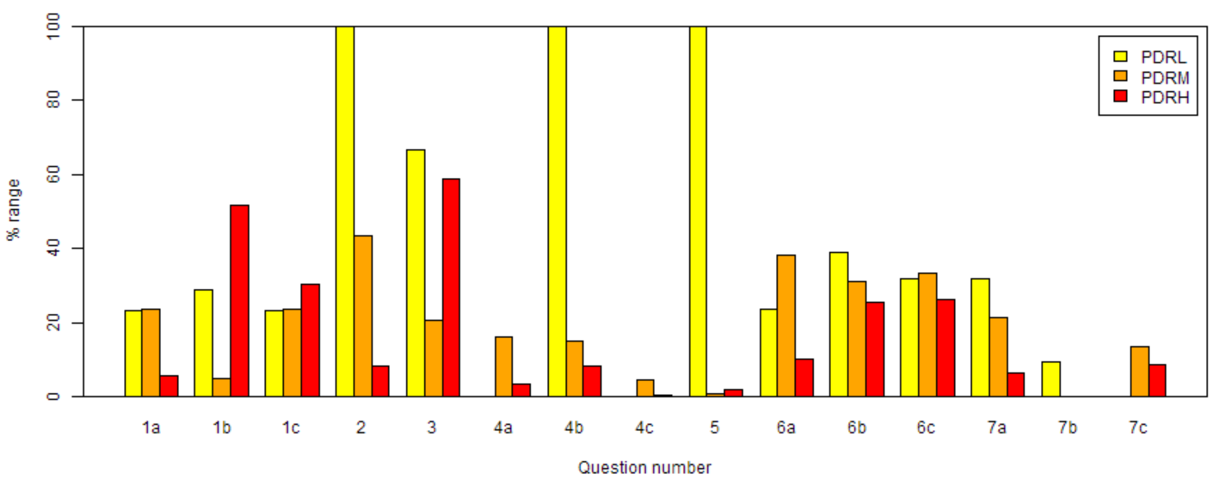

Figure 9: Percentage difference between rounds for the high (PDRM), middle or best (PDRL) and low (PDRL) estimates for the in-person elicitation.

4.1.1 Comparison with original data

If we compare the in-person elicitation with the original data (see Table 4 in Appendix B), we can note the following:

1. Elicited parameters often agree quite well with the original parameter values (see the PDRO column). The exception being questions relating to the Phytophthora model.
2. The elicitation and discussion for Q4 and Q6 showed that many of the assumptions underlying the original Phytophthora model where highly doubtful. While the basic model structure was sound (given substantial examination and acceptance from experts at MPI, Scion and the forestry industry), the interpretation of the parameters needed to be reconsidered in light of the elicitation results. The parameters remained highly dubious even after this elicitation workshop, and so changes were made to the Phytophthora model in order to sidestep at least some of these issues. However, much more work needs to be done on the Phytophthora escape model, though this was already expected.[19]

These results demonstrate the value of eliciting knowledge from experts for improving our models, even in the absence of a formal validation.

4.2 Online elicitation

The questions for the online elicitation focused mainly on the AGM and wood borers among pests and used vehicles, sea containers, used machinery and wooden furniture among items.

The results for the online elicitation can be summarised as follows. Figure 10 shows the percentage difference from round two to round one for the high medium and low estimates for each of the 19 questions. Some of the main points are:

1. The best estimates do not change substantially between rounds. The graph suggests large changes, however the numbers elicited during the online session were all substantially smaller than for the in-person elicitation (so small changes appear larger in this graph). Also of note:
    a. There is a slight increase for wood borers entering on furniture (Q12, Q13, Q14)
    b. There is an increase for the AGM on a used vehicle at a car yard (Q7)
2. The high estimates change the least. However:
    a. There is an increase for the questions associated with time-spent at locations
3. There is a large percentage difference for the low estimates. However:
    a. This is not true for questions which are related to time-spent at locations
4. The estimates for Question 3, infestation rates for beetles leaving from a source country, have a very slight difference across all measures. This shows that all participants were very confident in all their estimates. (See Appendix B.2.)
5. The TimeAtSite questions (Q16, Q17, Q18, Q19) all have very high estimates. This is not surprising since the longest time an item can spend at many of these locations could be very high.

[19] This is not fatal to including Phytophthora in the full model, since item pathway factors dominate for Phytophthora, particularly for the relative exposure probabilities.

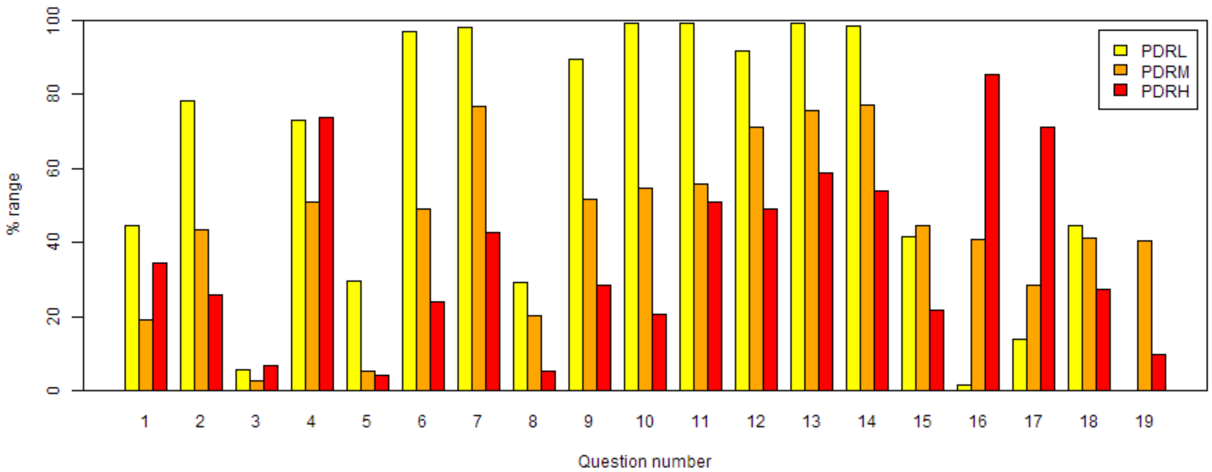

Figure 10: percentage difference between rounds for the high (PDRM), middle or best (PDRL) and low (PDRL) estimates for the online elicitation.

4.2.1 Analysis of the online discussion

Looking at the discussion closely, we see that several participants have changed their estimates in response to the discussions. Here are some of the clearer examples:

1. Question 2: What is the infestation rate for wood boring and bark beetles on wood packaging (percentage of tonnes infested) when leaving port from the source country?
    a. User 11 provides different sources, which could be defined as wood packaging. With this knowledge and further reasoning, the user increases his/her estimates.
    b. User 9 accepts this point but is influenced more by user 15, who has practical experience and data associated with this problem. As a result, he/she further reduces their estimates.
2. Question 6. Assume the presence of a single unhatched AGM egg mass on a used vehicle. What is the probability that this egg mass will be detected at a registration site?
    a. User 16 suggests lowering their estimates based on the comments but does not explicitly say so. (Furthermore, the estimate actually increases)
3. Question 9. Assume the presence of a single unhatched AGM egg mass on a container. What is the probability that this egg mass will be detected at a transitional facility?
    a. User 17 lowers their estimate as a result of a comment from user 9.
4. Question15. Assume a single WBB beetle is located somewhere on a tonne of wooden furniture that has entered the country. What is the probability that this beetle will be detected at any time after leaving the furniture shop?
    a. User 17 lowers his/her estimate and user 16 is influenced by the change of user 17 to lower their estimate as well.
5. User 11 lowered their estimate for all questions related to transitional facilities as they realised that there are very few inspections at these locations

4.2.2 Comparison with original data

If we compare the online elicitation with the original data (see Table 5 in Appendix B), we can note the following:

1. The detection rates are, across the board, lower than the original parameter values, and especially lower for WBB. For example, the WBB detection rate for furniture at port (Question 12): 0.02 (elicited) vs 0.85 (original). Note that many of these detection and treatment numbers were originally used as placeholders, so the large deviations here are not so unexpected.
2. Infestation rates are often higher than earlier, e.g. AGM and sea containers (Question 1): 0.037 (elicited) vs 0.0048 (original). The original number in this case was based on an estimate from the literature (which itself noted very low confidence), and likely constituted an underestimate, particularly as it was based on actual observed detections. Some effort was made to compensate for this in the original parameter, but the responses in the elicitation suggested a further increase still.

Like the in-person elicitation, these results demonstrate the value of eliciting knowledge from experts. The elicited parameters are likely to be much more accurate than the prior values for these parameters, since these particular parameters were not well supported by data or literature.

4.3 Summary

A modified Delphi elicitation was conducted to determine values for parameters that could not be settled by data, literature or a priori knowledge. The elicitation was conducted over two sessions: in-person and online. Both sessions proved effective and the elicited parameters have now been incorporated into the model.

The online session was as effective as the in-person session with respect to the answers provided. Although lacking the benefits of face-to-face communication, the approach has its own unique advantages, such as fewer restrictions on location and time, ease of presenting questions and collecting results, and a reduction in the influence of cognitive biases. Since the in-person elicitation was conducted before-hand with the same participants, it is quite possible this familiarity led to a smooth and effective online session.

The discussions in both sessions proved useful, both as a way of producing better responses to the questions asked, and also as a way of shedding further light on the nature of the models themselves. As is often the case, the in-person elicitation could potentially have benefitted from more time devoted to the discussion.

Two key issues arose. Firstly, the lack of consensus on the Phytophthora made it clear that a different approach was required to how Phytophthora was tracked by the model. This has led to improvements in the Phytophthora model. Secondly, participants' responses for questions with very small values often highlighted substantially different interpretations of what constituted a small value by the participants. This could perhaps be addressed more directly in future elicitation sessions by making use of logarithmic scales.

In Section 6.3, we will look at the impact of the elicitation on the results, by applying the new parameter values and conducting experiments to determine differences between the original and post elicitation parameters.

5 Sensitivity Analysis

We conducted a sensitivity analysis on several key parameters in the BNs, focussing especially on the AGM and Phytophthora escape models. The reasons for selecting the AGM are twofold. Firstly (and importantly), the AGM is a good example of how pests may enter New Zealand. Secondly, the AGM has been the most thoroughly studied model during the course of this project and hence the models and parameters are the most robust and accurate. The reason for selecting the Phytophthora model is quite simply because it is the least certain of all the models. For simplicity and tractability, we restrict the sensitivity analysis to the effects on output variables within a single BN submodel. Since we had much greater confidence in the model structure (due to its simplicity and in parts logical necessity, as well as having developed it with the oversight and help of experts from both MPI and Scion), we chose to focus our available resources on the analyses of parameters.

The sensitivity analysis serves two purposes here, which are closely related. Firstly, we would like to get a better picture of which parameters and nodes influence the results of the model most. Depending on which parameters are most influential, we may wish to spend more effort on improving the estimates for those parameters if doing so is feasible (such as through a Delphi elicitation process as described in Section 3.5.3). Secondly, we would like to understand how robust the model results are in the face of changes to parameters that are considered less central and to discover the key parameter ranges over which the model results are likely to remain roughly the same.

5.1 Approach

Ideally, conducting an exhaustive sensitivity analysis would provide a comprehensive (if summary) view of the interactions between all factors in the model, and the impact on the very final outputs of the model (i.e. in the present case, the exposure maps). However, it is computationally intractable to assess all possible interactions and dependencies here, even using approximate techniques. We instead focus on an analysis of the submodels.

There are a wide range of approaches to performing sensitivity analyses, both in general, and with BNs specifically. Here we focus on a small set of one-at-a-time (OAT) and two-at-a-time (2AT) analyses. OAT involves varying a single parameter while keeping all other parameters constant, and then recording the impact on the output or target variable. This is the simplest form of interaction that can be measured and while it can provide useful insights it can also be too simplistic. 2AT is very similar, but instead involves varying two parameters. This is broader than OAT and leads to better insights, but can be more complex, both to do and to communicate through (for example) graphs and other visualisations. Variance based sensitivity analysis (VBSA) (Saltelli *et al.*, 2008) is a more comprehensive approach that is suited to modelling more complex interactions, but is currently beyond the scope of this project. It is hoped that such an analysis will be carried out in a future stage.

Due to the continuous nature of all the BNs within the exposure model, the approach to the analysis here looks at the parameters that are used within the node equations. Often, these equations consist of simple values (such as for TimeAtSite), in which case the simple value of the node is varied for the analysis. In other cases, the equation for a node describes a distribution (e.g. HatchingAge), in which case the distribution is replaced by a simple value for the sensitivity analysis and then varied. In BN parlance, this type of analysis would normally be called a sensitivity to parameters analysis. The other common type of BN analysis is sensitivity to findings. In a sense, the analysis that we do here is a hybrid of the two, as the nodes that we look at are invariably root nodes that are set to specific values, which is equivalent to setting evidence on that node.

5.2 Main Pathway Model

The two main pathway models are the pathway point template and the pest quantity template. The parameters in these models are limited quite strictly by the structure of the model and there are very few interesting parameters to investigate (that are not themselves determined by external data). Nonetheless, we examined these models to ensure that this was indeed the case.

For the pathway point template, the following factors were considered:

- DetectionRate on ProportionTreated
- TreatmentEfficacy on ProportionTreated

Neither of these provided interesting insights and were directly proportional to the chain of functions that connect them.

Regarding the pest quantity template, the following factors were considered:

- NumberUnits on PestQuantity
- UnitInfestationRate on PestQuantity

Here again, the results were straightforward and did not warrant further analysis.

5.3 Pest Models

Here we focus primarily on the relatively well understood AGM escape model (called "AGM – Escape At Site Template.xdsl"), along with the simpler and highly uncertain Phytophthora model. For AGM, we explore the following factors and the influences they have on the probabilities of the specified output variables:

1. TimeAtSite on EscapesAtSite & Survives
2. HatchingAge on EscapesAtSite
3. AvgMinTemp & AvgMaxTemp on EscapesAtSite & Survives
4. EggAge on EscapesAtSite & Survives

For Phytophthora, we look at the following relationship:

1. TimeAtSite & SurfaceExposure on ProportionThatEscapeAtSite

5.3.1 AGM: Time at site on probability of escape at site and survival

Figure 11 shows that the probability of survival gradually decreases from 1 down to ~0.07 over a range of about 90 days, after which it steadies. The probability of escape at site moves in the opposite direction, increasing from 0 to ~0.66 over about 80 days and then flattens out.[20] Since the range of output probabilities is large, and the range of TimeAtSite values is clearly within plausible variation, this suggests that (for the given hatching ages, egg ages and so forth) TimeAtSite is a highly significant factor for AGM that is likely to impact the results of the model in important ways.

While a TimeAtSite estimate that is off by a day or two would not affect the results, larger differences (on the order of 10-20 days) may produce more significant changes in the results. To make this more concrete, if the current estimate of 2.7 days for time spent at the registration site were to increase or decrease by approximately half, the probability of escape at the registration site would sit somewhere between $[0.01,0.05]$ (i.e. a maximum change of 0.04). By contrast, if the current estimate of 20 days at the car yard were to increase or decrease by approximately half, the probability of escape at the car yard would range between $[0.12,0.38]$ (i.e. a maximum change of 0.26). Thus, it may be important to get a reasonably accurate estimate of these larger TimeAtSite values. Of course, once TimeAtSite is very large (over 100 days or so), such changes no longer have quite the same impact.

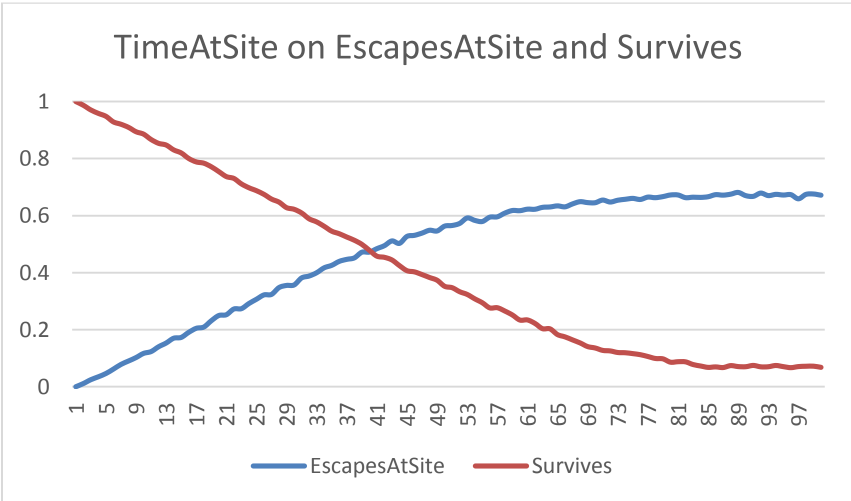

Figure 11: The effect of time at site on the probability of escape and survival. The Y-axis represents the probability while the X-axis is in days.

[20] To be clear, there is nothing at all that connects these two variables. It just happens to be the case that time affects them both in similar, but opposite, ways.

5.3.2 AGM: Hatching age on probability of escape at site

Figure 12(a) shows that hatching age has no notable impact on the probability of escaping at site for quite a large range of possible values. The only exceptions are for very low hatching ages (which are unrealistic for AGM) and for hatching ages which are above the maximum egg age defined in the model of around 140 degree days. It's likely that this is true due to the uniform distribution over egg ages in the yearly model.

If we look at the interaction between HatchingAge and TimeAtSite in Figure 12(b), we see somewhat more interesting behaviour. The colours on this graph indicate the probability of escape, which ranges from 0 (red) through 0.5 (orange) and up to 1 (white). When the TimeAtSite is low (say up to 5 days), then HatchingAge has very little effect on the probability of escape. By default, the escape model includes a value for TimeAtSite which is indeed low (having a mean of 2.5 days), hence Figure 12(a).[21] However, as the TimeAtSite increases, HatchingAge begins to have a more significant effect. If TimeAtSite is very high (say around 80-90 days), then HatchingAge has a significant impact on the probability of escape. Thus, hatching age may be important if the pests stay at particular sites for long enough.

Note also that the reverse is also true. Were the hatching age low, TimeAtSite would have very little impact. But as the hatching age increases, so too does the impact of TimeAtSite. And since HatchingAge in the model is high (with a mean of 95 days), TimeAtSite is a significant factor.

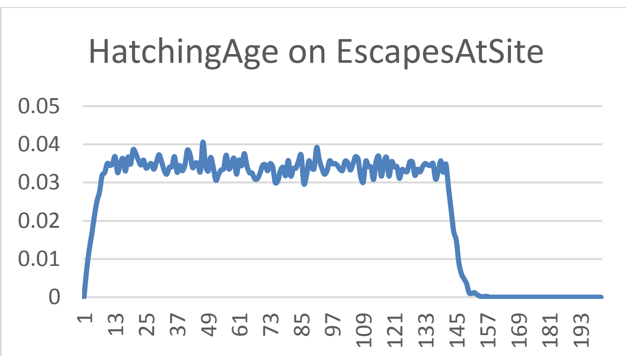

(a)

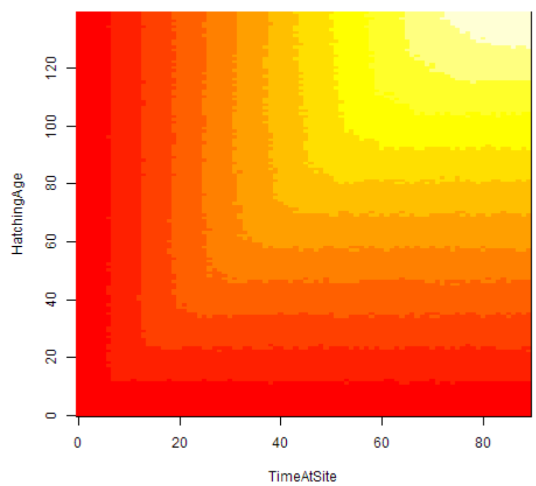

(b)

Figure 12: (a) The effect of hatching age on the probability of escape at site. The Y-axis represents the probability while the X-axis is in days. (b) The effect of hatching age and time at site in combination on the probability of escape. The red-orange-white gradient of colours represents different probabilities of escape, with 0=red; 0.5=orange; and 1=white.

[21] This prior is *always* overridden with pathway point-specific TimeAtSite values when the model is actually run, but it is useful to have a prior here, when doing analyses such as this.

5.3.3 AGM: Average minimum and maximum temperature on probability of escape at site and survival

The average minimum and maximum temperatures have very similar effects on the probability of escape and survival (Figure 13). Higher values for either of these variables leads to decreases in the probability of survival at the site, and an increase in the probability of escape at the site. The effect of the base temperature as a means of calculating the degree days of development can be seen in these graphs, with the effect of temperatures remaining unchanged up until 10 degrees.

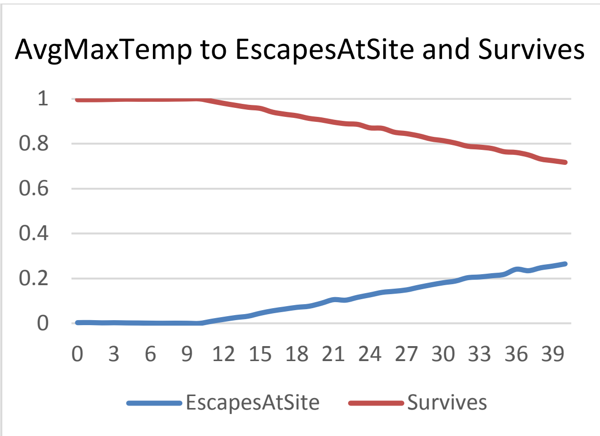

(a)

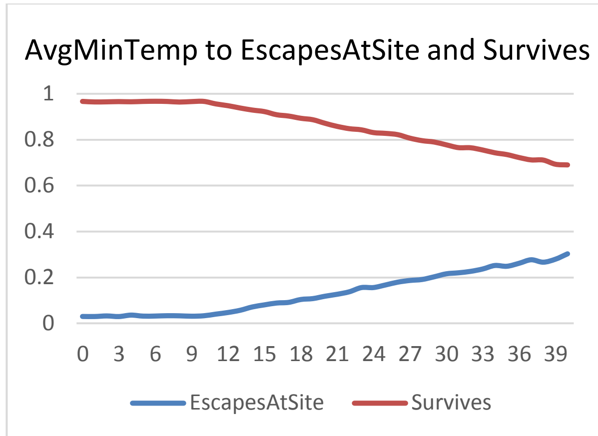

(b)

Figure 13: The effect of the average (a) maximum and (b) minimum temperature on the probability of escape at site and survival. The Y-axis represents the probability while the X-axis is in degrees Celsius.

Figure 14 shows that the effects of AvgMinTemp and AvgMaxTemp are mostly directly additive. That is, any increase in AvgMinTemp or AvgMaxTemp increases the probability of escape, within the feasible ranges for these temperatures. (But not below 10 degrees, of course.)

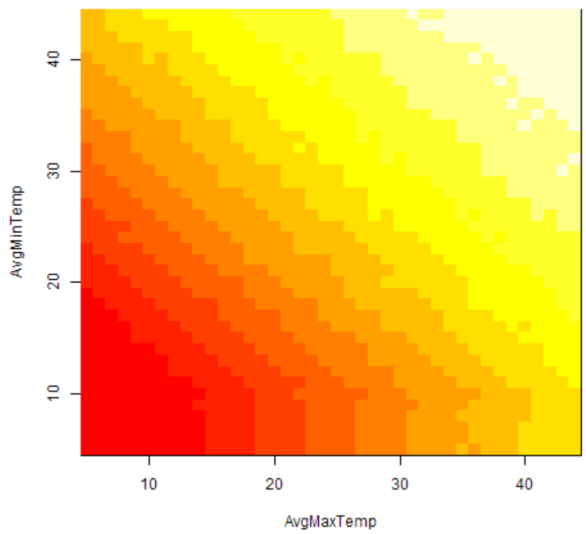

Figure 14: The effect of average minimum and maximum temperatures on the probability of escape. The red-orange-white gradient of colours represents different probabilities of escape, with 0=red; 0.5=orange; and 1=white.

5.3.4 AGM: Egg age on probability of escape at site and survival

Figure 15 shows that as the egg age increases to just shy of 95 days, there is a gradual increase in the probability of escape. Beyond this, the probability gradually drops off. This corresponds with the distribution over HatchingAge, which has a mean of 95 days. Indeed, the peak in this graph occurs slightly before 95 days, since hatching (and therefore escape) occurs while the pest is at the site, not necessarily directly upon arrival. The probability of survival is always high, but drops off after 95 days and continues dropping to 130 days. After this point, the survival rate jumps back to 1.0. Since the maximum age for an egg is around 130 days, the reason for this may not be obvious – however, survival here in reality means “does not die” (rather than “remains alive”). Thus, if an egg is (say) 200 days old it “does not die”, because it is already dead.

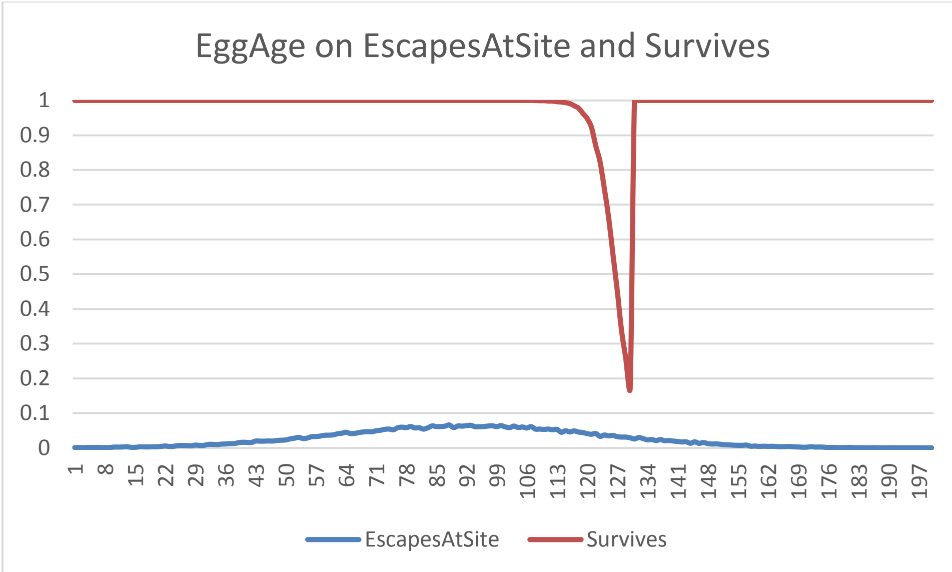

Figure 15: The effect of egg age on the proportion that escapes at site and the proportion that survives. The Y-axis represents the proportion while the X-axis is in days.

5.3.5 Phytophthora: Time at site and surface exposure on probability of escape at site

For Phytophthora, we focus on two variables of particular interest: TimeAtSite and SurfaceExposure. Figure 16 shows the pairwise effect that these two variables have on the proportion that escape at site. What is of interest to note here is that low values for TimeAtSite compress the impact of SurfaceExposure. Since SurfaceExposure is an *extremely* uncertain parameter, the model predictions may be more reliable for pathway points that have very short times (such as the passenger pathways). However, also note that for moderate to large values of SurfaceExposure, the proportion that escape becomes highly sensitive to low values of TimeAtSite – so that the decreased sensitivity to the highly uncertain SurfaceExposure is perhaps defeated by the significant increase in sensitivity to TimeAtSite.

The Phytophthora model and its parameters are known to be highly speculative. Had the sensitivity results shown more robustness in the face of parameter changes, it may have bolstered our confidence. Unfortunately, these results do not provide us with any further confidence in the Phytophthora model.

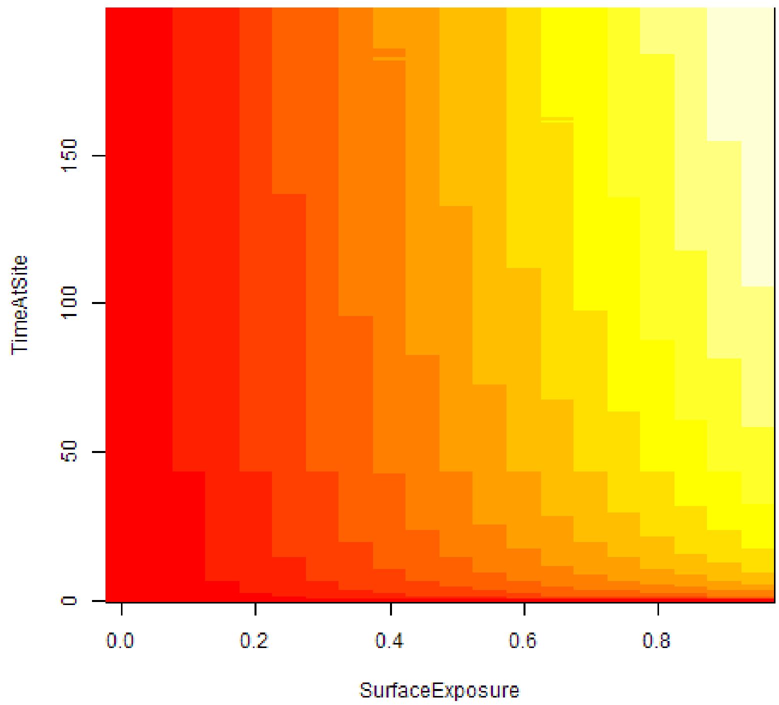

Figure 16: The 2AT effect of time at site and surface exposure on the proportion that escape at site for Phytophthora.

5.4 Summary

The sensitivity analysis has proven quite useful, not just with the results presented above, but with the additional checks (not described here) that it has provided on model correctness – checks that have proven useful in the development of the Stage 2 model, quickly allowing us to identify some model predictions that were clearly wrong for less commonly used parameter values.

The analysis above gives us greater confidence that the model is working as intended and that its results are indeed informative. It also provides us additional reasons to be cautious of the results produced by the Phytophthora model. The brief analyses of the pathway point template and the pest quantity templates confirm that these models are operating in a straightforward and correct manner, and bolster our confidence in their (significant) contribution to the final exposure maps. Analysis of the AGM escape model showed that:

- TimeAtSite is a significant factor
- HatchingAge is significant over a wide range, but within the ranges of confidence, its impact is less significant
- Temperatures have a predictable and robust impact on pest development
- Egg ages are indeed important, but not so much that the current uniform prior over egg ages needs to be improved with any urgency

None of the results suggest that major changes to the model parameters are required. However, if certain model assumptions no longer hold – for instance, if average temperatures were to fall substantially outside their usual ranges for some particular pathway – the models would certainly require updating.

6 Exposure Results & Analysis

We now present a broad sample of results across a number of pests and items. We begin with an overview of the exposure maps produced for each of the pests (or, more accurately, pest classes) currently modelled by the system. These results are not intended to be comprehensive or definitive, but rather an indicative snapshot of what the current model (with its current data and parameters) produces now that Stage 2 is complete. We expect that the parameters and data in the model will be subject to regular change in future, not just due to improvements in understanding but also because the real world processes themselves necessarily undergo change. After this (in Sections 6.3-6.5), we present results that compare some of the key changes and features introduced in Stage 2 with their associated Stage 1 counterparts.

In the upcoming sections, we specifically focus on the following results:

- **Pest specific exposure results.** We obtain exposure maps for each pest aggregated across all items compatible with the pest. These results provide a high-level view of how and where a particular pest may enter New Zealand.[22]
- **Aggregate results.** This exposure map aggregates exposure events across all pests and items and aims to measure the threat of all pests as a whole when entering New Zealand.
- **Monthly and yearly models.** This comparison is designed to explore the value of adding monthly factors to the models. In particular, we would like to know if the results from monthly models are substantially different from the yearly models. We chose the AGM with used vehicles for this comparison. As discussed earlier, this is a thoroughly tested pathway and provides an ideal test case.
- **AGM and Generic.** This comparison is designed to investigate how well a very simple model of pest escape performs against a much more detailed model. Again, used vehicles provides a good basis for comparison.
- **Original vs. post elicitation results.** This comparison aims to determine the significance of differences seen from the original to post-elicitation results.

Note that all of the exposure maps in the following sections display the expected number of exposures per square kilometre per year – for brevity, this will be referred to as the "exposure rate". Each map uses area unit polygons, which were chosen as a more appropriate resolution for these maps. This is in contrast to the higher resolution mesh blocks used in the maps for the

[22] Note that Chatham Island is not included in maps as pest entry into this island is not modelled.

Stage 1 report. Area units and mesh blocks are statistical units defined by Statistics New Zealand.[23]

6.1 Pest-specific Exposure Results

In this section, we examine the exposure results for each pest as they stand at the end of Stage 2. This includes improvements to parameters that were made after the elicitation sessions, as well as an array of small tweaks and improvements across model parameters and data. The results are aggregated across items applicable to each pest. For example, the results for the AGM include all possible combinations with associated items, including used vehicles, sea containers and sea vessels. Each item is given an equal weighting in the resulting maps.[24]

6.1.1 AGM

Aggregated results for the AGM can be seen in Figure 17. We see that the exposure rates are very low across the country, except for more populous regions as well as for areas surrounding ports including Auckland, Tauranga, Napier, Wellington, Christchurch, Dunedin and Bluff. Outside of these zones, the exposure rate drops off quite quickly. The top 100 area units (out of a total of 1911) by number of exposures for AGM account for 69% of the exposures, so exposures are not spread out very much across the country. Note, however, that *every* area unit is subject to some level of exposure (which is not always the case for other pests).

These results are consistent with the results from the Stage 1 model, and are not surprising since the first point of entry for AGMs on every item is the port. Furthermore, the time spent at the different pathway points are of the same order of magnitude leading to the ports seeing the largest exposure rates. This is further accentuated by sea vessels, where the item only includes the single port pathway point, thereby increasing the weight of exposures at these locations.

Focusing on the Auckland region (Figure 18), we see that the port (the strong red area, just above and to the left of the centre of the map) has the highest exposure rate. Generally, exposures are concentrated around arrival points and other key pathway points, especially where pests may spend more time while still in a viable stage (e.g. car yards for used vehicles). Going beyond these regions, the exposure rate drops off substantially but is still more concentrated around the cities.

[23] See the definitions for area unit and mesh block at http://www.stats.govt.nz/methods/classifications-and-standards/classification-related-stats-standards.aspx
[24] More detailed breakdowns are of course possible in Spear.

Figure 17: Map of the aggregated exposure rates for the AGM across used vehicles, sea vessels and sea containers for the North and South Islands.

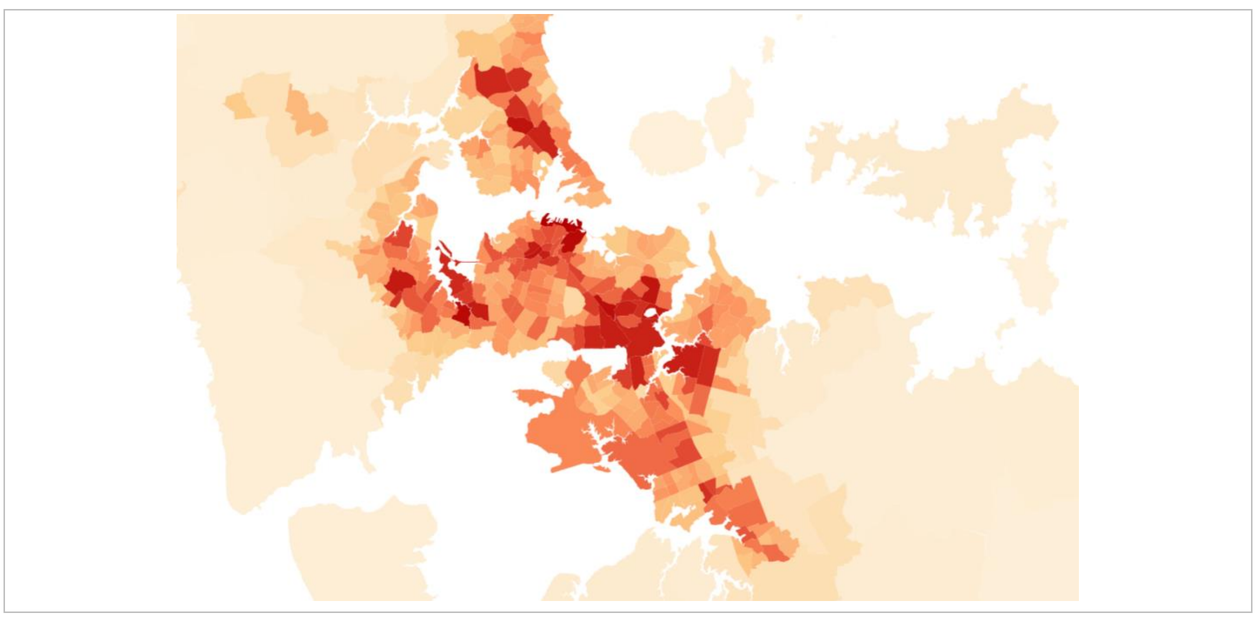

Figure 18: Map of exposure rates for the AGM aggregated across used vehicles, sea vessels and sea containers for the Auckland region.

6.1.2 WBB

Figure 19 shows aggregated results for WBB. Like the aggregated exposure results for the AGM, the exposure rates across the country are low except for a moderate increase in regions with greater population density. In comparison to AGM, overall exposure rates are lower (note the differing scale) and exposures are much more concentrated. For example, the top 100 area units by number of exposures for AGM accounted for 69% of the exposures, while the top 100 area units for WBB account for almost 99% of the exposures. Also of note is that there are many area units that are at no risk of exposure. Only about half of all area units (specifically, 1041 out of 1911) are subject to some level of exposure. This is because all the WBB have either escaped or died before reaching many of the endpoints in the pathways.

The two items associated with WBB are wood packaging and furniture. They move from the ports to transitional facilities, and in the case of furniture, to furniture shops and end points. Hence, WBB is mainly exposed in city areas.

Again, focusing on the Auckland region (see Figure 20), the area with the highest exposure rate is the port. Additionally, these numbers drop off with increasing distances from the port. There appears to be greater variability amongst regions within Auckland than is the case for AGM.

Figure 19: Map of the aggregated exposure rates for WBB across furniture and wood packaging for the North and South Islands.

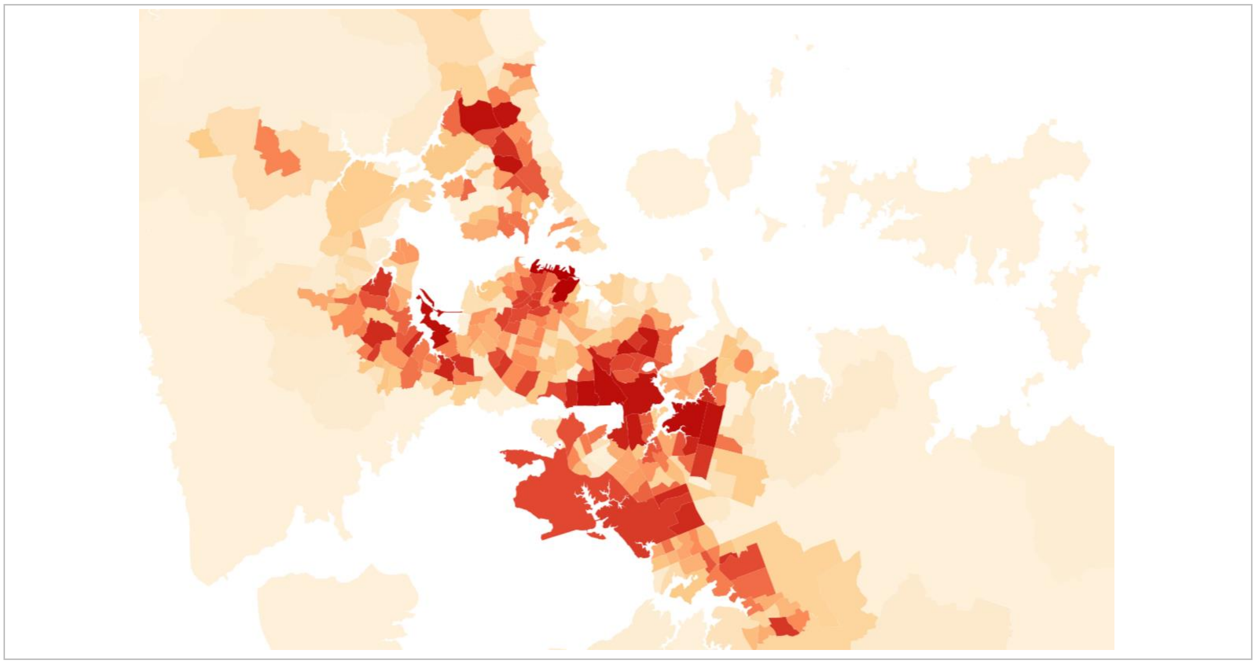

Figure 20: Map of exposure rates for WBB aggregated across furniture and wood packaging for the Auckland region.

6.1.3 Phytophthora

As noted previously, we have very little solid knowledge about Phytophthora, which creates various challenges in producing a model for it. Indeed, the uncertainty is sufficiently great that it is unclear how best to even measure or describe the quantity of the pest – whether that be a count of oospores (or some other life stage) or a measure of mass or some other measure. As noted in Section 3.2.4, since almost any sample of soil from a region will contain the Phytophthora from that region, we have tried to sidestep this issue by focussing instead on the distribution of soil. Essentially, this is akin to treating soil as the (substitute) pest. If we gain a better idea in the future of how much Phytophthora any unit of soil from a source country will likely contain, we can very easily add that to the model through simple multiplication. Note that "exposure" here can be read as referring to the escape of 1 gram of soil containing live phytophthora. Note also that the mortality rate of Phytophthora, for which we have reasonably good estimates, is applied to the soil – so the model only tracks "live" soil.

Figure 21 shows the aggregated exposure maps for Phytophthora. The exposure rate across the country is again generally low, except around ports and population centres. However, there are also a number of areas inland and in non-urban areas that also contain pockets of high exposure rates. Exposures are quite concentrated, more similar to WBB than AGM, with the top 100 area units accounting for 91% of all exposures. Interestingly, this is entirely due to the container pathway. If we omit the container pathway, then the top 100 area units only account for 35% of all exposures. This is likely due to air travellers (mainly visitors) being far more likely to visit areas outside the main cities.

The sea containers pathway here is by far and away the most significant contributor to the number of exposures. The passenger and returning resident pathways account for about 0.6 million exposures per year across all of New Zealand, while the container pathway alone accounts for about 8.8 million total exposures per year (i.e. 8.8 tonnes of live soil). These results are *highly* dependent on the grams of soil estimated to be on containers and people. The container pathway is so significant because the average number of grams per container entering the country has been estimated at 160, against an average of only 1 gram or so for air travellers. (Both estimates were derived from the elicitation sessions.)

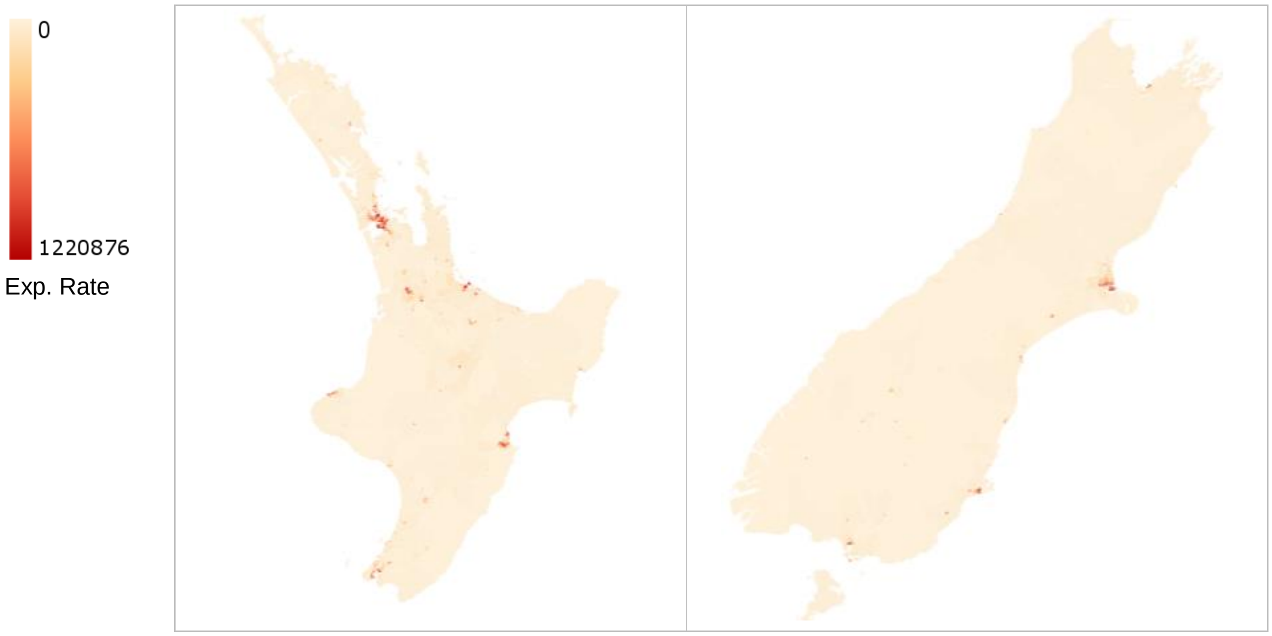

Figure 21: Map of the aggregated exposure rates for Phytophthora across sea containers, returning residents and visitors for the North and South Islands.

Figure 22 shows a close-up of the Auckland region for the aggregated results for Phytophthora. We see that the port (darkest red polygon) is the area subject to the highest exposure rate. The airport is located just to the south-south-west of the centre of this picture. Despite the airport being an important entry point for Phytophthora (due to returning residents and visitors), the exposure rate is not especially high. This is likely because, while it has a relatively large absolute number of exposures, they are spread over a wide enough area that its exposure rate is only slightly higher than for surrounding areas.

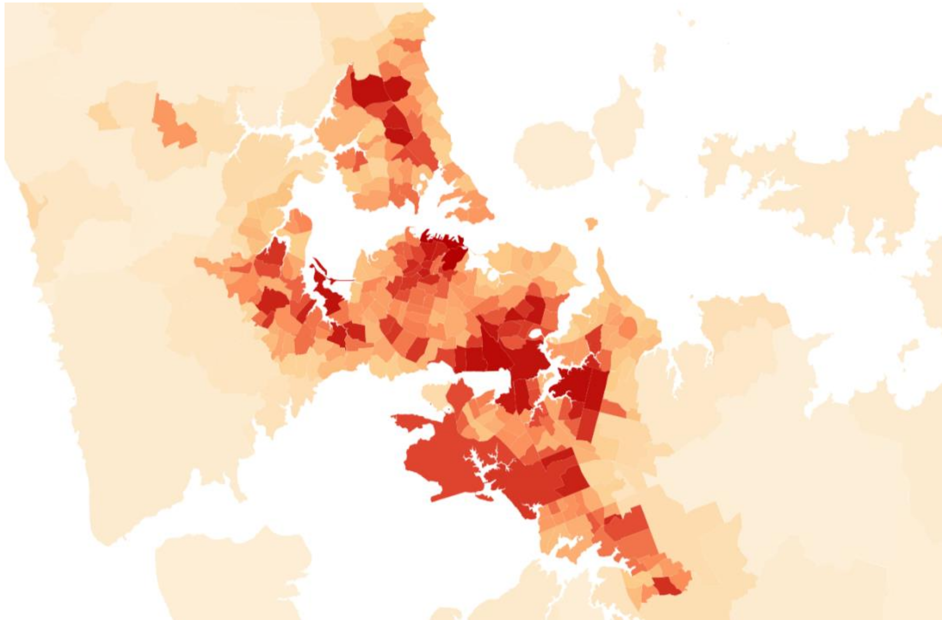

Figure 22: Map of exposure rates for Phytophthora aggregated across sea containers, returning residents and visitors for the Auckland region.

6.1.4 Fusarium

Similar to Phytophthora, we see the aggregated exposures appearing in more locations across New Zealand (Figure 23). In this case, it also translates to a wider spread of the pest, with the top 100 area units accounting for 45% of all exposures. While Fusarium makes use of both the WBB and Phytophthora escape models, the maps are largely dominated by Phytophthora-like exposure rates. In this case, returning residents is the most significant pathway, accounting for ~240,000 exposures. Visitors is the second most significant pathway, accounting for ~190,000 exposures, while sea containers comes in third, with ~34,000 exposures. Used machinery accounts for just 8 exposures on average, due to the very low estimated beetle infestation rate (relative to soil infestation rate) combined with the very low number of used machinery items entering the country. If soil were to be taken into account on the used machinery pathway, this number would almost certainly increase by a large amount.

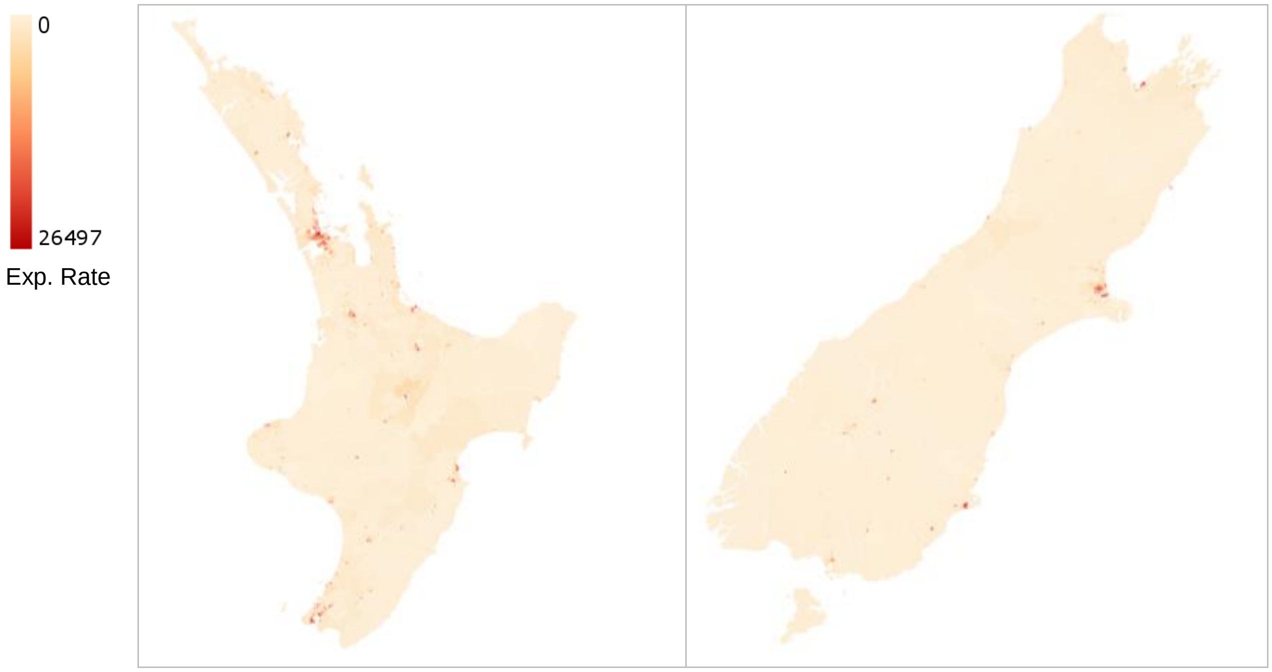

Figure 23: Map of the aggregated exposure rates for Fusarium across used machinery, sea containers, returning residents and visitors for the North and South Islands.

Figure 24 shows the aggregated exposure rates across the Auckland area. We see here that the concentration of exposures is very broadly spread out across Auckland, more so than for any of the previous pests. Despite air travellers being the most significant pathway for Fusarium, the region around the airport does not display an especially high exposure rate relative to the sea port.

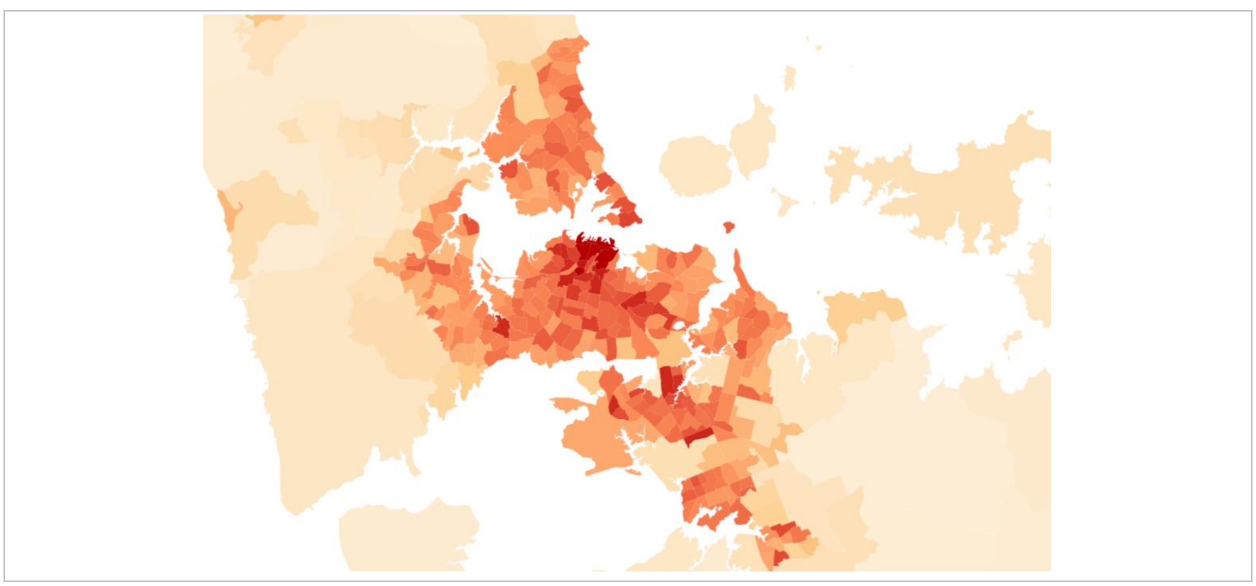

Figure 24: Map of exposure rates for Fusarium aggregated across used machinery, sea containers, returning residents and visitors for the Auckland region.

6.1.5 PSM

Figure 25 shows the aggregated results for the PSM. Exposures are spread somewhat differently than for earlier pests, with a moderate number of high exposure areas. The top 100 area units account for 74% of all exposures – however, in this case, this misses much of the story, as around 30% of all area units contain no exposures at all. (By contrast, all area units exhibit some number of exposures for AGM.) The total number of exposures is very low – around 20 compared to ~1,000 for AGM. This is due to the much lower number of pests per infestation (between 1-5 larvae for PSM versus 1,000 on average per egg batch for AGM) on items entering the country.

While the PSM as a pest is of course more similar to the AGM than it is to Fusarium, the item pathways are much more similar to Fusarium. As noted, it is the item pathways that tend to influence the resulting exposure maps the most. To an even greater extent than for Fusarium, there are high exposure rates scattered around the Hamilton-Tauranga-Taupo region (Figure 26), which largely results from the exposures due to returning residents and visitors.

Figure 25: Map of the aggregated exposure rates for PSM across used machinery, sea containers, returning residents, visitors and live plants for the North and South Islands.

Figure 26: Map of exposure rates for PSM aggregated across used machinery, sea containers, returning residents, visitors and live plants in and around the Hamilton-Tauranga-Taupo region.

6.1.6 Generic

Figure 27 shows the results for the Generic pest combined with all possible items. The exposure maps for the Generic pest follow the main pattern as seen previously, with urban and port areas being the areas of greatest concern. As with PSM, the areas around Auckland, Hamilton and Tauranga show more exposures relative to other areas across New Zealand, but Taupo does not figure so substantially. It’s likely that the accumulation of exposures at sites common to multiple pathways (such as transitional facilities and high population end points) has become a significant factor in these maps.

Figure 27: Map of the aggregated exposure rates for Generic across used vehicles, used machinery, sea containers, sea vessels, returning residents, visitors, furniture and wood packaging for the North and South Islands.

Focusing on the Auckland region (Figure 20), we see that the port and, in particular, airport are subject to very high exposure rates. However, there are a significant number of exposures nearby as well, especially in the north-west of Auckland.

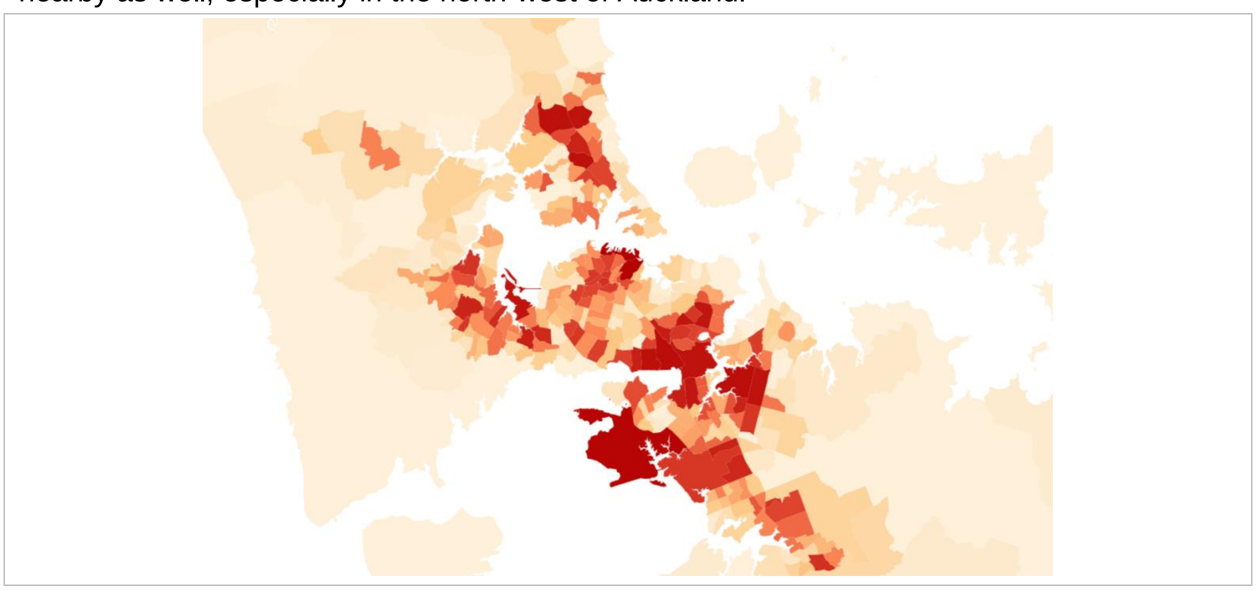

Figure 28: Map of exposure rates for Generic aggregated across used vehicles, used machinery, sea containers, sea vessels, returning residents, visitors, furniture and wood packaging surrounding the Auckland region.

6.2 Aggregated Exposure Results

The aggregated exposure results include all pests and all their associated items (see Appendix C), not including the Generic pest. Each unique pest-item pathway has been assigned equal weight. While this can be changed from within Spear to any custom weighting, we use equal weights here for simplicity.

This map attempts to give an overall sense of the exposure risk across the country, for all pests and items. Strictly speaking, this is incorrect and (as mentioned in the Stage 1 report), extreme caution should be used in interpreting the following map. Firstly, risk is a quantity that combines probability and utility (often called likelihood and consequence in this context). This map shows only probability. Secondly, an "exposure event" for each pest is defined differently, depending on the pest. For example, an exposure event for AGM occurs whenever an egg that has entered the country successfully hatches and separates from its item. In contrast, an exposure event as currently defined for Phytophthora occurs whenever some gram of live soil becomes separated from its item. These are clearly not the same kind of event. Most importantly, the consequences of the two events are very different. However, the two events are treated the same for the purposes of the following map. The weights available in Spear provide for some flexibility as they allow one to essentially specify the relative consequences of an exposure event for each pest (or even each pest-item combination).

With that warning in mind, Figure 29 shows the aggregated exposure map for New Zealand. The map seems to most closely resemble the Phytophthora exposure map, which is unsurprising for reasons described shortly. The top 100 area units account for 89% of all exposures, suggesting the spread is similar to that for Pythophthora and WBB. The total number of exposure events across New Zealand is 9.8 million, making it clear that Phytophthora (which had some 8.8 million exposure events) contributes the largest number of events. Interestingly, the results are quite different to those for the Generic pest across all items – underlining that the two should not be treated as interchangeable.

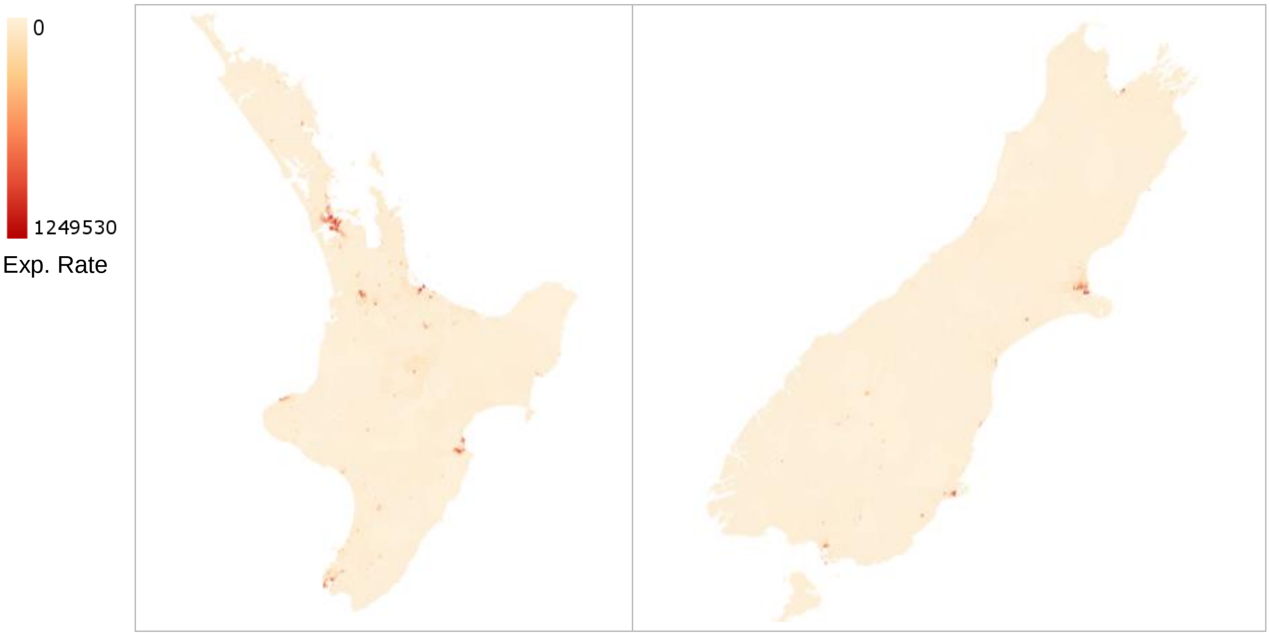

Figure 29: Map of exposure rates for all pest-item combinations for the North and South Islands.

If we zoom into the Auckland region (Figure 31), we can see that the pattern of exposure rates is again very similar to that for Phytophthora, with most exposures occurring at the port, but a significant number also occurring in the north and south-east.

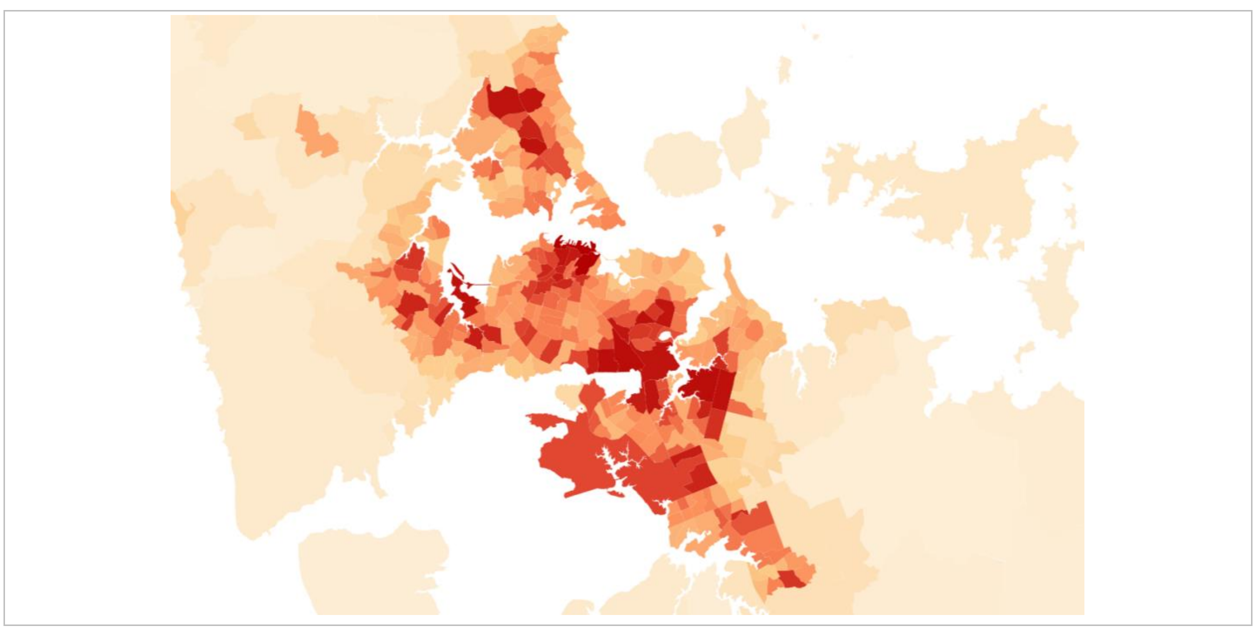

Figure 30: Map of exposure rates for all pest-item combinations surrounding the Auckland region.

6.3 Comparison of Monthly and Yearly AGM Models

We use the AGM – Used Vehicle pathway to compare the use of monthly versus yearly resolution for the models. This involves creating two models: a yearly version of the AGM model, that contains no monthly factors; and a monthly version, that does. In particular, in the latter version the influence of month on the average maximum and minimum temperatures and egg age have been included. To be clear, the monthly model places prior distributions over those monthly factors (and hence does not make use of monthly specific data). Prior distributions were taken from the literature for the respective source countries (in the case of egg ages) and from average weather data (from New Zealand's MetService) for average maximum and minimum temperatures.

Figure 31 shows the results of the runs on the two models. On the left, the yearly model shows increased exposure rates compared to the monthly model (right) where the exposures are at medium to high levels. It appears that the monthly model predicts a much greater spread of exposures. We can infer that the average minimum and maximum temperatures from the yearly model tend to lead to an overestimate in the number of exposures, which would cause more of the pests to escape *earlier* in their pathways, making the exposures more concentrated. By contrast, if the norm is for temperatures to be too low, AGMs will take longer to develop, and be more likely to arrive at an endpoint (and thus be more spread). The top 100 area units account for 43% of the exposures with the monthly model (as compared to 69% with the yearly model).

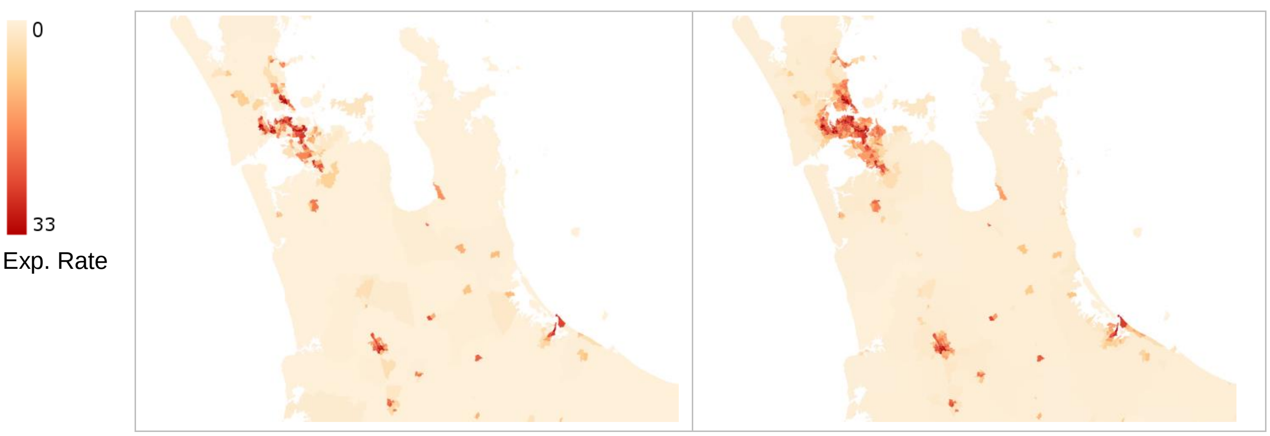

Figure 31: Expected exposures for the AGM - used vehicles yearly (left) and monthly (right) pathways.

6.4 Comparison of the AGM and Generic Pests (Used Vehicles)

The AGM escape model is a moderately complex model of AGM behaviour. It is entirely possible that a simpler model, like the Generic model, would work just as well for the purposes of producing our exposure maps. Hence, we examine the difference between the AGM and the Generic pest model on the used vehicles pathway to determine if the differences produced by the two escape models warrant the more complex AGM model.

Before generating the exposure maps, the (initial) parameters for the Generic model were altered such that the conditional probability of escape at site (and survival) given time at site is as close as possible to the AGM model, when TimeAtSite is set to 2.5 days. There are two important points to note:

- The probability of escape (or survival) does not match the AGM model very well when TimeAtSite is *very* different to 2.5 days – for example, at 100 days, the AGM model gives a probability of escape of 0.66 versus the generic model's probability of 0.77. For the most part, however, there is a very good fit for values of TimeAtSite in the range $[0,80]$.
- It is only because we have created a more detailed BN for AGM that we can properly specify the relationship between TimeAtSite and the probability of escape (survival) in the Generic model. In the absence of such a model, we would need to estimate the relationship in some other way.

Figure 32 shows a map of the expected number of exposures for the AGM (left) and Generic pest (right). In general, the maps match each other very well. The main difference is in populated areas (e.g. Auckland), where the Generic model predicts somewhat greater exposure rates. This is highly likely due to the different responses of the models to large TimeAtSite values.

Figure 32: The expected number of exposures for AGM - used vehicles (left) vs. Generic - used vehicles (right).

6.5 Comparison between original and post-elicitation parameters

In order to determine the impact of the elicitation, a complete set of runs of all possible item-pest combinations were conducted with both pre- and post-elicitation parameters.[25]

[25] Appendices A-C are particularly useful for the following section. Table 4 and Table 5 in particular contain the pre-elicited and post-elicited values for each of the parameters that were addressed in the elicitation. Note that, in general, the average of the round 1 and round 2 averages were used.

This is not a perfect test of how the elicitation changed the exposure maps – some of the new parameter values that were chosen after the elicitation are not just summaries of elicitation answers, but also included (non-elicited) changes that were needed in order to address some of the issues that were raised. The most important example of this is Phytophthora. Note also that the newer models (i.e. PSM, Fusarium and Generic) were not settled at the time of the elicitation, and so their pest-specific parameters were not directly addressed by the elicitation questions.[26] We have therefore not included these models in the comparison.

Table 2 contains (for various pest-item combinations) the average number of exposures across all area units for both the pre- and post-elicitation parameters. Of the three pests analysed here, the exposure numbers for all pests have increased after elicitation. The exposure numbers for AGM have changed the least (with an average log ratio of 0.7), and those for Phytophthora have changed the most (with an average log ratio of 4.0). The AGM – Used Vehicles unique pathway underwent the smallest change (log ratio of 0.3), which is unsurprising given it is the most stable of all the models. There was a moderately large change for the average exposure numbers for AGM – Sea Containers (log ratio of 1.1). This is largely due to the change in infestation rate, which was originally set at ~0.0008 (based on interception data, which was likely an underestimate) and was changed to ~0.004 after the elicitation.

Note that changes in the exposure numbers do not necessarily affect our conclusions, particularly if the changes are uniform across area units. For example, if we consider a single unique pathway (such as AGM – Used Vehicles) and just increase the infestation rate, the exposure numbers will increase, but the order over the area units, and possibly even the ratio of exposures in one area unit to those in another, may not change. We summarise this by saying that, while the *absolute* exposure numbers may change, the *relative* exposure numbers may not change. We will look at this again in Section 6.5.1 below.

Returning to the table, we see that there are fairly large differences in the pre- and post-elicitation average exposures for the WBB pathways, with the new parameters resulting in higher exposure rates for both pathways. There is an especially large change for the furniture pathway. This is mostly due to increases in the infestation rates (of about 10 times for both pathways) and the very large drop in the detection rates at the port and the transitional facilities (from 85% originally, down to about 5% after the elicitation). The detection rate numbers were mainly used as placeholders in the Stage 1 model (reused from AGM for all pests), so the large change there is no surprise. By contrast, the original infestation rates were based mostly on literature, however since the information in the literature did not fully cover what was required, some extrapolation was also necessary. Given the range of possibilities is best measured on a logarithmic scale, the difference in the infestation rates before and after elicitation is relatively small (see Q2 & Q3 in Table 5, Appendix B). Nonetheless, these new infestation rates combined with the lower detection rates produced a substantial increase in the total number of exposures (with an average log ratio of 3.0).

[26] We hope to be able to address these in a later elicitation session for at least PSM and Fusarium.

Phytophthora played host to the largest changes in exposure numbers on average. While the changes to the two air traveller pathways could be considered more modest, the sea container pathway saw an extremely large increase in the expected number of exposures (with a log ratio of 6.6). This is due to a change in several of Phytophthora's parameters that were made as a result of the elicitation (either directly or indirectly), but perhaps most significantly the change in the infestation rate (to on average 100 grams of soil per container, from a previous setting of 1 unit of infestation per container).

Table 2: The average exposure count (and standard deviation) across area units for pre- and post-elicited parameters. The log ratio column is the log (base 10) of the ratio of post-elicitation to pre-elicitation parameters. Rows with a substantial difference (log ratio > 2) have been highlighted.

Item	Pest	Pre-Elicitation		Post-Elicitation		Log Ratio (Post to Pre)
		Average	SD	Average	SD	
Sea Containers	AGM	2.011	15.090	27.388	188.531	1.1
Sea Vessels	AGM	1.049	1.629	5.035	7.656	0.7
Used Vehicles	AGM	0.183	0.363	0.388	0.937	0.3
Returning Residents	Phytophthora	0.683	1.362	163.702	327.365	2.4
Sea Containers	Phytophthora	0.002	0.026	8485.629	111943.646	6.6
Visitors	Phytophthora	0.124	0.597	133.793	646.637	3.0
Furniture	WBB	0.000	0.001	1.475	34.935	4.4
Wood Packaging	WBB	0.056	0.790	1.912	27.366	1.5

Table 3 shows how the exposure numbers change on average at the level of the area unit. Specifically, the table displays the average of the percentage difference from the new data to the original data across all area units. In some cases, this may show differences that we may not catch in the first table. For example, if the absolute number of exposures does *not* change, but the relative location of exposures *does*, we would not see any difference in the first table, but would in the second. Conversely, if the number of changes is very significant for only a few area units, then we will likely miss seeing this in the second table, but catch it in the first.

For the most part, Table 3 looks much like Table 2. However, one notable change is that the WBB – Furniture unique pathway is now showing virtually no difference between pre- and post-elicitation. The reason is that most of the area units show no level of exposure (as noted in Section 6.1.2), and that has not changed between the pre- and post-elicited results.

Table 3: The average percentage difference (plus standard deviation) from the new to the original data across all area units. Rows with a substantial difference have been highlighted.

Item	Pest	Perc. Diff.	SD
Sea Containers	AGM	13.546	1.139
Sea Vessels	AGM	3.184	1.490
Used Vehicles	AGM	0.700	0.727
Returning Residents	Phytophthora	236.049	65.648
Sea Containers	Phytophthora	21314328.541	393251221.937

Visitors	Phytophthora	1074.058	249.121
Furniture	WBB	0.167	2.374
Wood Packaging	WBB	672.623	11192.067

6.5.1 Further Investigation of Specific Pathways

Given the differences between the original and new results, we now pick a couple of example unique pathways to investigate further. In particular, we choose AGM – Used Vehicles (the most robust unique pathway) and Phytophthora – Visitors (one of the least robust pathways, and the one that exhibited the second greatest amount of change).

6.5.1.1 AGM - used vehicles

Figure 33 shows a map of the rank percentage difference of the exposures from the new to the original maps for AGM – Used Vehicles. The rank percentage difference is the percentage change in an area unit's rank in terms of expected exposures.[27] If we take note of the scale, we can see that the ranks over area units do not change overly much – we do not see any area unit move up or down by more than 3.1% of all places, and most move much less than that.

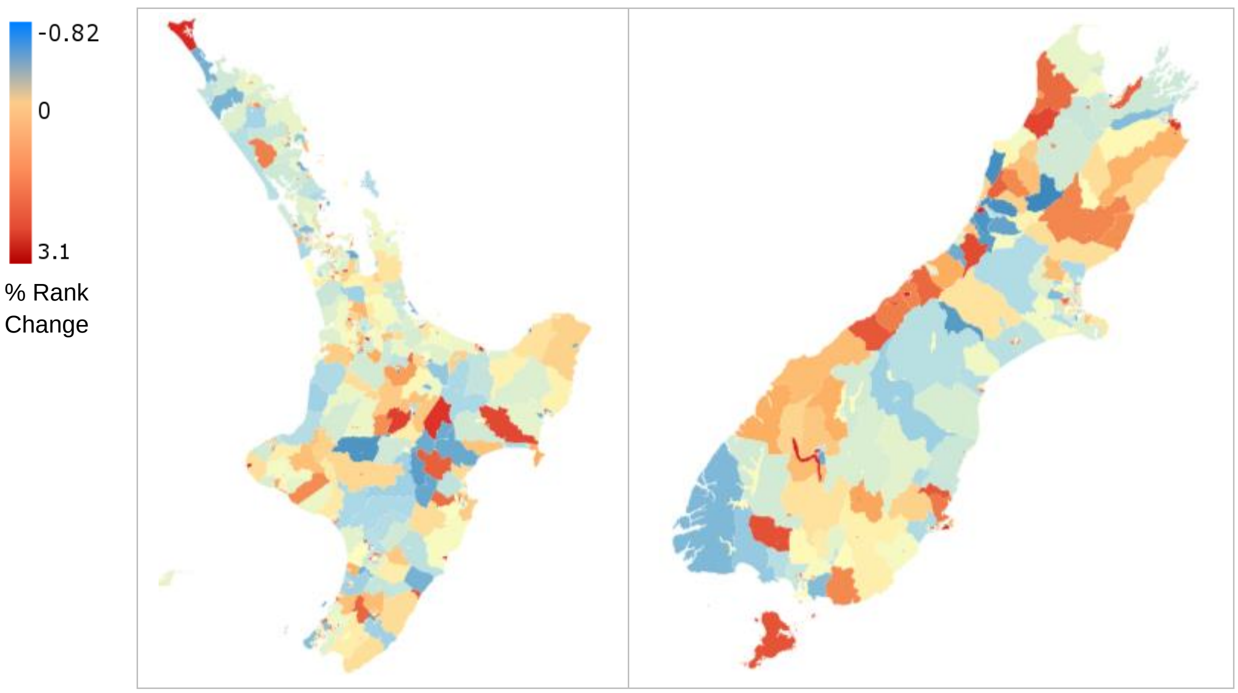

Figure 33: Map of the percentage difference between exposures of the new and original models for AGM – used vehicles for the North and South Islands.

If we consider the Auckland region (Figure 34) and overlay the various pathway points (such as registration sites and car yards), we can generally see that there are less significant changes

[27] So the area unit with the highest number of exposures would receive rank 1, the next highest, rank 2, etc.

near the intermediate pathway points (i.e. ports, car yards and registration sites). Larger changes tend to have occurred in non-urban areas.

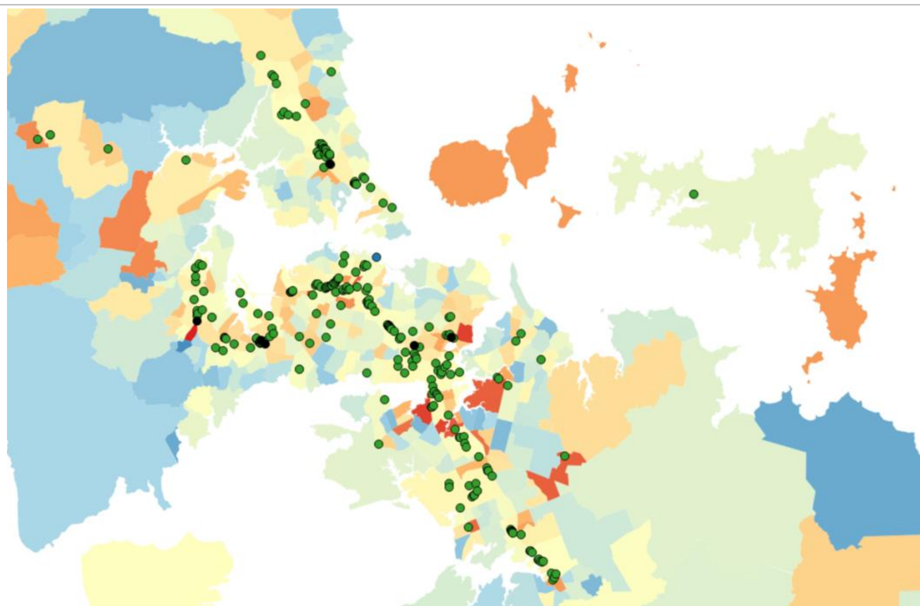

Figure 34: Map of the percentage difference of exposures for AGM – used vehicles focusing on the Auckland region. The pathway points including port (blue circles), car yards (green circles) and registration sites (black circles).

6.5.1.2 Phytophthora – visitors

Figure 35 shows a map of area unit rank changes for Phytophthora on the Visitors pathway. The story here is a fair bit different to AGM, as we might expect. Some area units show very large changes in their rank indeed. As with AGM, these tend to be in non-urban areas. In addition, the changes in urban areas tend to all be on the negative side. Almost all the area units in Auckland, for example, have dropped slightly in rank (Figure 36), which is also mostly true for areas such as Hamilton, Christchurch and Wellington.

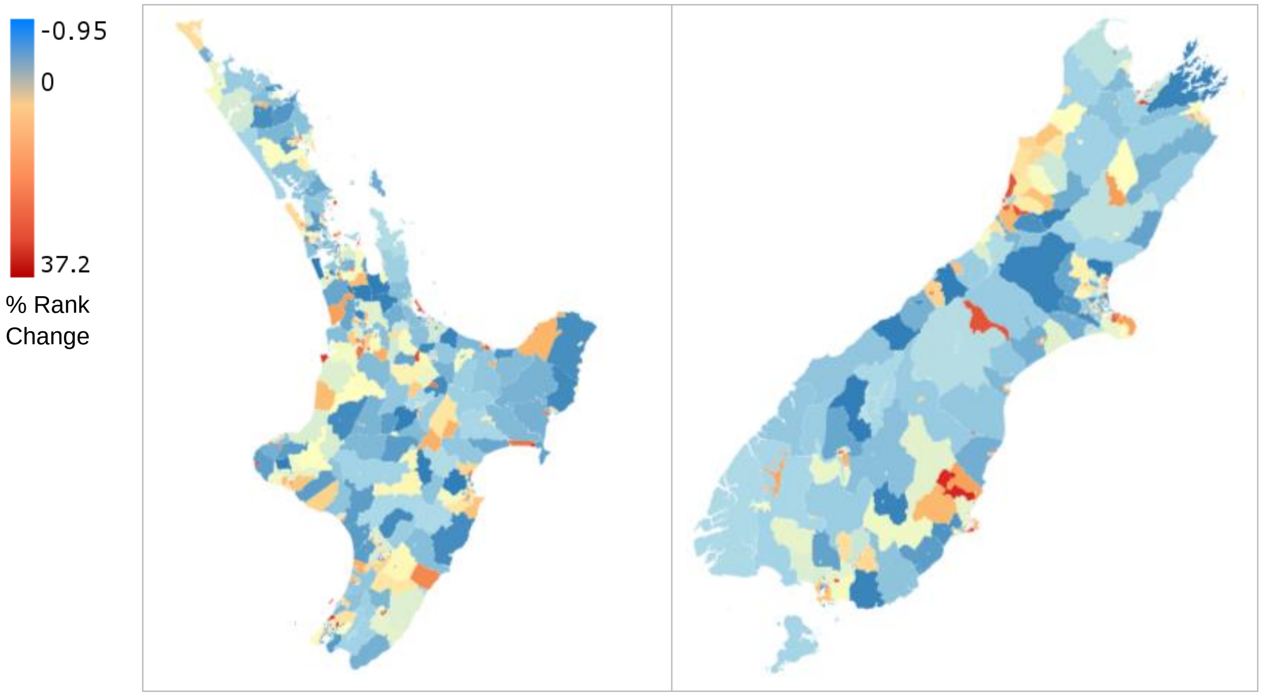

Figure 35: Map of the percentage difference between exposures of the new and original models for Phytophthora – visitors for the North and South Islands.

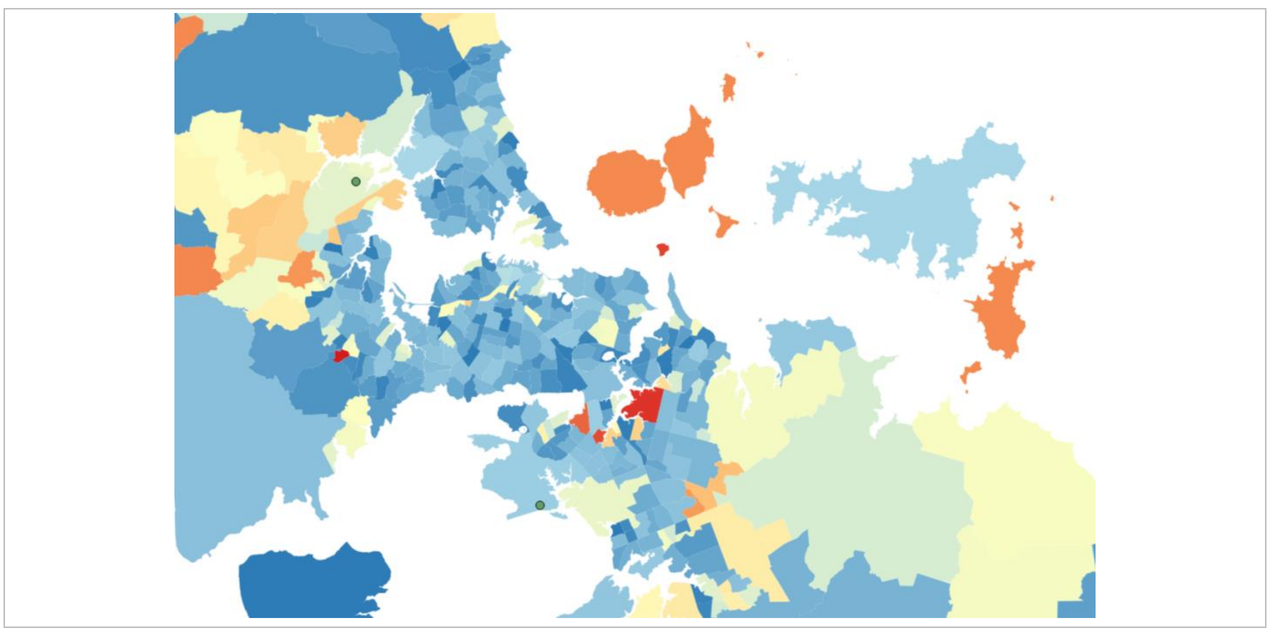

Figure 36: Map of the percentage difference of exposures for Phytophthora – visitors focusing on the Auckland region. The pathway points include airports (green circles).

6.6 Summary

The pest-specific aggregated results show a substantial amount of consistency between pests. The most highly dispersed pests appear to be Phytophthora and Fusarium, as exposures often only ever occur at the end points of these pathways. AGM and PSM produce similar results, though on very different scales, while WBB exposures are highly concentrated around pathway points.

The aggregated results tell a story that is quite consistent with all of the individual pest-specific exposure maps. However, care must be taken when interpreting these results. In particular, Phytophthora dominates due to the much larger number of exposure events that occur with this type of pest than with others. At the very least, an appropriate pest or pathway weighting should be used when creating aggregated results in future.

We also made a series of comparisons between versions of the model that reflected different changes made in Stage 2. These included comparisons of the monthly and yearly model, of the AGM against a suitably tailored generic pest and of the original parameters versus the post-elicited parameters.

A comparison of the AGM – Used Vehicles yearly and monthly models showed that the yearly model often overestimates the exposures at the pathway points, leading to a greater level of concentration around pathway points as compared to the monthly version of the model. This suggests the monthly model will generally be more accurate than the yearly model, particularly if appropriate monthly data can be incorporated.

The Generic model fared remarkably well against the more complicated AGM model. This bodes well for modelling less well-understood pests in future. However, there is a limitation: the key parameters for the Generic model could only be determined here because we already had the AGM model in hand. This is obviously an obstacle for handling other pests this way.

The comparison of the pre-elicited and post-elicited models showed the value of the elicitation sessions. While the results for AGM did not change substantially, particularly when the relative results (i.e. area unit ranks) were considered, WBB and, especially, Phytophthora did show some notable changes.

7 Conclusion and Future Directions

The Stage 2 model is now in a much improved state over the Stage 1 model. Additional pests have been added, including a generic pest that can be used for modelling all pests at an abstract level, or as a simple model for specific pests. All of the placeholder parameters have been replaced with expert estimates, and many of the key parameters have been updated based on the results of the expert elicitation. Despite significant changes to many aspects of the model, the results have proven to be relatively robust. While the absolute number of exposures has undergone quite a bit of change in some cases, the *relative* exposure rates remain quite similar to the results of the Stage 1 model, resulting in exposure maps that share the key features with their Stage 1 counterparts. The Phytophthora model is a notable exception. The model is subject to significantly more uncertainty and variability than is the case for the other pest models and so less faith can be put in its results. However, as with other models, the item pathway factors dominate and, as a consequence, the exposure maps do not deviate as much as one might imagine.

Alongside the Stage 2 model improvements, Spear has been created to make it easy to run and re-run exposure models with different parameter sets, as well as to keep track of past model runs. The software constitutes a key outcome of the project, since the threats from invasive pests is evolving over time and will require a model that can be re-run easily as new information comes to light.

There is still much that can be done to improve the exposure model. While many of the parameters have been refined through the elicitation sessions, further improvements to parameters is still a key area for future work. The data used by the model for item pathways is already of quite a high quality, however if reliable monthly data were to become available, that may be useful to incorporate. Pest model structures could of course be improved in many ways, however, based on the lesser impact of the pest models on the final exposure maps, such improvements should come later than more critical improvements, such as those to item, infestation and detection parameters. The models, of course, need a great deal more in the way of validation; we expect MPI to conduct assessments of the system in parallel with and against its existing systems, but a formal validation should also be conducted in future.

A key avenue for improving the model, of course, is through sensitivity analysis. We have performed one-at-a-time and two-at-a-time analyses on several parameters within submodels, such as the pest models and the pathway point models. These showed in particular some interesting effects in the AGM and Phytophthora escape models. Exploring more complex interactions can provide further insights, which is only achievable with a global sensitivity analysis technique such as variance-based sensitivity analysis. This analysis is left to explore in future. In addition, a full model sensitivity analysis may be especially worthwhile, though making such an analysis tractable is a major difficulty. If the aim was solely to set limits on a plausible range of possible outputs, a simple approach would be take advantage of the BN structure and place uncertainty distributions over parameters that are currently specified as constants. However, we would also like to identify influential parameters, and this is more difficult. Nonetheless, we believe it would be well worth the effort and would provide greater insight into the robustness of the model.

Finally, the software is now ready for serious use, but still needs to be evolved significantly to adapt to the needs of the forest health surveillance system. There are many ways in which the software could be extended to make model runs and exploration simpler, and in particular to allow for simpler data and parameter updates that can be managed entirely within the software, but nonetheless easily exported and integrated into other software and systems. Integration with the cost optimisation model unfortunately fell just out of reach for this project, but some form of integration is necessary in future to ensure the integrity of the overall forest health surveillance model. This applies not just at a software level, but also at the level of the underlying model. In particular, uniting the two models brings the system closer to a full and proper risk analysis.

8 Glossary

area unit	A geographical unit (polygon) defined by Statistics NZ. Area units are formed from aggregations of mesh block polygons. There are approximately 2,000 area units that cover New Zealand.
Bayesian network (BN)	A graphical representation of the probabilistic influences that hold between a set of variables.
exposure	An escape of a single pest (or unit of pest) from an item into the environment.
exposure rate	The number of exposures per square kilometre per year (for some geographical area).
item	Whatever a pest can be attached to and transported on. In this report, that can be a used vehicle, a container, a vessel, wood packaging, wooden furniture or a passenger.
mesh block	The smallest geographical unit (polygon) defined by Statistics NZ. There are approximately 47,000 mesh blocks that cover New Zealand.
model	In the context of this report, a model always refers to a BN.
pathway	A pathway is a set of points along which an item and pest can travel.
pathway point	A set of points that make up a pathway. A point can be a type of location that is mostly well-defined (e.g. car yards) or locations that are less well-defined (e.g. the path between the port and the registration site)
pest	Any undesirable organism or insect that is capable of entering the country.
risk	The *combination* of probability and (negative) consequences. For instance, strongly negative consequences can be low risk if they have a low enough chance of occurring.
unique pathway	A *unique* pathway is a combination of a specific pest (e.g. AGM) and an item pathway via which that pest travels (e.g. used vehicles through ports, registration sites, car yards and end points).

9 Data Dictionary

Most of the variables that appear in the Stage 2 models also appear in the Stage 1 prototype. Definitions for these variables were first given in the Stage 1 report, but are reproduced here for convenience. New variables or variables with updated definitions have been highlighted in pink.

9.1 Pest arrival

Node	Type	Formula	Node Description
UnitInfestationRate	Input	EpidemicLevel*ExportMonth	Values: $[0,1]$ Unit: <proportion> The infestation rate (per item unit) on arrival The avgRate is calculated by multiplying the rate of infestation in each source country, by the proportion of all units arriving from that country. The sd specifies noise around this baseRate, which is typically calculated as 10% of the magnitude of the baseRate for the prototype.
NumberUnits	Input	Normal(<avgUnits>,<sd>)	Values: $[0,\infty]$ Unit: quantity/year The number of item units that arrive per year. Units can be whatever is appropriate for the pest and item (e.g. number, mass, volume, etc.). sd is the variation around this.
PestsPerInfestation	Input	Typically: Normal(<avg>,<sd>)	Values: $[0,\infty]$ Unit: pest quantity/unit quantity The expected number of pests per infested unit In the case of larval batches, this would be the distribution over the number of eggs per batch. In other cases (such as adult wood borers) this might be just 1 or some low number plus noise.
NumberUnitsInfested	Interm.	NumberUnits*UnitInfestationRate	Values: $[0,\infty]$ Unit: quantity/year The expected number of units with infestations This is a deterministic multiplication of the parent nodes.
PestQuantity	Output	PestsPerInfestation*NumberUnitsInfested	Values: $[0,\infty]$ Unit: quantity/year The expected number of pests that arrive at the pathway entry point This is a deterministic multiplication of the parent nodes. It is the output node for this submodel, and its distribution is copied to the PreviousPestQuantity node in the Pathway Point

			submodel.
ExportMonth	Input	Uniform(1,12)	Values: [1,12] Unit: month The month of export for the item Depending on the export month, the infestation rates can be quite different. (Month-based nodes are not currently used in the Spear software.)
Month	Interm.	(ExportMonth+TransitTime<12 ? TransitTime+ExportMonth : TransitTime+ExportMonth-12)	Values: [1,12] Unit: month The month of arrival of the item Essentially accounts for transit time on top of the export month. (Month-based nodes are not currently used in the Spear software.)
TransitTime	Input	Typically: <number>	Values: [0,34] Unit: months Transit time from the source port Time take for the item to arrive from the source port. It is measured in days and the large values (e.g. 34) is generally the time required to arrive by sea. (Month-based nodes are not currently used in the Spear software.)
(EpidemicLevel)	Input	Typically: <number>	Values: [0,∞] Unit: <scalar> The level of infestations expected in the current year relative to the model average Currently unused node that may be dropped. Could be used to temporarily scale the infestation rate from year to year.

9.2 Treatment

Node	Type	Formula	Node Description
DetectionRate	Input	Typically: <number>	Values: [0,1] Unit: <proportion> The proportion of pests that are detected Typically, this is just a point estimate, but can also be a truncated normal or beta distribution.
TreatmentRateForUndetected	Input	Typically: <number>	Values: [0,1] Unit: <proportion> The proportion of all undetected pests that nonetheless get treated This can be used when treatment occurs in the absence of inspection, or when ‘over-treating’ occurs (for example, if a detection leads to an entire ship’s cargo being treated).
TreatmentEfficacy	Input	Typically: <number>	Values: [0,1] Unit: <proportion> The proportion of pests that are truly eradicated when an attempt is made to eradicate them
TreatmentRate	Interm.	DetectionRate+(1-DetectionRate)*TreatmentRateForUndetected	Values: [0,1] Unit: <proportion> The proportion of pests to which treatment is actually applied It is assumed that all detected pests are treated (whether or not successfully) and this is combined with the proportion of undetected pests that are treated.
ProportionTreated	Output	TreatmentRate*TreatmentEfficacy	Values: [0,1] Unit: <proportion> The proportion of pests *successfully* treated This is calculated from the rate of treatment multiplied by the efficacy of the treatment.

9.3 Pest Activity (Escape Models)

Node	Type	Formula	Node Description
TimeAtSite	Input	Typically: Normal(<number>,<sd>)	Values: $[0, \infty]$ Unit: days The time spent at the site This is the amount of time that the pest spends at a particular pathway point (at the port or the transitional facility, for instance).
ProportionThatEscapeAtSite or EscapesAtSite	Output	<varies>	Values: $[0,1]$ Unit: <proportion> The proportion of pests that escape from the site or Values: 0 or 1 (False/True) Unit: Boolean Whether or not the pest escapes from the site A pest model must include either ProportionThatEscapeAtSite and ProportionSurvive (if a population model), or otherwise EscapesAtSite and Survives (if an individual based model) as output nodes.
ProportionSurvive or Survives	Output	<varies>	Values: $[0,1]$ Unit: <proportion> The proportion of pests that do not die due to natural causes or Values: 0 or 1 (False/True) Unit: Boolean Whether or not the pest survives beyond the current site A value of 1 indicates that the population neither declines nor grows. Growth can occur due to conception/birth; for instance, insects may lay eggs and fungi may propagate. Reduction can occur due to mortality. The reduction in this node *does not* include the proportion that escape (which is taken out at another point of the model) nor deaths due to human intervention. A pest model must include either ProportionThatEscapeAtSite and ProportionSurvive (if a population model), or otherwise EscapesAtSite and Survives (if an individual based model) as output nodes.
BaseTemp (AGM)	Input	Typically: <number>	Values: 10 Unit: degrees Celsius

			The baseline temperature used in the computation of degree days The baseline temperature can be different for different pests, whose development can be modelled with degree days.
AvgMinTemp (AGM)	Interm.	Typically: <number>	Values: [4.7,13.6] Unit: degrees Celsius The average minimum monthly temperatures in New Zealand
AvgMaxTemp (AGM)	Interm.	Typically: <number>	Values: [13.1,22.9] Unit: degrees Celsius The average maximum monthly temperatures in New Zealand
DegreeDaysAtSite (AGM)	Interm.	((AvgMaxTemp+Max(BaseTemp,AvgMinTemp))/2-BaseTemp) *TimeAtSite	Values: [0,200] Unit: degrees Celsius An estimate of the degree days spent by pests at this site A revised measure of days based on temperatures to compare with the hatching age. Typically, higher temperatures lead to an increased number of degree days leading to a higher probability of escape at site
HatchingAge (AGM)	Interm.	Typically: Normal(95,30)	Values: [0,200] Unit: days The hatching age of the AGM egg
EggAge (AGM)	Interm.	Trim((Month-RelativeOvipositedMonth)*30*6,0,140)	Values: [0,140] Unit: days The egg age relative to the oviposited month
MaxAge (AGM)	Interm.	Typically: <number>	Values: 130 Unit: days The maximum egg age which is set to be slightly higher than the maximum hatching age
Month (AGM)	Input	Typically: Uniform(0,12)	Values: [1,12] Unit: months Current month of the year (in the model)
OvipositedMonth (AGM)	Interm.	Triangular(5,7,9)	Values: [5,9] Unit: months Month of the year when the AGM eggs were laid
ProportionBalloon (AGM – sea vessels)	Interm.	Trim(Normal(0.012,0.1),0,1)	Values: [0,1] Unit: proportion Proportion of the larvae that balloon to shore

DailyEscapeRate (Generic)	Input	Typically: <number>	Values: $[0,1]$ Unit: probability Probability of escape on a single day This is later combined with TimeAtSite to determine the probability of escape if the item remains at site for TimeAtSite days.
MortalityRate (Generic, Phytophthora, Fusarium)	Input	Typically: <number>	Values: $[0,1]$ Unit: probability Probability that the pest dies (for a single day) This is later combined with TimeAtSite to determine the probability of survival if the item remains at site for TimeAtSite days.
SurfaceExposure (Phytophthora, Fusarium)	Input	Typically: <number>	Values: $[0,1]$ Unit: proportion Average proportion of an item's surface covered in soil This node is highly uncertain. The surface in this case can refer to any part of the item that is exposed to the outside world, and therefore provides a means of escape.
Disturbance (Phytophthora, Fusarium)	Interm. & Input	Typically: Trim(Normal(Ln(6*TimeAtSite+1)/8,0.1),0,1)	Values: $[0,1]$ Unit: <number> How much disturbance the item undergoes This increases as TimeAtSite increases. As disturbance approaches 1, the amount of soil exposed that escapes approaches 100%. This is an intermediate node, but is also intended as an input node, with the input being the disturbance relation to time.
LarvaeAge (PSM)	Interm.	<dependent on month>	Values: $[0,90]$ Unit: degree days The larvae age in degree days The PSM would enter as larvae, mainly between the months of March and April. Hence, the larvae age is included in the BN as opposed to an egg age in the case of the AGM.
MaturationAge (PSM)	Interm.	45*BaseDegree	Values: $[0,\infty]$ Unit: degree days The age of maturation of the larvae The age which the larvae matures, typically at the end of May. Adults emerge within two to three weeks.
AvgTemp (WBB)	Input	Typically:	Values: $[0,50]$

		<distribution>	Unit: degrees Celsius Distribution over average temperatures for New Zealand
AvgLifetimeTemp (WBB)	Input	Typically: <distribution> Normal(18,3)	Values: $[0,50]$ Unit: degrees Celsius Distribution over average temperatures that WBB will have witnessed over the course of its lifetime
LifeSpan (WBB)	Interm.	Normal(41434*AvgLifetimeTemp^-1.298,100)	Values: $[0,\infty]$ Unit: days Life span in days of an individual WBB The formula here was fitted to data for the life spans of anoplophora based on their average life time temperatures
Mortality (WBB)	Interm.	1/LifeSpan	Values: $[0,1]$ Unit: probability Probability that the WBB dies (for a single day)
Activity (WBB)	Interm.	<number>	Activity is a rough indication for how much the beetles move about, from 0 to 1 The relation to temperature is taken from the data collected on Hylurgus activity (hylurgus_activity.csv), from the Scion woodchip project.

9.4 Pathway Point

Node	Type	Formula	Node Description
PreviousPestQuantity	Input	Binomial(n,p)	Values: $[0,\infty]$ Unit: pest quantity The quantity of pests from the previous step Typically, the binomial distribution is copied, either from the Pest Arrival PestQuantity node, or from the previous Pathway Point's NextPestQuantity node
ProportionToHere	Input	<number>	Values: $[0,1]$ Unit: <proportion> The proportion of pests arriving at this particular *location* A location differs to a pathway point, in that pathway points consist of multiple locations. (See the text for an explanation.)
ProportionTreated	Input	<number>	Values: $[0,1]$ Unit: <proportion>

			The proportion of pests eradicated via treatment
ProportionSurvive	Input	<number>	Values: $[0,1]$ Unit: <proportion> The proportion of pests that do not die due to natural causes It is possible that future versions of this model will allow this to account for reproduction as well as death
ProportionThatEscapeAtSite	Input	<number>	Values: $[0,1]$ Unit: <proportion> The proportion of pests that escape from the item at the site
PestQuantityArrives	Interm.	Binomial(PreviousPestQuantity,ProportionToHere)	Values: $[0,\infty]$ Unit: pest quantity The quantity of pests that arrive to this specific location
PestQuantity	Interm.	Binomial(PestQuantityArrives,1-ProportionTreated)	Values: $[0,\infty]$ Unit: pest quantity The quantity of pests that arrive at this location and survive treatment
NextPestQuantity	Output	(PestQuantity-Exposures)*ProportionSurvive	Values: $[0,\infty]$ Unit: pest quantity The quantity of pests that leave this pathway point This is the PestQuantity minus the number of exposures, and reduced by the pests that die.
Exposures	Output	Binomial(PestQuantity,ProportionThatEscapeAtSite)	Values: $[0,\infty]$ Unit: pest quantity The quantity of pests that escape from the item at the site
AnyExposure	Output	Exposures>0	Values: True or False Unit: Boolean Whether or not there is any exposure at this point

10 References

Brown, B.B., 1968. *Delphi process: a methodology used for the elicitation of opinions of experts* (No. RAND-P-3925). RAND Corp Santa Monica CA.

Hanea, A. M. et al. "Classical meets modern in the IDEA protocol for structured expert judgement." Journal of Risk Research (2016): 1-17.

Linstone, H.A. and Turoff, M. eds. The Delphi method: Techniques and applications. Vol. 29. Reading, MA: Addison-Wesley, 1975.

Mascaro, S., 2015, *Forest Health Surveillance Exposure Mapping Prototype*, Bayesian Intelligence

Mascaro, S. & Thiruvady, D. (2016), *Spear: Exposure modelling software for invasive forest pests*, Bayesian Intelligence

Mascaro, S. & Thiruvady, D., (forthcoming), *Spear: Technical documentation*, Bayesian Intelligence

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S, 2008. *Global Sensitivity Analysis: The Primer*. John Wiley & Sons.

Wedde, S., Waite, C., Hyde, N., Whyte, C., 2005. *Used Vehicles Monitoring Survey* (No. BMG/05-06/10). Biosecurity Monitoring Group, Biosecurity New Zealand.

Appendix A Elicitation Questions

A.1 In-person elicitation questions

- Q1a: How many days does a used vehicle spend at the port?
- Q1b: How many days does a used vehicle spend at the registration site?
- Q1c: How many days does a used vehicle spend at the car yard?
- Q2: Consider a single egg mass attached to a sea vessel that is docked at port. (The egg mass is located somewhere on the ship in open air.). What proportion of eggs in the mass would balloon to shore? (0-100%)
- Q3: How many kilograms of soil would be attached to a container's surface when it arrives at a port?
- Q4a: What proportion of the imported soil on a tourist's clothing would fall off at the airport?
- Q4b: What proportion of the imported soil on a tourist's clothing would fall off at a nature-based attraction?
- Q4c: What proportion of the imported soil on a tourist's clothing would fall off at any other point?
- Q5: What is the infestation rate for AGM on used vehicles in countries that have AGM during active months? (I.e. the average number of egg masses per used vehicle during active periods).
- Q6a: What is the probability that a kilogram of soil attached to containers that originates from a country with phytophthora contains phytophthora?
- Q6b: What is the probability that a kilogram of soil attached to a newly arrived tourist's clothing that originates from a country with phytophthora contains phytophthora?
- Q6c: What is the probability that a kilogram of soil attached to a returning resident's clothing that originates from a country with phytophthora contains phytophthora?
- Q7a: Assume the presence of a single AGM egg mass on a used vehicle. What is the probability that this egg mass will be detected at the port?
- Q7b: Assume a single AGM egg mass has been detected on a used vehicle. What is the probability that treatment (including any re-inspection and re-treatment) will be effective in neutralising the egg mass?
- Q7c: Assume a single AGM egg mass is not detected. What is the probability that it will be treated anyway? (For example, due to detection of another egg mass.)

A.2 Online elicitation questions

- Q1: What is the infestation rate for AGM on containers when leaving port in countries which have AGM during active months?
- Q2: What is the infestation rate for wood boring and bark beetles on wood packaging (percentage of tonnes infested) when leaving port from the source country?
- Q3: What is the infestation rate for wood boring and bark beetles on furniture (percentage of tonnes infested) when leaving port from the source country?

- Q4: On average, how many grams of soil would be attached to a tourist's clothing and equipment on entering the country?
- Q5: Assume the presence of a *single unhatched* AGM egg mass on a used vehicle. What is the probability that this egg mass will be detected at a *registration site*?
- Q6: Assume the presence of a *single unhatched* AGM egg mass on a used vehicle. What is the probability that this egg mass will be detected at a *car yard*?
- Q7: Assume the presence of a *single unhatched* AGM egg mass on a used vehicle.
  What is the probability that this egg mass will be detected at any time *after* the used vehicle has left the car yard?
- Q8: Assume the presence of a *single unhatched* AGM egg mass on a container.
  What is the probability that this egg mass will be detected at the *port*?
- Q9: Assume the presence of a *single unhatched* AGM egg mass on a container.
  What is the probability that this egg mass will be detected at a *transitional facility*?
- Q10: Assume a *single* WBB beetle is located somewhere within a tonne of wood packaging entering the country.
  What is the probability that this beetle will be detected at the port?
- Q11: Assume a *single* WBB beetle is located somewhere within a tonne of wood packaging that has entered the country.
  What is the probability that this beetle will be detected at a transitional facility?
- Q12: Assume a *single* WBB beetle is located somewhere within a tonne of wooden furniture entering the country.
  What is the probability that this beetle will be detected at the port?
- Q13: Assume a *single* WBB beetle is located somewhere within a tonne of wooden furniture that has entered the country.
  What is the probability that this beetle will be detected at the transitional facility?
- Q14: Assume a *single* WBB beetle is located somewhere on a tonne of wooden furniture that has entered the country.
  What is the probability that this beetle will be detected at a furniture shop or other retail outlet?
- Q15: Assume a *single* WBB beetle is located somewhere on a tonne of wooden furniture that has entered the country.
  What is the probability that this beetle will be detected at any time *after leaving* the furniture shop?
- Q16: How much time on average does a container spend at the port?
- Q17: How much time on average does a container spend at the transitional facility?
- Q18: How much time on average does a piece of furniture spend at a furniture shop?
- Q19: How much time on average does used machinery spend at a transitional facility?

Appendix B Elicitation Results

B.1 Average results for the in-person elicitation

Table 4: Average across all participants for each question in the in-person elicitation. Averages have been calculated for the High, Best (or Middle) and Low question responses for both rounds 1 and 2. The last column (% Diff) is the percentage difference of the Round 2 average best estimate to the original parameter value (major differences have been highlighted). ** Indicates parameters whose elicitation question differed somewhat from the model parameter (e.g. different units). A comparable value has been calculated for the original parameter in these cases.

							Round 1			Round 2		
Question	**Pest**	**Item**	**Pathway Point**	**Parameter**	**Original**	**High**	**Best**	**Low**	**High**	**Best**	**Low**	**% Diff**
Q1a (days)		Used vehicle	Port	Time at site	0.50	8.00	2.14	0.65	7.55	1.64	0.50	3.27
Q1b (days)		Used vehicle	Reg. site	Time at site	0.63	9.00	2.73	0.71	18.55	2.59	0.51	4.13
Q1c (days)		Used vehicle	Car Yard	Time at site	20.00	135.00	27.45	1.27	194.09	21.00	0.98	1.05
Q2 (%)	AGM	Sea vessel	Port	Balloon to shore	8.00	67.27	15.24	2.27	61.82	8.60	0.00	1.08
Q3 (gms)	Phy.	Sea container	Port	-	1.00	22.32	0.22	0.00	54.09	0.18	0.00	0.18
Q4a (%)	Phy.	Visitor	Airport	-		82.27	5.09	0.09	85.00	4.27	0.09	-
Q4b (%)	Phy.	Visitor	Nature attrac.	-		82.00	24.55	3.18	89.55	20.85	0.00	-
Q4c (%)	Phy.	Visitor		-		87.73	25.45	1.00	87.27	26.66	1.00	-
Q5 (%)	AGM	Used vehicle	-	Infest. Rate	**2.50	38.91	7.57	0.00	39.64	7.52	1.56	3.01
Q6a (%)	Phy.	Sea container	-	Infest. Rate	0.01	65.45	19.38	5.56	72.73	31.41	7.28	3140.91
Q6b (%)	Phy.	Visitor	-	Infest. Rate	0.09	56.36	28.27	10.64	75.45	40.91	17.36	454.55
Q6c (%)	Phy.	Resident	-	Infest. Rate	0.09	57.27	30.27	12.45	77.73	45.45	18.27	505.05
Q7a (%)	AGM	Used vehicle	Port	Detect. rate	85.00	89.45	58.36	27.00	95.73	74.09	39.55	0.87
Q7b (%)	AGM	Used vehicle	-	Treat. rate	97.00	99.99	98.45	80.00	99.99	98.45	88.09	1.01
Q7c (%)	AGM	-	-	Treat. rate (u)	0.00	67.73	27.55	6.45	61.82	23.82	6.45	-

B.2 Average results for the online elicitation

Table 5: Average across all participants for each question in the online elicitation. Averages have been calculated for the High, Best (or Middle) and Low question responses for both rounds 1 and 2. The last column (% Diff) is the percentage difference of the Round 2 average best estimate to the original parameter value (major differences have been highlighted). ** Indicates parameters whose elicitation question differed somewhat from the model parameter (e.g. different units). A comparable value has been calculated for the original parameter in these cases.

							Round 1			Round 2		
Question	**Pest**	**Item**	**Pathway Point**	**Parameter**	**Original**	**High**	**Best**	**Low**	**High**	**Best**	**Low**	**% Diff**
Q1 (%)	AGM	Sea containers	-	Infest. Rate	**0.48	32.51	4.57	0.05	21.33	3.70	0.03	7.71
Q2 (%)	WBB	Wood packaging	-	Infest. Rate	0.30	23.20	3.30	0.11	17.17	1.86	0.02	6.20
Q3 (%)	WBB	Furniture	-	Infest. Rate	**0.78	26.56	4.76	0.12	24.78	4.64	0.11	5.99
Q4 (gms)	Phy.	Visitors	Port	Pests/Infest.	1.00	220.00	5.17	0.70	57.78	2.55	0.19	2.55
Q5 (%)	AGM	Used vehicle	Reg. site	Detect. rate	85.00	74.43	42.89	15.00	77.67	40.56	10.56	0.48
Q6 (%)	AGM	Used vehicle	Car yard	Detect. rate	0.00	55.56	18.28	6.78	42.22	9.33	0.22	-
Q7 (%)	AGM	Used vehicle	End points	Detect. rate	0.00	37.78	10.17	5.67	21.67	2.39	0.11	-
Q8 (%)	AGM	Sea container	Port	Detect. rate	85.00	73.22	41.11	16.12	69.43	32.78	11.44	0.39
Q9 (%)	AGM	Sea container	Trans. facility	Detect. rate	85.00	71.67	33.67	10.57	51.33	16.34	1.12	0.19
Q10 (%)	WBB	Wood packaging	Port	Detect. rate	85.00	42.33	6.23	0.12	33.54	2.84	0.00	0.03
Q11 (%)	WBB	Wood packaging	Trans. facility	Detect. rate	0.00	41.10	5.56	0.12	20.22	2.46	0.00	-
Q12 (%)	WBB	Furniture	Port	Detect. rate	85.00	40.00	6.56	0.68	20.34	1.91	0.06	0.02
Q13 (%)	WBB	Furniture	Trans. facility	Detect. rate	0.00	45.10	12.34	1.13	18.57	3.00	0.01	-
Q14 (%)	WBB	Furniture	Furniture shop	Detect. rate	0.00	49.22	10.95	0.68	22.78	2.51	0.01	-
Q15 (%)	WBB	Furniture	End points	Detect. rate	0.10	49.99	10.44	0.01	39.22	5.80	0.02	58.00
Q16 (days)	-	Sea container	Port	Time at site	0.50	87.44	4.50	0.72	12.78	2.67	0.73	5.34
Q17 (days)	-	Sea container	Trans. facility	Time at site	0.63	64.11	4.50	0.76	18.67	3.22	0.65	5.14
Q18 (days)	-	Furniture	Furniture shop	Time at site	20.00	326.67	31.56	5.68	237.78	53.78	10.22	2.69
Q19 (days)	-	Used machinery	Trans. facility	Time at site	0.63	42.00	8.00	0.96	37.89	4.78	0.96	7.62

Appendix C Pest-item combinations

Table 6: Pest-item combinations.

Pest/Item	Used Vehicles	Sea Containers	Sea Vessels	Wood Packaging	Furniture	Returning Residents	Visitors	Used Machinery	Windborne	Live Plants
Generic	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AGM	Yes	Yes	Yes	No	No	No	No	No	No	No
WBB	No	No	No	Yes	Yes	No	No	No	No	No
Phytophthora	No	Yes	No	No	No	Yes	Yes	No	No	No
Fusarium	No	Yes	No	No	No	Yes	Yes	Yes	No	No
PSM	No	Yes	No	No	No	Yes	Yes	Yes	No	Yes