Value Added — Modelling the marginal return on investment within and across pathways

CEBRA project 21D

Final Project Report

Edith Arndt[1], Anca Hanea[1], Chris Baker[1,2,3], Thao P. Le[1,2,3], Andrew Robinson[1], Aaron Dodd[1], David Rolls[1], and John Baumgartner[1]

[1]The Centre of Excellence for Biosecurity Risk Analysis, The University of Melbourne
[2]School of Mathematics and Statistics, The University of Melbourne
[3]Melbourne Centre for Data Science, The University of Melbourne

March 26, 2024

cebra
Centre of Excellence for
Biosecurity Risk Analysis

Contents

List of Figures

List of Tables

1. Executive summary

Project 21D delivers a proof-of-concept resource allocation tool that the department can use to optimise investment into risk controls at the border. Underlying this tool are detailed simulation models of two biosecurity threat pathways (sea containers and containerised goods; cut flowers and foliage), based on conceptual diagrams produced through extensive consultation with departmental experts and analysis of actual import management and inspection data. The tool integrates the two import pathways to allow comparison of investment strategies between different pathways. Users can adjust import volumes and inspection sensitivities in order to assess the impact of changes – in terms of economic damage due to reduced asset yield – which is computed via integration of the pathway models with CEBRA's Value model.

Background

The objective of project 21D was to demonstrate a proof-of-concept for a whole of department approach to the setting of border risk control levels on the basis of risk-return. The project developed a tool (i.e., a basic user interface connected to a bespoke variant of the Value model) for determining the relative benefits of changed levels of investment (i.e., marginal benefits) in two, or more, competing interventions – on one or more pathways – so that the optimum level of investment in each intervention can be derived. Such a resource allocation tool has the potential to assist the department with determining the optimum level of investment in each intervention when resources are limited, or to achieve an appropriate level of protection against contaminant leakage and subsequent reduction in asset monetary values.

In the first phase of the project a user requirements analysis identified the functionality that end users would like to see implemented in the resource allocation tool. Subsequently, the scope of the project was limited to functionalities covered under user stories 1–3, which are plain language statements of the expected requirements:

- US01: As a Risk Owner, I need to understand the relative costs and benefits of different risk controls, so that I can determine the level and type of controls required to achieve the appropriate level of protection (ALOP).
- US02: As a Pathway Manager, I need to understand the relative costs and benefits of different risk controls, so that I can determine the level of investment in two competing controls.
- US03: As A Risk Owner or Pathway Manager, I need to understand how the design of a control influences overall risk, so that I can optimise its reduction (in context).

Concurrently with the user requirements analysis we developed a simple version of the resource allocation tool, the Minimal Reproducible Example, and tested it with end users. Then we conducted a targeted desktop review of how past research has dealt with resource allocation problems in biosecurity and related fields. Our review concluded that the use of Bayesian Networks (BNs) would be the best way to solve the decision problem, which was to optimise investment into risk controls that are dependent and aimed at protecting multiple assets that are affected by multiple threats. BNs can model complex probabilistic relationships.

To be able to develop a proof-of-concept resource allocation tool, we needed to select at least two distinct pathways and model them. The selected pathways were *sea containers and containerised goods* and *cut flowers and foliage*. Using the two pathway models we could split investment across the pathways and achieve optimal investment strategies with limited funding.

Methodology

CEBRA used the following overall approach to model the each of the two pathways:

1. *Create the conceptual diagram*: identify and visually present the process of importing sea containers/cut flowers into Australia, including the interactions of containers/cut flowers with a range of biosecurity risk controls and activities at the border. The diagram should include key determinants of decisions and requires incorporating a range of DAFF documents and correspondence with operational experts.

2. *Implement the simulation model*
    a) *Derive simulation model structure*: translate the conceptual diagram into a simulation model structure that consists of variables (nodes) and links connecting the nodes to demonstrate dependence and causality.
    b) *Parameterise simulation model*: Use a statistical and data-driven approach to estimate model parameters. The model should simulate containers/cut flowers moving through the biosecurity system and calculate the amount of undetected contamination (i.e., the output of the model is the leakage rate). Where possible, parameters in the simulation model should be informed by data received from the department. The simulation model includes options to adjust policy, allowing us to simulate the changes in undetected contamination under a range of candidate policies.

The sea container pathway

The first pathway we considered was the sea container and containerised goods pathway, which we refer to as the container pathway.

Through in-depth consultation with the department, we created a conceptual diagram showing the current controls in sea container risk management. The diagram lays out how containers are inspected and released given their attributes, such as whether they came from a port on the Country of Action List (CAL), whether that container was subject to the Sea Container Hygiene System (SCHS), whether the container contained goods, and whether it had a rural destination.

In order to fit the structural model and estimate current contamination approach rates, the department provided a large and complex dataset that included data from the S-Cargo and AIMS systems for sea containers and wharf-gate inspection data. AIMS data include cargo compliance verification (CCV) data, and data regarding the containerised goods.

In total, the provided dataset had 3,582,363 unique containers, of which 2,050,726 were deemed suitable. We matched these containers to the conceptual diagram of sea container movement to form the structural model, resulting in the `pathways` R package[1]. The package provides probability estimates of leakage of contamination and how that probability depends on the attributes (including CAL, SCHS, type of goods, and rural destination).

Note that CEBRA project 22B, *Sea Cargo Risk Management*, also looked at imported sea containers. To improve efficiency and streamline the work between projects 21D and 22B, both projects and both sets of teams worked together to produce the structural model of the sea container pathway.

The cut flowers pathway

The second pathway we considered was the cut flowers and foliage pathway, which we refer to as the cut flowers pathway. We consulted with the department and used a range of provided documents to create a conceptual diagram of the management decisions and actions on the cut flowers pathway, including documentation assessment and physical inspection. We translated the conceptual diagram into a simulation model structure, also included in the `pathways` package. The department provided a dataset with 101,022 unique entry-lines of cut flowers and foliage consignments, of which 100,756 lines were used. We fitted this data to the cut flowers structural model, and estimated contamination approach rates based on attributes (consignment type and country risk level).

Resource allocation tool

The proof-of-concept resource allocation tool is the main output delivered by this project. The innovative tool is a Shiny web app that allows end users to assess the relative benefits of different strategies of resource allocation into biosecurity risk controls. The app consists of a user interface with a number of tabs for specifying simulation scenarios. In the *Settings* tab users can define the spatial extent of the simulation, the duration, the number of iterations, the volume of containers or the number of lines arriving, and the global inspection effectiveness of each pathway. Inspection effectiveness can also be optionally defined at inspection method level. The *Analyses* tab has an option to compare investment scenarios within or across risk controls, involving one or both pathways and a custom selection of species. This tab also provides an interface to the original Value model comparison of biosecurity system *on* versus *off*, where users can select functional groups and simulate their arrival, spread and impact using a newly developed more time-efficient trajectory sampling approach. Tool outputs are in tabular and graphical format and can be exported for further use.

---

[1] `https://gitlab.unimelb.edu.au/cebra/cebra-r-packages/pathways`

Link to the CEBRA Value model

CEBRA previously developed a spatially explicit, bio-economic simulation model (the Value model) for estimating the net present value of Australia's biosecurity system (Dodd *et al.*, 2020). The resource allocation tool integrates the pathway models with the Value model to estimate the value gained when changing biosecurity policy. The tool calculates the amount of contamination passing through the border biosecurity system undetected, and how the amount of contamination depends on investment into risk controls. The amount of contamination that enters Australia is translated into establishment rates and then used as an input to CEBRA's Value model. The Value model estimates the damages of simulated pest and disease spread as the reduction in the monetary value of economic, environmental and social assets.

Recommendations

Based on our experience with the container and cut flowers data sets we offer the following recommendations to improve the quality of model outputs and to support the uptake of the resource allocation tool:

- **Improve data quality** by recording additional relevant information (e.g., proportion of stems inspected in a line or consignment of cut flowers — volume data is needed for this) and by curating database entries (e.g., introducing automatic data entry constraints so that missing data or incorrectly filed directions can be avoided; reviewing direction categories). Better data streams can make it easier to match up specific containers/goods with their attributes (source country, offshore/on shore treatment etc.), applied actions/directions (inspection, inspection type, treatment etc.), and outcomes (contamination found, release or export etc.).

- **Estimate inspection sensitivities**. To demonstrate the utility of the proof-of-concept resource allocation tool we assumed inspection sensitivities for the different types of border inspections undertaken on both pathways. However, when operationalising the project outputs, sensitivity parameters should be estimated through a formal expert elicitation process because inspection sensitivity directly influences the inferred approach rate and thus the leakage rate. Reducing uncertainty around inspection sensitivity parameters would improve the confidence in model outputs.

- **Hold workshop(s)** for use of the R code base. Successful integration of our research findings into departmental processes requires a solid understanding of the R packages developed in this project. CEBRA proposes to support the transition with tailored workshops for relevant departmental staff.

2. Introduction

Australia's biosecurity system faces increases in the volume, diversity, and complexity of introduction pathways for pests and diseases (Sardain *et al.*, 2019). This puts environmental, economic and social assets at risk of being damaged should an incursion or outbreak of a pest or disease occur. Government agencies have different options for responding to this pressure: to do nothing, to invest more resources into biosecurity activities, or to improve the way how resources are allocated. The first option is not feasible because regulators have a legislative requirement to address the threat of potential damages caused by exotic pests and diseases. The second option may be difficult to achieve because government agencies are confronted with ongoing resourcing challenges and there is indication that increasing investment is not effective in significantly reducing residual biosecurity risk (Craik *et al.*, 2017). Consequently, the preferred option for decision-makers would be to better allocate operational resources across the range of interventions (i.e., risk controls) they implement to mitigate biosecurity risk.

The objective of project 21D was to develop a proof-of-concept tool for allocating resources across the border biosecurity system. This final project report summarises the distinct elements of work that contributed to the development of the tool. Project 21D had many milestones, some of which were associated with a deliverable (Table 2.1). The first phase in this project identified user requirements (i.e., functionality) for the resource allocation tool through workshops with participants from the Department of Agriculture, Fisheries and Forestry (DAFF, the department, Chapter 3.1). In the next phase the project team undertook a review of existing resource allocation literature to identify the most suitable methodological approach for this project (Chapter 3.2). Concurrently, a simple version of the resource allocation tool, the Minimal Reproducible Example was developed as a R shiny app Chang *et al.* (2022) and tested with end users. To demonstrate proof-of-concept of a simple resource allocation tool, two pathways were selected to allow comparison of investment strategies between at least two pathways.

The first pathway selected in this project was the sea containers and containerised cargo pathway, from hereon referred to as the container pathway (Sections 4.1, 5.1, and Chapters 7, 8). The container pathway is of particular interest to the DAFF because of the high volume of sea containers imported by Australia annually. The container pathway incorporates many risk controls (e.g., external and internal inspection types, sampling regimes, and treatments) and provides sufficient data for quantitative description because the department routinely records the outcome of inspections and treatments. CEBRA project 22B, *Sea Cargo Risk Management*, also examined imported sea containers. Due to overlapping requirements, the structural model of the sea container pathway is shared between both projects.

The second pathway selected was the cut flowers and foliage pathway (Sections 4.2, 5.2, and Chapter 9), from hereon referred to as the cut flowers pathway. This pathway is of concern to the department because of its known risks (namely, frequent arthropod

Table 2.1.: Project 21D Milestones (MS)

MS	Description	Chapter(s)
1	Project start meeting	NA
2	User requirements analysis	3.1
3	Desktop review	3.2
4	Minimal reproducible example (Deliverable 1)	NA
5	Container pathway simulation model	4.1, 5.1, 6.1,7, 8
6	Validation of container pathway model (Deliverable 2)	NA
7	Cut flowers pathway simulation model	4.2, 5.2,6.2, 9
8	Validation of cut flowers pathway model (Deliverable 3)	NA
9	Improvements to the CEBRA value model (Deliverable 4)	Appendix H
10	Tool (Deliverable 5)	10
11	Draft report (Deliverable 6)	All

interceptions — the cut flowers pathway is among the top three known commodities for organism interceptions; IGB, 2019) and because it offers a sufficient amount of import data for this project. The cut flowers pathway was also a timely choice because the department commenced a three-part pest risk analysis for cut flower and foliage imports, starting in 2017, with the second report released in October 2022 and a third report expected some time in 2023.

Chapters 8 and 9 show the simulation results for the two selected pathways. They substantiate our modelling approach by demonstrating the utility of the simulation models for interrogating data and manipulating variables of interest.

To use the outputs of the pathway simulation models (namely, the leakage rates under different risk control strategies), the code base of CEBRA's Value Model for calculating expected damages was updated (Appendix H). In its final phase, the project developed a basic user interface (the tool) that allows end users to compare the relative costs and benefits of different levels of intervention within and across the two pathways (Chapter 10).

3. Methodology

3.1. User requirements analysis

The first phase of this project focused on establishing a baseline specification of expectations. To this effect, CEBRA conducted a series of two-hour workshops with participants from different divisions within the department. The purpose of the workshops was to identify "User stories", which are articulations (plain language statements) of what type of functionality the project's resource allocation tool would need to have to deliver value to departmental end users.

The workshops identified five user stories. User stories 1-3 were deemed in scope, while user story 4 was out of scope and user story 5 considered a potential extension.

Table 3.1.: User Stories

ID	User story
US01	As a Risk Owner, I need to understand the relative costs and benefits of different risk controls, so that I can determine the level and type of controls required to achieve ALOP
US02	As a Pathway Manager, I need to understand the relative costs and benefits of different risk controls, so that I can determine the level of investment in two competing controls
US03	As A Risk Owner or Pathway Manager, I need to understand how the design of a control influences overall risk, so that I can optimise its reduction (in context)
US04	As a Departmental Executive, I need to understand the relative costs and benefits of different risk controls, so that I can implement controls that minimize commercial impacts
US05	As a Risk Owner, I need to understand the relative benefits and costs of different risk controls, so that I can determine the appropriateness of controls when conducting assurance

In other words, user story 1 investigates how to minimise cost when investing in different risk controls to achieve ALOP (the appropriate level of protection). User story 2 looks at optimising investment across two risk controls and user story 3 at maximising the value of investment across multiple pathways / controls.

Details on the user requirements workshops and the full list of 40 functional requirements (i.e., things the tool must be able to do) and 7 non-functional requirements (i.e., how the tool needs to operate) are provided in an interim project report (Dodd *et al.*, 2021). The user requirements specifications guided the development of the simple resource allocation tool described in Chapter 10. However, at the time of project delivery, complications associated with defining the costs corresponding to varying levels of controls prevented these costs from being incorporated into the tool. For example, it was unclear how cost scales with inspection effort, the latter being the example control class provided in the proof-of-concept tool. Rather than implement dummy costs, which were deemed misleading, a decision was made to omit this component, leaving cost-benefit comparisons as a post-processing exercise or avenue for future tool improvement.

3.2. Methods search (targeted literature review)

Milestone 3 of project 21D delivered a desktop review of how the common resource allocation problem of investing finite resources across multiple risk controls and targeting multiple threats has been solved in different fields: biosecurity, invasive species management, and the operations research literature in the context of bushfire risk management (Dodd *et al.*, 2022). The review report defined the generic decision problem as a (multiple-choice, multiple-objective) quadratic knapsack problem. The optimal solution for knapsack problems is often derived by using integer programming or piecewise linear programming. However, the report concluded that using dynamic programming techniques to solve an optimisation problem with many dependent risk controls and multiple threats that affect multiple assets would not be possible computationally. Objective functions are often built on simplifying assumptions and can be solved through brute force (i.e., calculating the benefits of every single combination directly, and selecting the one with the highest value). However, the time required to calculate the optimal strategy makes this option unfeasible when multiple levels of investment are simultaneously considered for many competing controls (e.g., 27×10^{24} years when considering 5 levels of investment for each of 60 controls, assuming the benefit of each strategy can be computed in 1 nanosecond). To avoid oversimplifying the decision problem to a point where brute force computation is possible, but instead stay true to the complexity of the risk control network at the border, Dodd *et al.* (2022) proposed to use Bayesian Networks (BNs) to model the container and cut flowers pathways.

BNs allow risk analysts to model complex probabilistic relationships (e.g., dependencies, Pearl, 1988) and have increasingly been used in import risk analysis (e.g., Holt *et al.*, 2018; Jamieson *et al.*, 2021). With BNs capturing the complexity of biosecurity interactions, the optimisation process can then be delegated to users who choose the set of alternative investment options to compare.

3.3. Conceptual diagram

The first step for milestones 5 and 7 in this project was to develop conceptual diagrams of the container pathway and the cut flowers pathway, as a basis for building a model structure (Figure 3.1). The department currently has no overarching conceptual diagram for the container pathway that shows the combination of risk controls implemented at the border. The department's diagram of the cut flowers pathway is high-level. Thus, CEBRA developed a conceptual diagram for both pathways based on written material and meetings with experts. The reading material provided by DAFF included work instructions, policy documents, and e-mail content (see Appendix A).

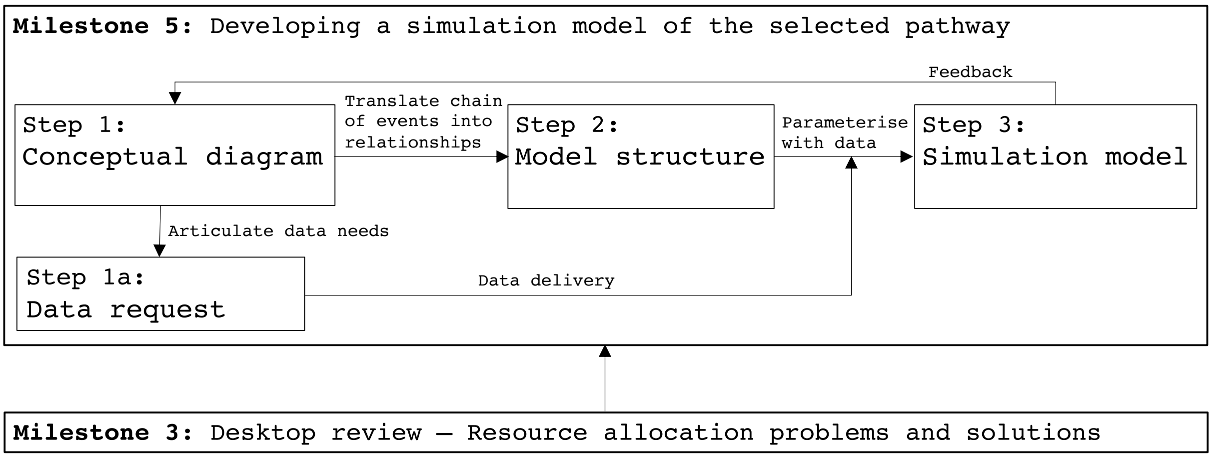

Figure 3.1.: Steps involved in developing a simulation model of the selected pathway.

Based on the material provided and meetings with staff from pathway policy (cargo and conveyances), operational areas (docks and the airport), and cut flowers pathway managers, CEBRA developed first drafts of the conceptual diagrams. Subsequently, questions for follow-up meetings were collated to refine the interaction of containers or goods with risk controls and to confirm that the models reflect the department's understanding of the processes. CEBRA further refined the models based on feedback from the analysis of the container and cut flowers data in step 3, and more consultation with DAFF experts (Table B.1).

In agreement with the department, risk controls related to Brown Marmorated Stink Bug and Khapra beetle were not included in the final conceptual diagram of the container pathway. The main reason for the exclusion was that these risk controls are mainly commodity rather than container focused.

After experts had validated the main elements of the conceptual diagrams, CEBRA submitted a data request to the department (Step 1a). In the second step, the chain of events for both pathways was translated into a model structure that captured all relationships and dependencies among variables (Figure 3.1). In the third step, the models were parameterised with data provided by the department to obtain a simulation model of the pathways.

3.4. Model structure

The conceptual diagram constitutes the foundation for developing a simulation model underpinned by a probabilistic framework. In such models we represent probabilistic relationships between uncertain variables, first qualitatively (corresponding to the model selection step) and then quantitatively (corresponding to the parametrisation step). We used the graphical advantages of Bayesian networks (BNs) (e.g., Korb & Nicholson, 2010) for the qualitative, model selection and structuring step. BNs provide a transparent and efficient way to model complex relationships, while allowing integration of data from different sources and with varying degrees of uncertainty (e.g., Pearl, 1988).

The qualitative part of a BN (the BN's structure) consists of a set of variables, represented as nodes of a directed acyclic graph, and a set of arcs connecting these nodes. The arcs represent dependence, causality, and interactions between the variables. These relationships can be probabilistic (e.g., if X happens there's a chance Y will happen as a consequence) or deterministic (i.e., functional, logical).

Figure 3.2 shows an example of a BN model drawn from the Integrated Biosecurity Risk Analysis Model (IBRAM) (Jamieson *et al.*, 2021) developed for New Zealand's Ministry for Primary Industries (MPI) for a context similar to ours. The same example is discussed in Dodd *et al.* (2022) and we repeat part of that discussion for illustrative purposes.

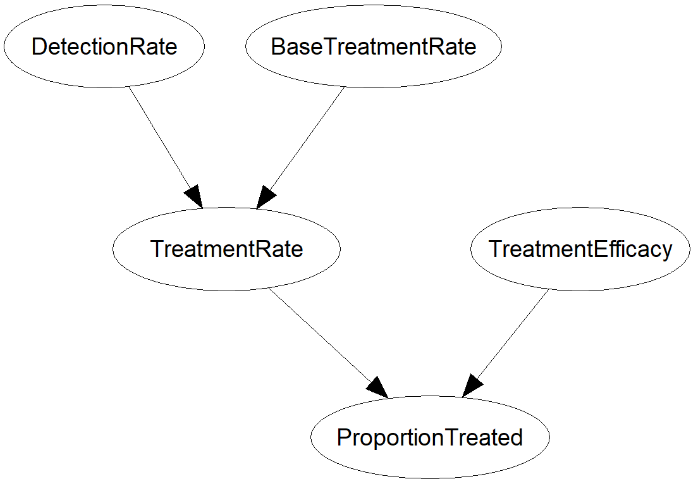

Figure 3.2.: A Bayesian Network (BN) example from Jamieson *et al.* (2021)

The BN in Figure 3.2 is used in Jamieson *et al.* (2021) to model the effectiveness of pre-border management interventions. In general, in BNs:

- Variables (e.g., the probability that biosecurity risk material is detected in a consignment) are shown as nodes (e.g., *Detection Rate*).
- Links, also called arcs, indicate relationships between nodes.
- Arcs point from a "parent" to a "child" node.

The *Treatment Rate* node from Figure 3.2 is the child node of the parents: *Detection Rate* and *Base Treatment Rate*. Parent nodes are parametrised as distributions. In the example above, *Detection Rate*, *Base Rate Treatment* and *Treatment Efficacy* each require a probability distribution (e.g., the parametrisation of the model step needs the probability of detection, the probability of treatment and the probability that treatment is effective). For their children, such probability distributions need to be specified conditional on the values of their parents. For example, the probability distribution of the proportion of treated items (*Proportion Treated*) has to be specified (and evaluated) conditional on two factors (corresponding to its two parents): the probability that an item will be treated AND the efficacy of that treatment. The nodes and arcs alone are called the qualitative part or the structure of a BN. The marginal and conditional probabilistic distributions are called the quantitative part of a BN.

For our two selected pathways, risk controls and properties of containers/cut flowers became nodes of the BN. The combination of e.g., containers' properties leading to a certain type of risk control was represented by pointing arrows from the former to the latter. The probabilistic nature of these influences will be captured by conditional probabilities of certain risk controls to be engaged, given the containers' properties. These controls influence the leakage rate which is the end point of the BN (not depicted in the conceptual diagrams). Variables that do not mirror nodes of the conceptual diagram were added to the BN structure, since the leakage rate is influenced by external factors (not representing the process captured by the conceptual diagrams), i.e., the arrival rate, the detection rate, etc.

Several iterations of the BN structure were necessary, reflecting the iterations and refinements of the conceptual diagram.

The structure of the model guides the data requirements for parametrisation. Available data rarely has the size format needed for reliable parametrisation. We discuss data processing in the next section.

3.5. Data cleaning and processing

The department provided two sets of data: one for the container pathway, and one for the cut flowers pathway. Each posed multiple challenges to overcome in order to extract the information needed to parameterise the models.

The aim of the cleaning and processing is to organise all the information about each container or cut flower line/consignment into a single location. Each container or cut flowers line/consignment should have information about their attributes (e.g. country source) and any decisions applied and their results (such as whether they were inspected, and any inspection result).

The data cleaning and processing was nontrivial.

The data for the containers pathway came in multiple parts, requiring us to combine them together. There was ambiguity between combining data sets, requiring us to make assumptions. Data cleaning included standardising variable names, country names, removing duplicate rows, and mapping port codes to countries. Entries with too much missing data were removed. We also applied 'rules' to sort out pieces of information (such as sorting the different contaminants and hazards listed into biosecurity-risk-material-type).

As the cut flowers data was relatively simpler, we made assumptions about the majority of missing data and management actions, rather than removing incomplete entries (such as for cut flower documentation assessment results when no documentation information is available in the data). The cut flowers data also contained ambiguous results and incorrectly filed data.

Once all the relevant information is gathered, cleaned, and organised, we can easily make summaries and extract information such as the proportion of containers that had a rural destination and had detected contamination. This allows us to parameterise the model.

3.6. Parametrisation of the model

The simulation model is the parametrised (quantitative) BN. The parametrisation of the BN structure involves calculating conditional distributions of child nodes (i.e., nodes pointed at by arrows) given their parents (nodes where the arrow originated from) in the graph. When relationships between parents and children are not probabilistic, they are driven by mathematical formulae where the child is a function of the parents. Uncertainty propagates through these relationships. The required conditional distributions were estimated as relative frequencies from the data provided by the department.

For example, the outcome of an external container inspection is a function of the true external contamination rate (i.e., the external approach rate) and the sensitivity of the inspection method. From the data, we can directly calculate the number of containers that had an inspection with no discovered contamination, and the number of containers with contamination. By estimating the inspection sensitivity (through departmental consultation), the BN will subsequently estimate the 'true' external contamination rate. The true contamination is a function of the observed contamination rate and the detection capability (i.e., the sensitivity) of the inspection method.

Chapters 8 and 9 demonstrate the utility of the simulation models for manipulating variables to test the effects on the leakage rate, thereby substantiating our modelling approach.

Finally, model reparametrisation is advisable if container attributes or inspection activities change in ways that modify these frequencies or indeed the structure of the network. This includes cases where the relationship between top level node states and contamination rates change. For example, adding a set of countries to the CAL (and hence removing from the non-CAL cohort), may impact the contamination rates of one or both groups, requiring reparametrisation of the model.

3.7. Uncertainty, counterfactuals, sensitivity analysis

The output of the simulation models is the leakage rate, which is determined by the top level variables of the pathway. Conditioning the simulation models on the different possible states of the variables and then running the models leads to different leakage results. To capture the *uncertainty* around model outputs simulations were run 1,000 times to obtain credible intervals around the point estimate. The uncertainty is well approximated by the sampling distribution of the probability parameter of binomial distributions centered on the point estimates of leakage rates.

Counterfactual analysis consists of investigating "what if" scenarios or different policy changes. Because these changes act on an "uncertain state" of the world we are modelling with the proposed simulation model (driven by the set of interconnected nodes in the BN) the effects/consequences of these changes on the values or distributions of certain model variables. For example: adding to, or removing countries from the CAL (is a decision which) changes the distribution of the node CAL (which represents one of the top-level input variables in the container pathway). In this project we assessed a number of counterfactual scenarios for each pathway. These counterfactual scenarios may not fully capture the changes to underlying risks because the simulation models do not consider behaviour changes in response to policy changes. Behavioural aspects are difficult to model due to 'information asymmetry' where the regulator is not privy to internal business decisions made by non-government pathway participants.

We view *sensitivity analysis* as a systematic investigation of how the impact of uncertainties of one or more model parameters can change the uncertainties of the output variables (i.e., we exclude sensitivity to the choice of model, variables and other modelling assumptions). A global sensitivity analysis examines sensitivity to each parameter and every possible combination of changes of parameters. This is almost never performed in practice. In turn a one-parameter-at-a-time (or local) sensitivity is much more common. Many of the situations investigated through the counterfactual analysis can be considered part of the sensitivity analysis. In other words, the decision to e.g., remove a country from the CAL changes the distribution (i.e., the parameters of this distribution) of the CAL variable (the proportion of countries on the CAL, therefore the number of containers arriving from CAL locations), hence such a decision is both part of the counterfactual and a sensitivity analysis.

In most cases, simulating a scenario whereby one or more biosecurity controls are modified will require altering the conditional probabilities of affected nodes and reconstructing the network. The `pathways` package includes functions that facilitate changing inspection effectiveness and/or treatment effectiveness when constructing a network, and assessing sensitivity of leakage to variation in either or both of these. Other controls that may be of interest are less easily adjusted at present; however, manual manipulation of CPTs is possible, e.g. via the `replaceCPT` function provided by the `gRain` R package, and wrappers could be added to the `pathways` package.

3.8. Proof-of-concept resource allocation tool

The simple resource allocation model was developed as an interactive Shiny (Chang *et al.*, 2022) app in R (version 4.2.1, R Core Team, 2022) and integrates the two pathways selected in this project with CEBRA's bio-economic Value model (Dodd *et al.*, 2020). It addresses the requirements outlined in user stories 1-3 (Table 3.1) by integrating the pathway models with CEBRA's spatiotemporal model of biosecurity arrival, spread, and impact. The tool is described in detail in Chapter 10.

4. Conceptual diagrams

The purpose of a conceptual diagram is to capture and visualise the interactions of imported containers or the cut flowers commodity with the border risk controls and processes currently implemented by the department. They are the basis for developing the model structures.

In this chapter, we describe the conceptual diagrams for the container (Section 4.1) and cut flowers (Section 4.2) pathways.

4.1. The container pathway

The scope of the conceptual diagram[1] for the container pathway is limited to the *Country of Action List (CAL) high-risk pathway* and the *low-risk pathway* of containers (Figure 4.1). The *no-risk pathway* will be captured separately by analysing the Cargo Compliance Verification (CCV) data, and is not incorporated into the conceptual diagram. Goods, the contents of containers, and transshipping containers are generally excluded, because the focus is on arriving containers. However, the diagram does include goods-related interventions such as full-unpack and rural tailgate. On rare occasions, these inspections can include biosecurity officers inspecting the inside of containers while also inspecting specific goods. In this sense, internal contamination includes both goods and internal container contamination.

The CAL pathway describes the regulatory process for containers and break bulk cargo originating or arriving from countries and ports that have been identified as having high historical levels of contamination with giant African snails, other pests, or soil (IGB, 2018). Influencing the container contamination rate on this pathway is the Sea Container Hygiene System (SCHS), an offshore quality assurance system that consists of approved facilities located at ports in some CAL countries. Shipping containers passing through these facilities are cleaned and treated with insecticide (internally and externally), and stored in designated areas at the port of loading (IGB, 2018). The onshore CAL biosecurity measures aim to manage external contamination risk by prescribing external inspections (six-sided) of full containers, and external and internal inspections of empty containers on the CAL high-risk pathway. The destination of containers and the type of goods determine the implementation of other types of inspection on the CAL pathway, namely external inspection, rural tailgate and full-unpack at an approved arrangement (AA).

---

[1] The first iteration of the conceptual diagram was developed by Aaron Dodd.

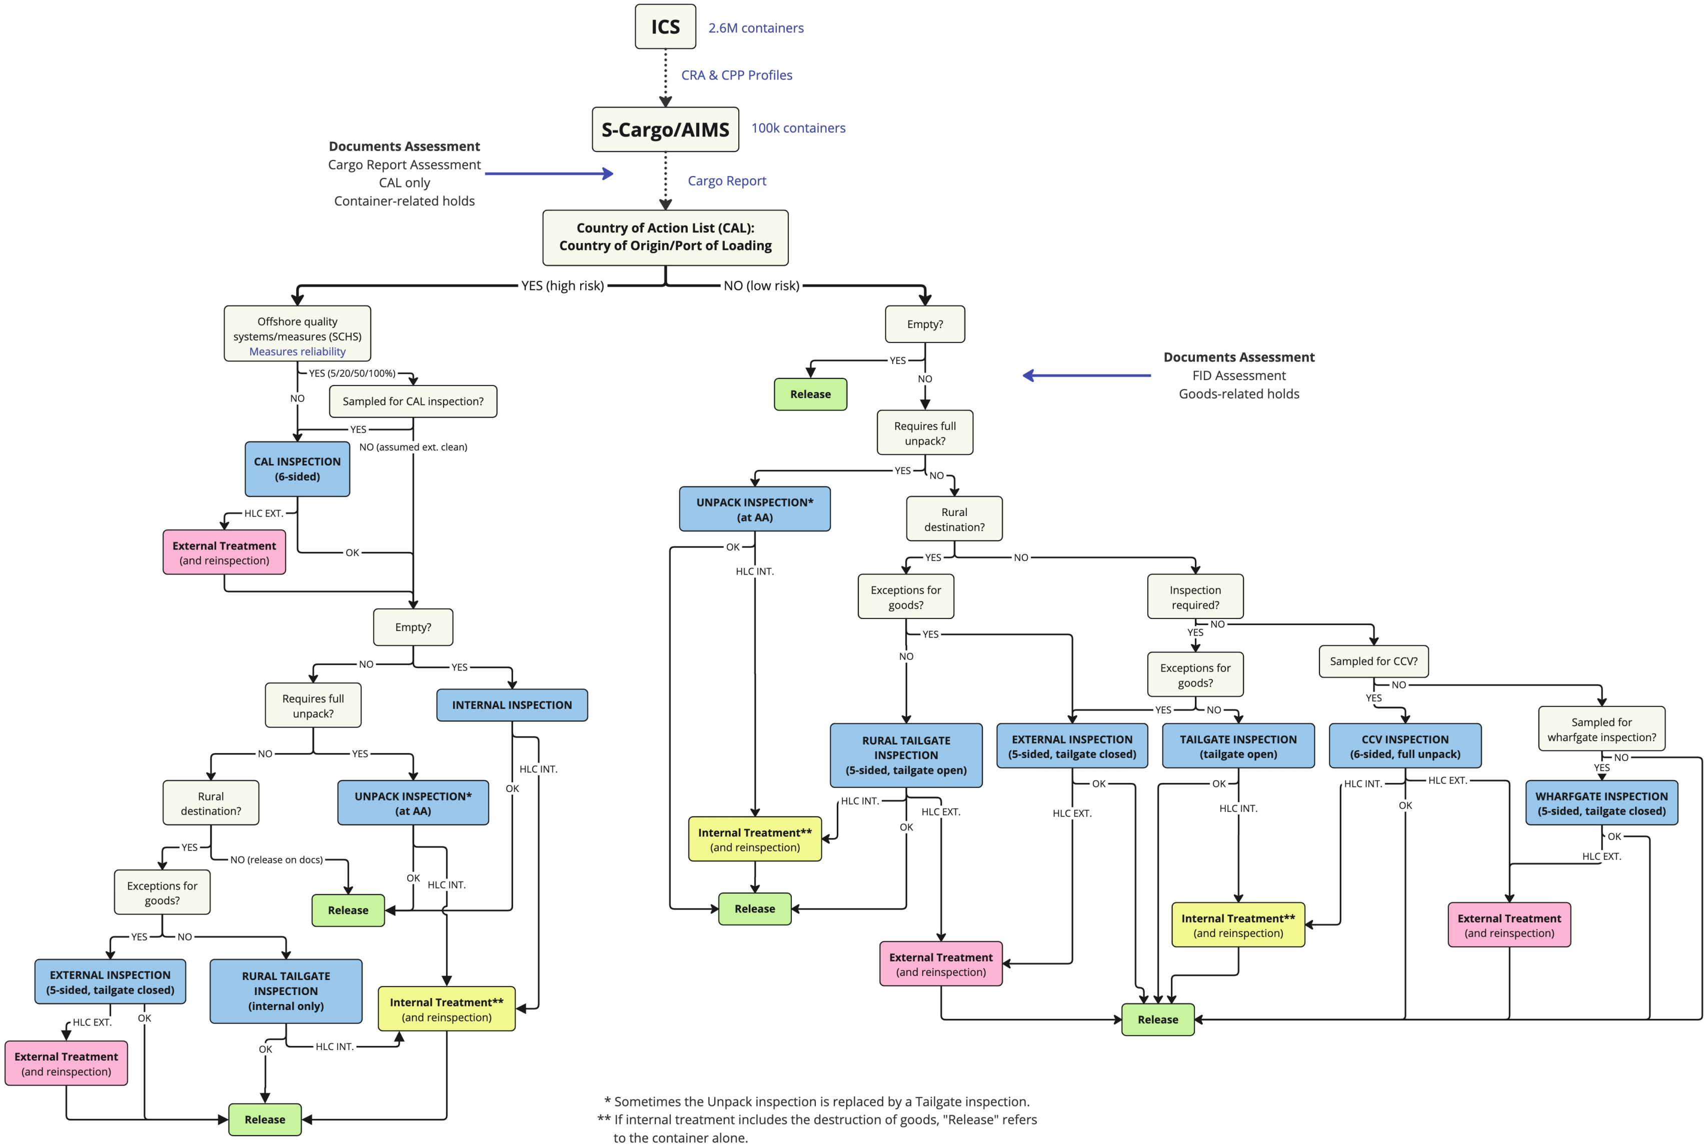

* Sometimes the Unpack inspection is replaced by a Tailgate inspection.
** If internal treatment includes the destruction of goods, "Release" refers to the container alone.

Figure 4.1.: Conceptual diagram of the sea container management pathway

The low-risk pathway refers to containers that are not from a CAL country or port. These containers can have a rural destination which triggers different inspections based on the type of commodity. The low-risk pathway also includes containers with a metro destination that are typically released on documents due to their very low risk status. However, metro-bound containers can be randomly selected for CCV or wharfgate inspections. CCV inspections are not a risk control but act as a verification tool for the department to confirm that containers not directed to the department's Agriculture Import Management System (AIMS) pose a very low risk to biosecurity. Wharfgate inspections are done rarely and, like CCV inspections, provide an estimate of leakage on the low- and no-risk container pathway.

The conceptual diagram does not always represent a true sequence of events at the border, because containers may be re-directed to other interventions for a particular reason. Examples of these edge cases are containers that initially were directed for CCV inspection may in fact be released from biosecurity control without inspection. Or, sometimes a direction to unpack cargo is replaced by a tailgate inspection.

The conceptual diagram includes four key questions at the top. The answers to these questions determine which of the border risk controls most of the containers will be subjected to. The questions are:

1. Does the container come from a CAL country of origin or port of loading?
2. Was the container treated offshore under the SCHS?
3. Does the container have a rural destination?
4. Is the container empty[2]?

Clean, full containers on the high-risk pathway that do not require a full unpack and have a metro destination are released from biosecurity control. However, empty high-risk containers which are inspected and are free from external contamination must be inspected internally and are only released when free of contamination. On the low-risk pathway, empty containers and full (i.e., non-empty) containers that do not require an inspection (and are not selected for CCV/wharfgate[3] leakage surveys) are released from biosecurity control. All other scenarios depicted involve a regulatory intervention prior to release of containers.

The biosecurity risk controls in the conceptual diagram include inspections and treatments, which can be applied externally and/or internally[4]. Based on the country or port of origin, the destination, and presence and type of goods, we distinguish between eight different types of inspections (Table 4.1) that are represented in the conceptual diagram. For example, CAL containers that have not been managed through the offshore SCHS require a *six-sided* external inspection, whereas full non-CAL containers with a rural destination undergo a *Rural Tailgate* (i.e., 5-sided inspection + tailgate

---

[2] A non-empty container covers FCL (Full Container Load), FCX (Full Container with Multiple House Bills of Lading) and LCL (Less than Container Load). Some of the goods from a non-empty container, (e.g., diplomatic goods, dangerous goods) may be classified as exceptions and trigger different interventions due to this status.

[3] Wharfgate inspections can also be undertaken on containers that have a rural destination, but the proportion of containers in this cohort is small (0.17% across wharfgate-inspection-eligible containers).

[4] Internal inspections represent only a small percentage of the total number of inspected containers. A proportion of those internal inspections are conducted by industry and inspection results do not get reported to DAFF.

Table 4.1.: Types of inspections on the container pathway.

Inspection type	Description	Outcomes
Six-sided	6-sided external inspection	External treatment and re-inspection, or full unpack
Rural Tailgate	5-sided external inspection with tailgate open	External and/or internal treatment and re-inspection, or release
External	5-sided external inspection with tailgate closed, for goods with exception	External treatment and re-inspection, or release
Tailgate	Internal inspection with tailgate open	Internal treatment and re-inspection, or release
Wharfgate	3-4 sided external inspection (performed on truck at exit gate) with tailgate closed, randomly selected	External treatment and re-inspection or release
CCV	4-sided external inspection with container on the ground and full unpack of goods, randomly selected	External and/or internal treatment and re-inspection, or release
Internal	Internal inspection of empty containers	Internal treatment and re-inspection, or release
Un-pack	Full unpack of goods at an AA	Internal treatment and re-inspection, or release

open), unless directed for a full unpack[5], or when contents preclude internal inspection (in which case the tailgate remains closed).

Contaminated containers or goods trigger a treatment. Contamination types include snails, ants, insects, seeds, soil, and other plant or animal material. If low-level contamination is found at an inspection, it is treated on the spot (less than 5 minutes cleaning). High-level contamination requires containers and/or goods to be treated and then re-inspected. In this project, we only differentiate between internal and external treatments (i.e., not between types of treatments) and assume that the effectiveness of the treatment is 100%. When treated containers and/or goods are compliant at re-inspection, containers are released from biosecurity control. If not compliant, we assume they will get treated again until compliant at re-inspection, but we do not capture these other iterations in the conceptual diagram. *Release from biosecurity control* is the endpoint in the diagram.

As mentioned in Section 3.3, the conceptual diagram does not include BMSB and Khapra risk controls.

---

[5] When containers require a full unpack at an AA it does not mean that the container itself will be inspected. Only a small proportion of containers require full unpack and inspection.

4.2. The cut flowers pathway

The cut flowers and foliage conceptual diagram is spread across Figures 4.2 and 4.3. It is constructed primarily from the cut flowers work instructions by taking each instruction and group of instructions and adding it to the conceptual diagram.

These are the main stages we distilled:

1. Pre-export actions. The majority of cut flowers and foliage consignments require some kind of pre-export process, such as treatment (fumigation) or NPPO-approved systems approach.
2. Pre-inspection and documentation checks. This involves checking that all relevant documentation (e.g. phytosanitary certificates) are in place, and that the goods are as described on the documentation.
3. Inspection.
    a) Sampling (determining how many samples to inspect). At most 600 stems are inspected per consignment.
    b) The inspection procedure. The outer packing is inspected first, followed by the actual goods inspection.
    c) Inspection outcomes and follow-ups, such as further identification of unknown seeds or treatment in the case of discovered live snails, slugs etc.
4. Final actions, be it release, re-export or destroy. If the goods are not released, then the importer can choose whether to export or destroy.

In general, almost all incoming consignments require some kind of offshore treatment and certification. All consignments should go through document checks, and almost all consignments should have some kind of inspection (with a maximum of 600 stems inspected in a consignment, based on ISPM 31 International Standards for Phytosanitary Measures (2008)). Not all contaminated lines trigger a treatment — goods can be released if for example there are only 1 or 2 contaminant seeds, or they can be directly exported or destroyed. If treatment does occur, then it is typically fumigation with methyl bromide or reconditioning.

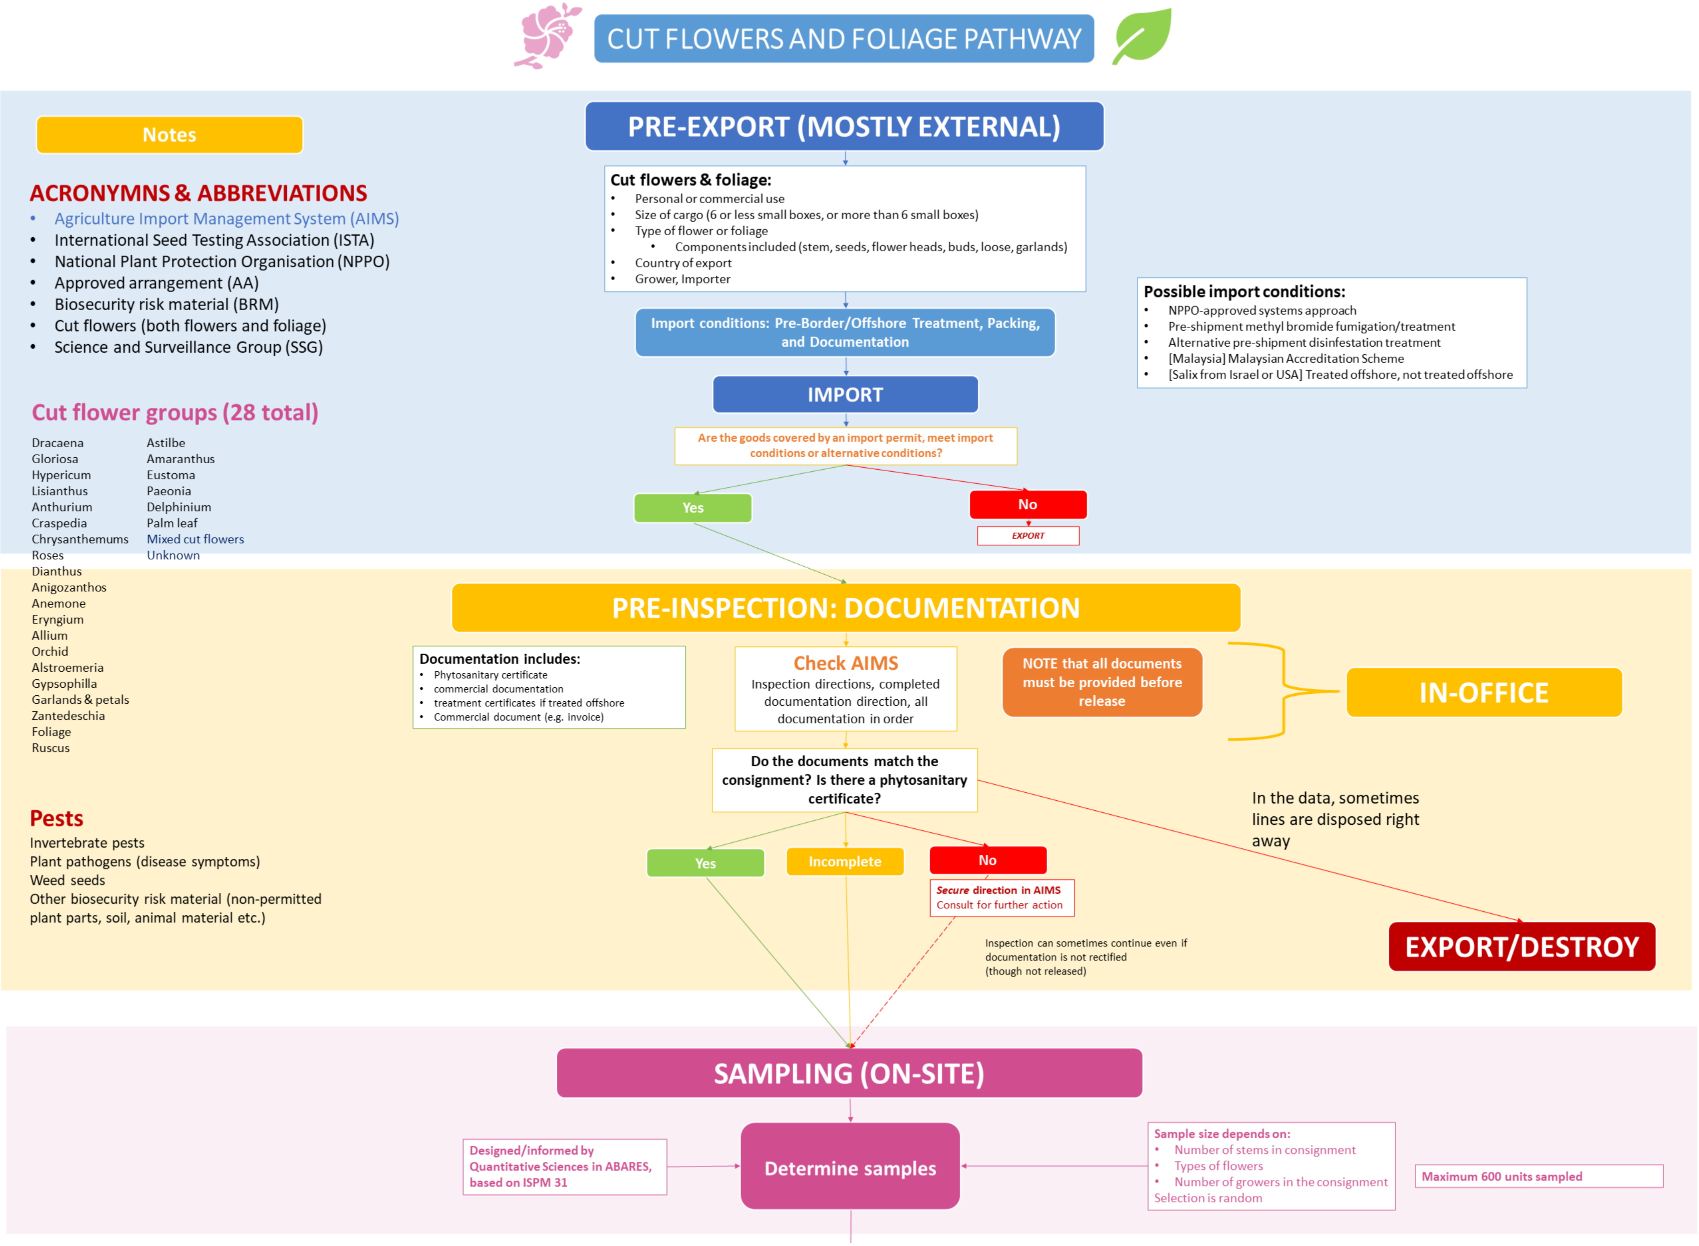

CUT FLOWERS AND FOLIAGE PATHWAY

Notes

ACRONYMNS & ABBREVIATIONS

- Agriculture Import Management System (AIMS)
- International Seed Testing Association (ISTA)
- National Plant Protection Organisation (NPPO)
- Approved arrangement (AA)
- Biosecurity risk material (BRM)
- Cut flowers (both flowers and foliage)
- Science and Surveillance Group (SSG)

Cut flower groups (28 total)

Dracaena
Gloriosa
Hypericum
Lisianthus
Anthurium
Craspedia
Chrysanthemums
Roses
Dianthus
Anigozanthos
Anemone
Eryngium
Allium
Orchid
Alstroemeria
Gypsophilla
Garlands & petals
Zantedeschia
Foliage
Ruscus
Astilbe
Amaranthus
Eustoma
Paeonia
Delphinium
Palm leaf
Mixed cut flowers
Unknown

Pests

Invertebrate pests
Plant pathogens (disease symptoms)
Weed seeds
Other biosecurity risk material (non-permitted plant parts, soil, animal material etc.)

PRE-EXPORT (MOSTLY EXTERNAL)

Cut flowers & foliage:
- Personal or commercial use
- Size of cargo (6 or less small boxes, or more than 6 small boxes)
- Type of flower or foliage
  - Components included (stem, seeds, flower heads, buds, loose, garlands)
- Country of export
- Grower, Importer

Possible import conditions:
- NPPO-approved systems approach
- Pre-shipment methyl bromide fumigation/treatment
- Alternative pre-shipment disinfestation treatment
- [Malaysia] Malaysian Accreditation Scheme
- [Salix from Israel or USA] Treated offshore, not treated offshore

Import conditions: Pre-Border/Offshore Treatment, Packing, and Documentation

IMPORT

Are the goods covered by an import permit, meet import conditions or alternative conditions?

Yes

No

EXPORT

PRE-INSPECTION: DOCUMENTATION

Documentation includes:
- Phytosanitary certificate
- commercial documentation
- treatment certificates if treated offshore
- Commercial document (e.g. invoice)

Check AIMS

Inspection directions, completed documentation direction, all documentation in order

NOTE that all documents must be provided before release

IN-OFFICE

Do the documents match the consignment? Is there a phytosanitary certificate?

In the data, sometimes lines are disposed right away

Yes

Incomplete

No

***Secure* direction in AIMS**
Consult for further action

Inspection can sometimes continue even if documentation is not rectified
(though not released)

EXPORT/DESTROY

SAMPLING (ON-SITE)

Designed/informed by Quantitative Sciences in ABARES, based on ISPM 31

Determine samples

Sample size depends on:
- Number of stems in consignment
- Types of flowers
- Number of growers in the consignment

Selection is random

Maximum 600 units sampled

Figure 4.2.: Conceptual diagram of the cut flowers and foliage management pathway, part 1 (start)

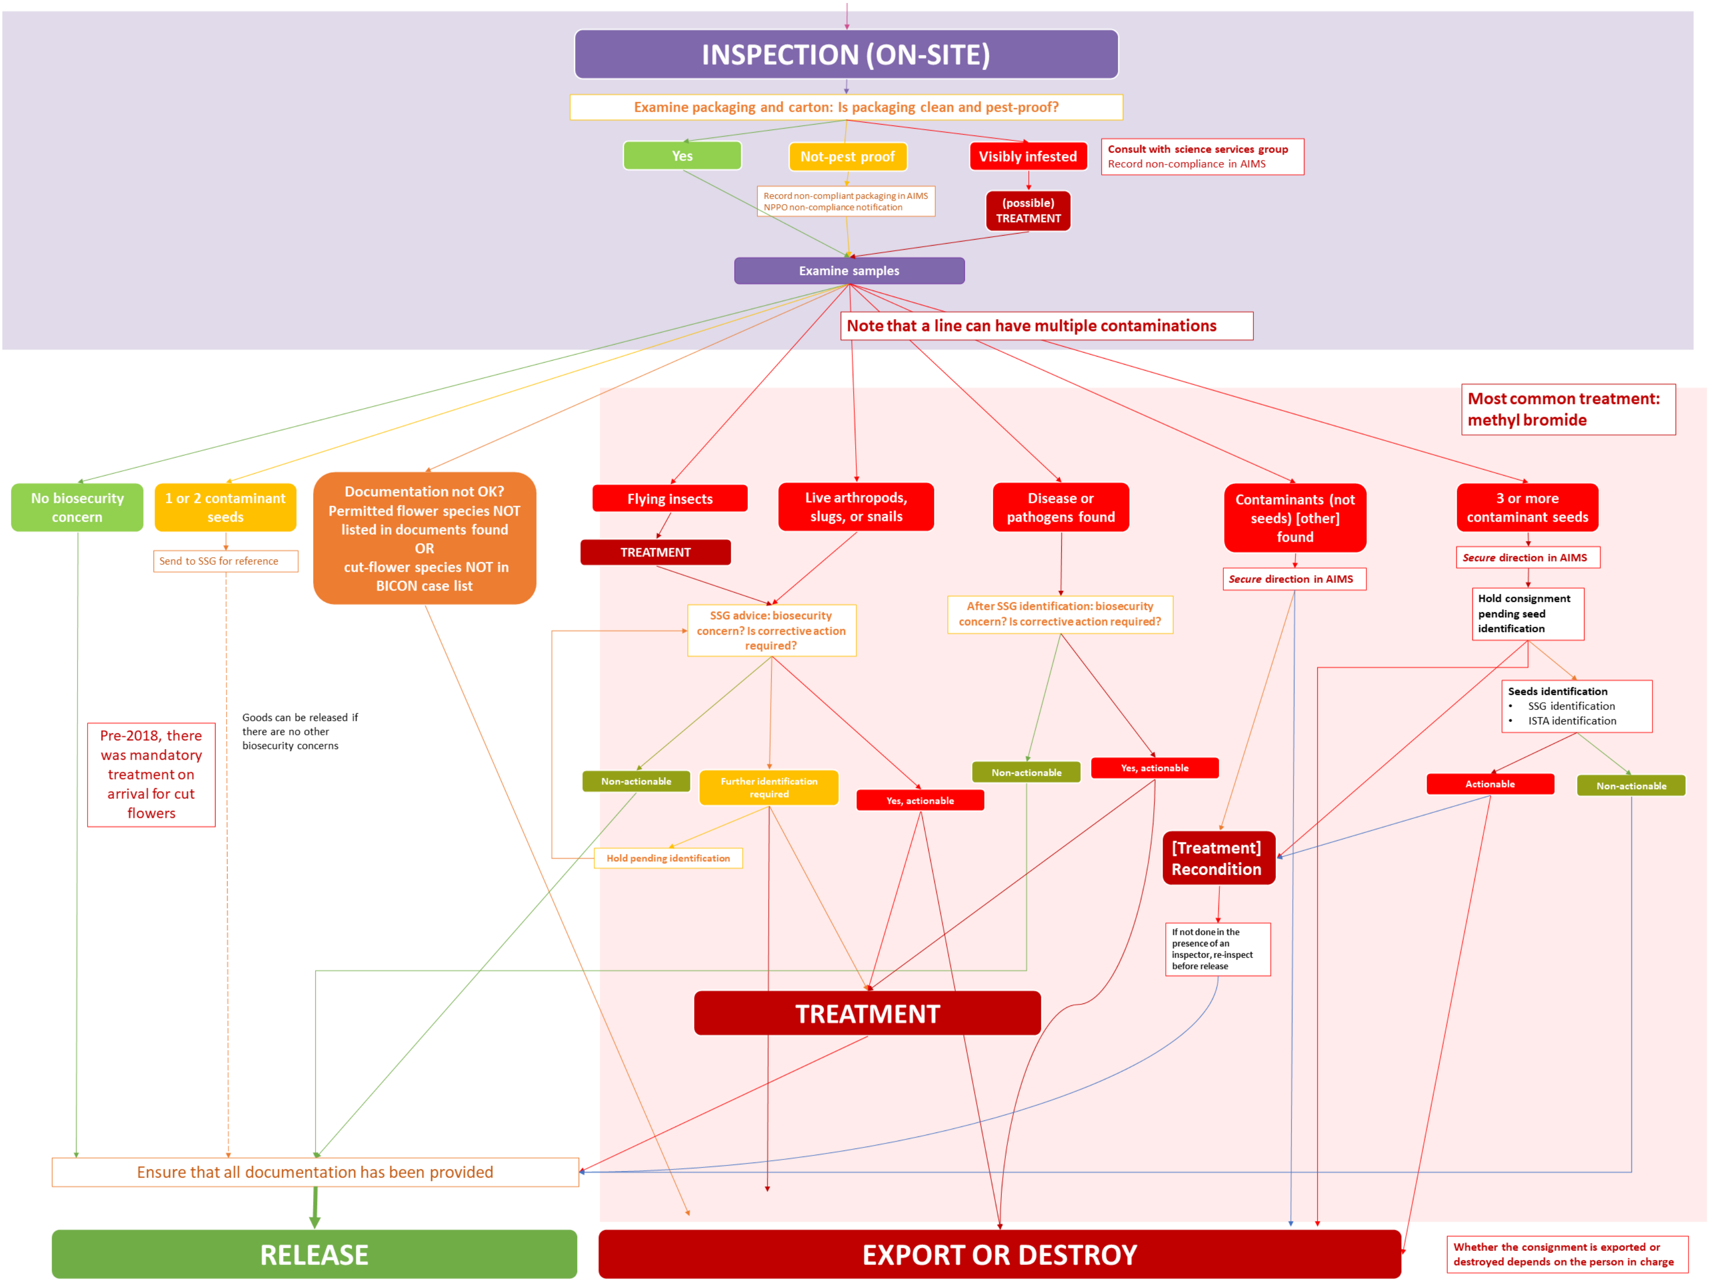

Figure 4.3.: Conceptual diagram of the cut flowers and foliage management pathway, part 2 (end)

5. Model structures

In this chapter we present and explain the model structures for the container (Section 5.1) and cut flowers (Section 5.2) pathways.

When translating the conceptual diagram into a BN, some of the nodes were better represented by constants (rather than distributions) and many of the arcs represented either logical relationships or were only present to signal that the child node has to be parametrised in different situations. Moreover, the dependence between some variables was better represented as undirected links. As a consequence, some parts of the model need to use the calculus behind a BN, whereas other parts can simply use mechanistic/mathematical relationships. Hence, we have built a BN-like model, rather than a “classical” BN. All parts however can be integrated to build a simulation model of the population of containers arriving at the border, and propagate it through the biosecurity system, to estimate a leakage rate, and similarly for the cut flowers pathway.

The full list of variables represented as nodes in the BN-like model for the container pathway is presented in Table C.1 and for the cut flowers pathway in Table C.2. The inter-dependencies between the nodes are depicted in Figures 5.1 and 5.2. Links amongst top-level nodes of each network (i.e., *CAL Country*, *Rural Destination*, *Type of Goods* and *SCHS* for containers, and *Source Country* and *Type of Goods* for cut flowers) represent dependence between those variables; this correlation structure is estimated during model parametrisation (Chapter 6).

5.1. The container pathway

The four key questions at the top of the conceptual diagram informed the top level variables in the BN which in turn influence the risk controls modelled. These are the variables: *CAL Country*, *Rural Destination*, *Type of Goods* and *SCHS* (Figure 5.1 and see annotated version in Appendix D, Figure D.1).

While three of the four key questions have a direct correspondent node in the model, the answer to the question: “was a container treated offshore under SCHS?” (yes/no) triggers different sampling protocols for inspection. These protocols depend on the reliability of the facility and the historical contamination rates. We captured these details in the node *SCHS reliability* which informs the *SCHS* node.

The first five nodes listed in Table C.1 can be considered containers’ properties. Different combinations of containers’ properties lead to different types of risk controls. This is represented by arcs pointing from the former to the latter. These controls influence the leakage rate which is the end point of the BN (last variables listed in Table C.1). Other variables which do not mirror “nodes” of the conceptual diagram were added to the BN structure, since the leakage rate depends on them. These are the true approach rate and sensitivity of different types of inspections.

The true approach rate is captured through two variables, which differentiate between internal and external contamination. The sensitivity of internal and external inspections is different for each inspection and this is captured by arcs pointing from each inspection type to the sensitivity nodes.

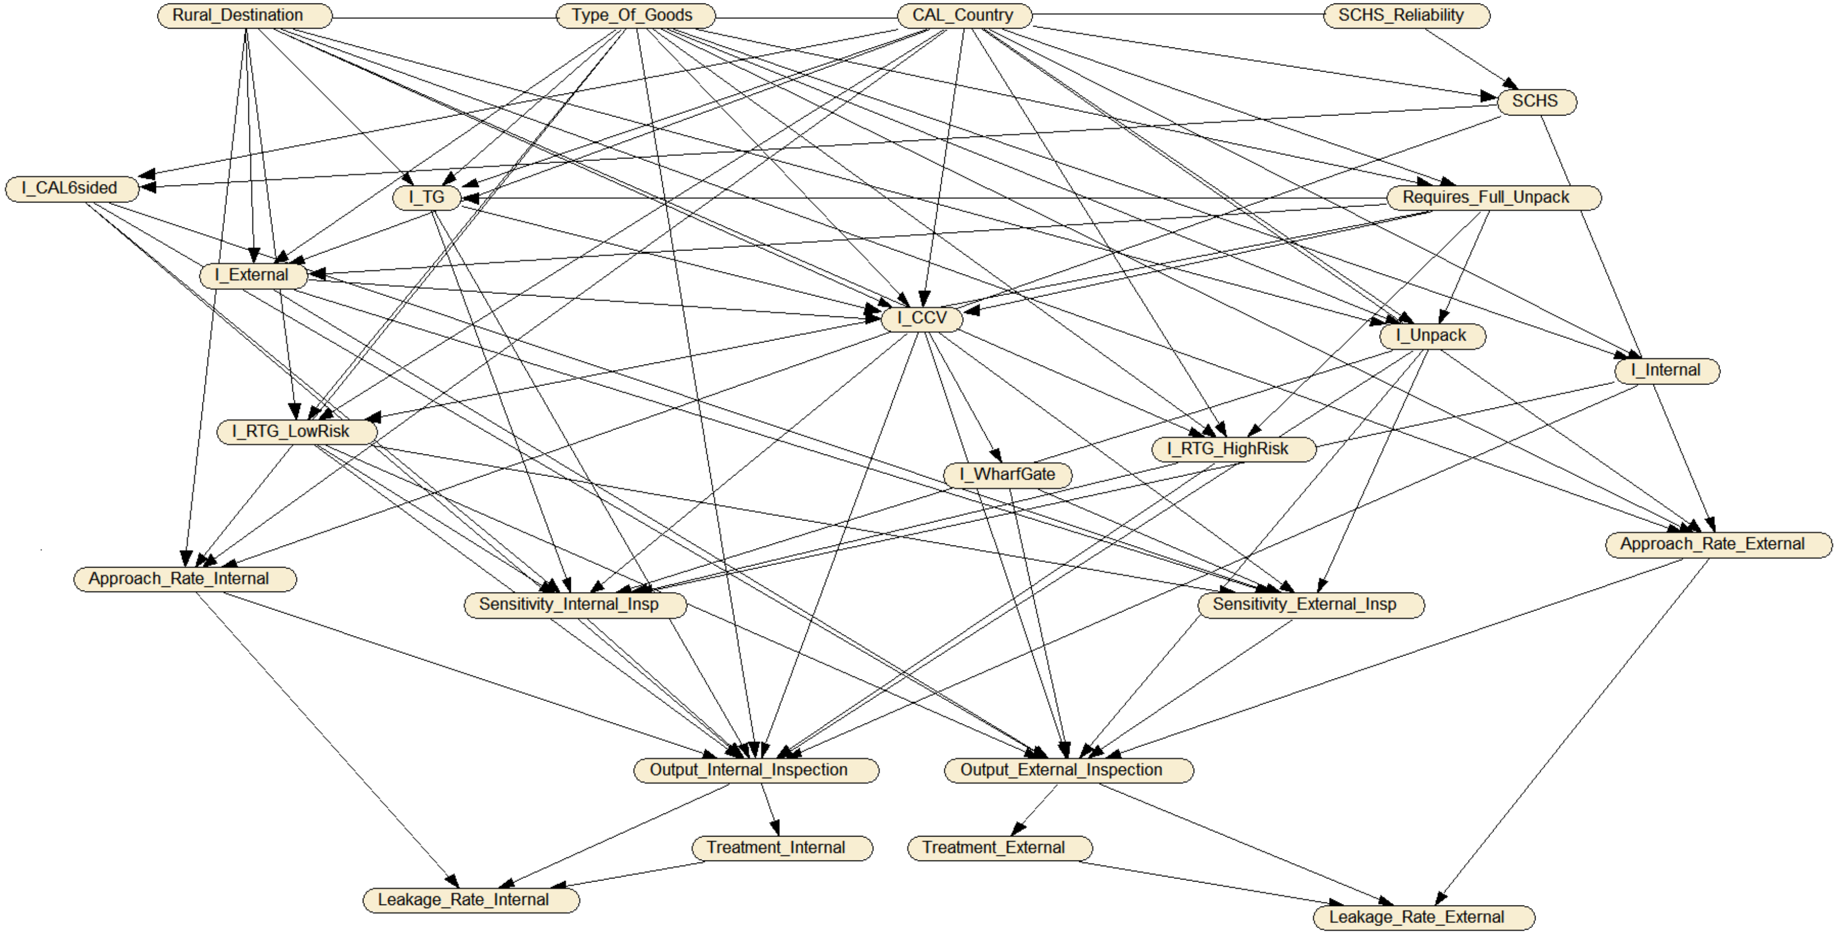

Figure 5.1.: The BN-like model of the container pathway.

5.2. The cut flowers pathway

The top-level correlated variables for the cut flowers pathway were *Source Country* risk and *Type of Goods* (cut flowers or foliage), as shown in Figure 5.2. To be able to use *Source Country* risk as a top level variable, the data on the potential risk of source countries needed to be discretised. CEBRA used information from two departmental documents[1] to determine contamination risk for each source country in the data. The tables provided information on how many pests were found in each country and on how many pathways. The number of pathways was taken into account. Countries that appeared more often as being potential locations/sources of pests were deemed more risky. Based on the frequency of appearances, countries were divided into roughly equal groups of 'high', 'medium' and 'low'. Figure 5.3 shows the level of contamination risk for all the countries cut flowers are imported into Australia from. Note that this map will be biased, as pest-risk knowledge across the world is uneven and we know more about risks from countries that export to Australia.

Documentation assessment is given as “InitialDocAssResult”, i.e. initial documentation assessment result. This is because consignments can have multiple documentation assessments (see the conceptual diagram in Figures 4.2, 4.3) near the beginning or just before release.

We consider two types of inspections: internal inspection (goods inspection) and external inspection (goods packaging inspection), each with their own contamination probabilities and inspection outcomes. Based on these outcomes, the line may be treated, and then released, or destroyed or sent back[2].

The inspection outcome can be: “clean” (i.e., no biosecurity risk material found), or “infested”(i.e., biosecurity risk material was found). “Clean” outcomes can be “compliant” or “non-compliant” typically representing compliance of the documentation and packaging.“Infested” outcomes can be “actionable” or “not actionable”, representing such flags in the data.

The true approach rate is captured in two variables for internal and external contamination and is estimated assuming 90% inspection sensitivity.

Note that the modelling did not take into account the percentage of stems in a line or consignment being inspected. We however expect this would be a straightforward modification, if data were available. This could involve either adding a new node into our BN-like network or lowering the inspection sensitivity to account for the fact that only a proportion of stems are inspected. However, determining the best method would require further work and suitable information.

---

[1] Final Pest Risk Analysis for Cut Flower and Foliage Imports - Part 1 (Table XVII) and Part 2 (Table VIII) (Department of Agriculture, 2019; Department of Agriculture, Water and the Environment, 2021)

[2] All the variables added below the “Treatment” and “DestroySendback” nodes represent simple deterministic calculations using the probabilistic variables described.

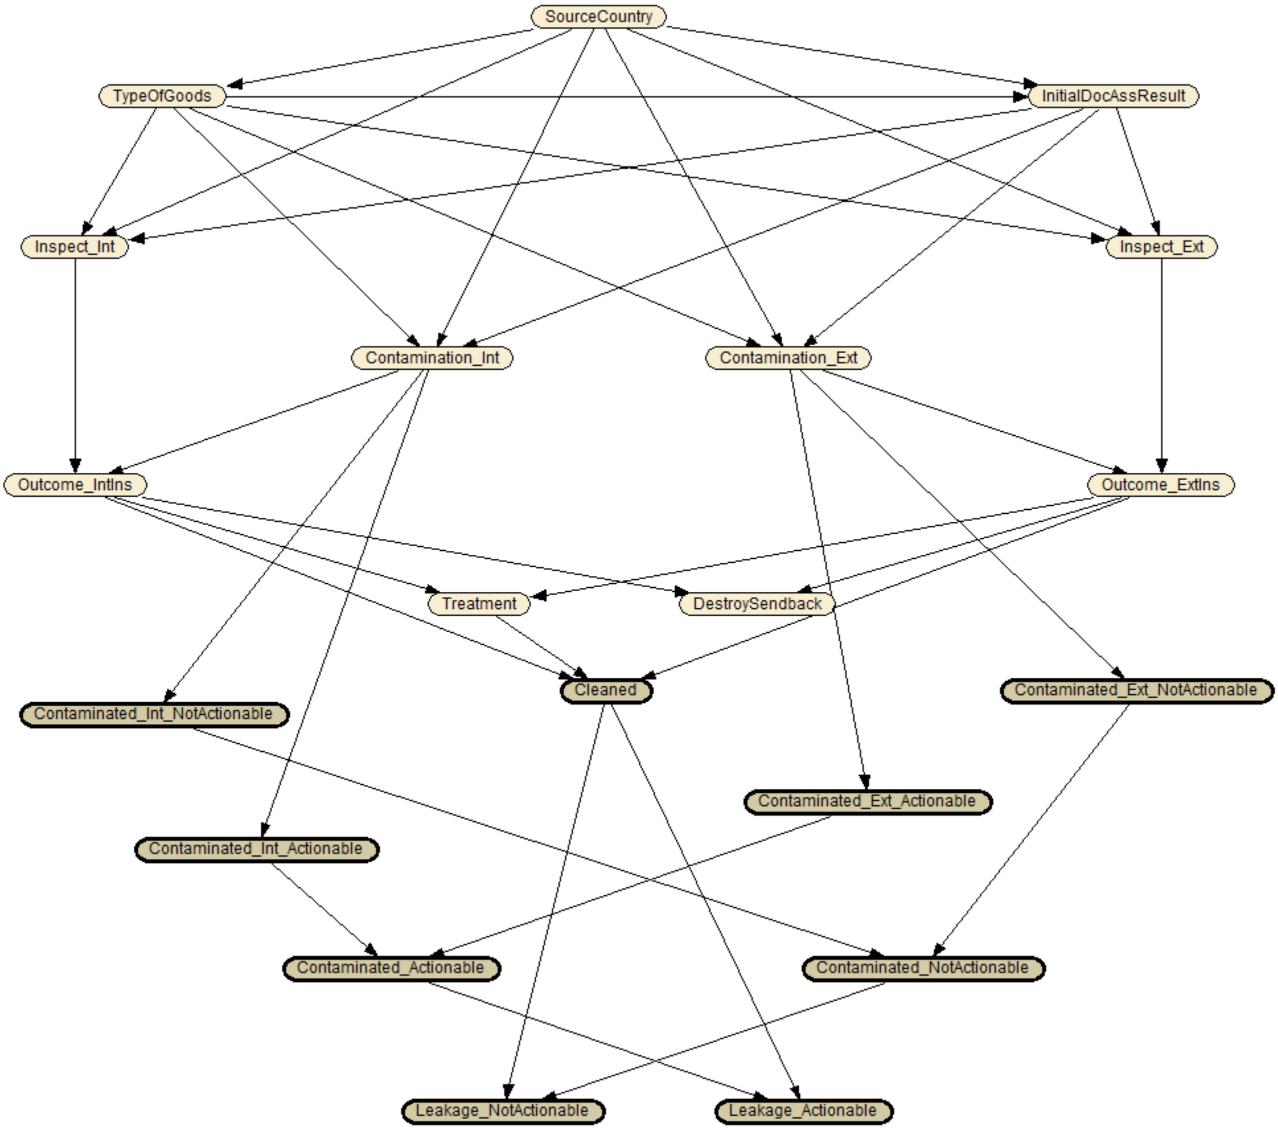

Figure 5.2.: The BN-like model of the cut flowers pathway.

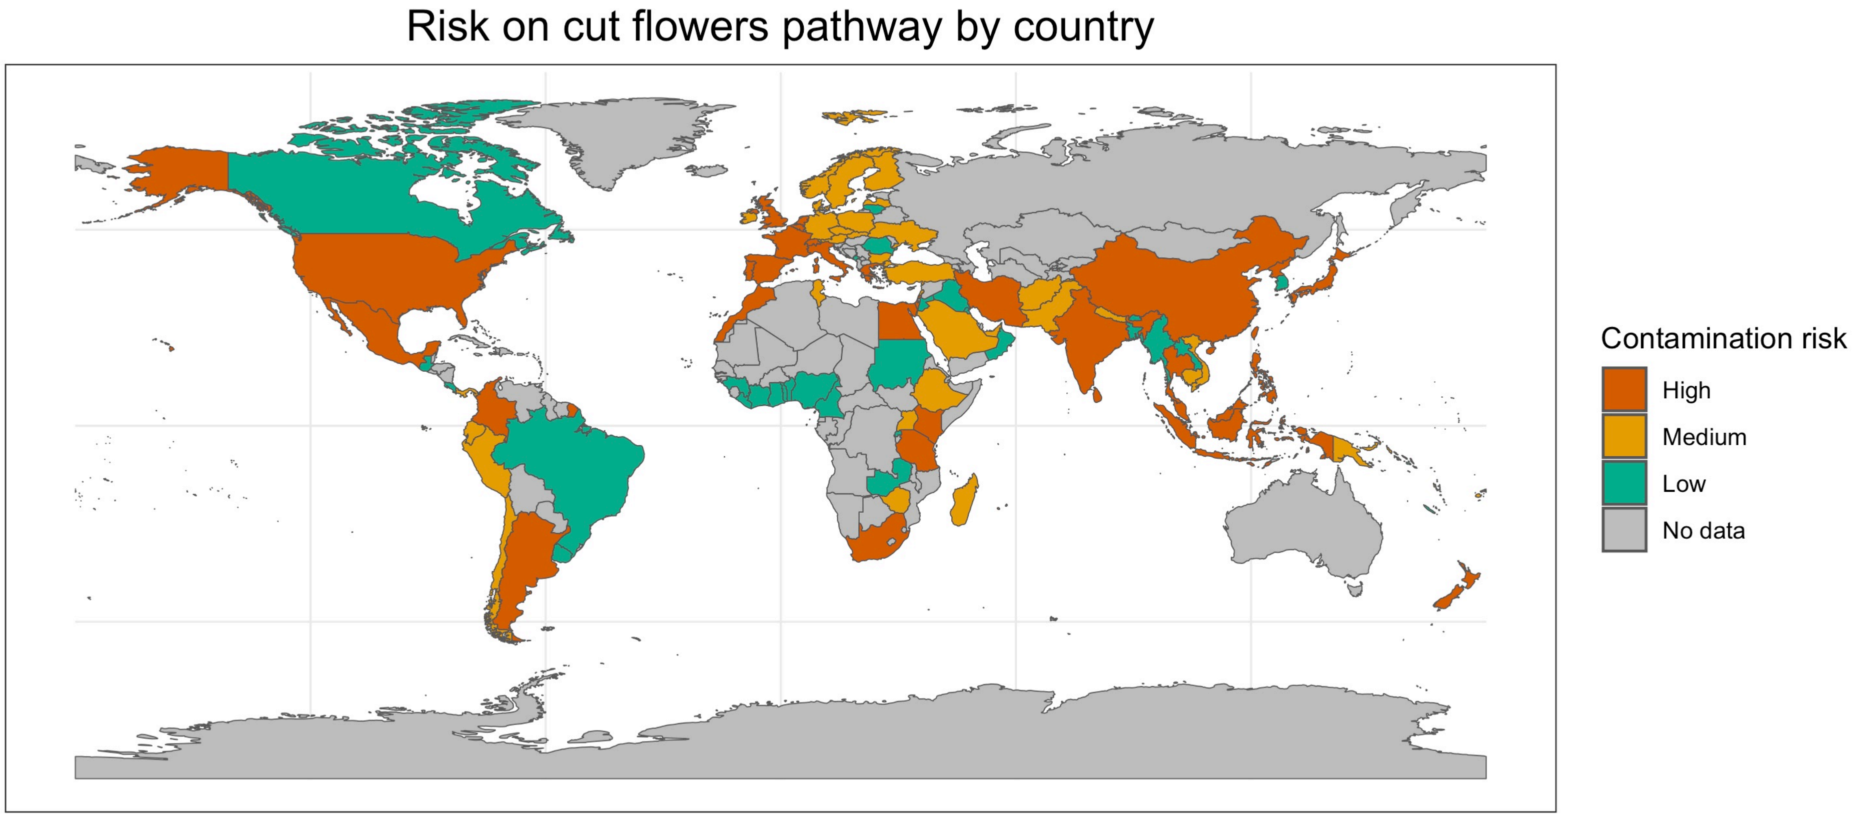

Figure 5.3.: Cut flower contamination risk by source country, based on Table XVII from Department of Agriculture (2019) and Table VIII from Department of Agriculture, Water and the Environment (2021).

6. Data processing and parametrisation

In this chapter we describe how the provided data were processed and then used to parametrise the models for the container (Section 6.1) and the cut flowers (Section 6.2) pathways. The parametrisation of both models is representative of the pathway configuration that existed at the time of model development.

6.1. The container pathway

The container pathway simulation model is produced by a set of R functions and related datasets[1], which together facilitate the construction of the parametrised BN.

Development of the simulation model of the container pathway requires parametrisation of the BN like model (i.e., estimation of the conditional distributions of child nodes' states given their parents' states) using data. The department provided multiple data sets, including data from S-Cargo and AIMS systems for sea containers, wharfgate inspection data, cargo compliance verification (CCV) data, and data regarding the containerised goods. These data were complemented with the list of countries on the Country Action List (CAL) from the DAFF website (DAFF, 2022). The model considers the cohort of containers referred to the Department, and does not capture leakage attributable to the non-referred container population.

All available datasets were combined, and data cleaning included: standardising variable names and country names, removing duplicate rows, and mapping port codes to countries. We only considered containers arriving in the period 2018 Q4 through 2022 Q3 (this excluded ∼530,000 containers arriving outside the considered time frame). This choice of period reflects our understanding that the border biosecurity system with respect to containerised goods underwent significant change in 2018, hence earlier data are not representative of the present system. Further containers were excluded, such as those lacking either an AIMS entry record or directions (∼900,000), as well as a relatively small number recorded as break-bulk or transhipment, or which lacked final direction. The attenuation in the number of containers remaining in the cohort, as inclusion criteria are applied, is shown in Figure 6.1.

A total of 3,582,363 containers were included in the raw extract provided by the department. After cleaning and filtering, 2,081,781 containers remained, and 2,050,726 (∼98.5%) of these conformed to the processes defined by the conceptual diagram and so were deemed suitable for parametrising the model. The cleaned data contains all attributes of interest, such as relevant port locations, rural/metro destination, type of containerised goods, whether any inspection was undertaken, and the inspection results and associated attributes/information needed for model parametrisation.

---

[1] These are available at `https://gitlab.unimelb.edu.au/jbau/pathways`

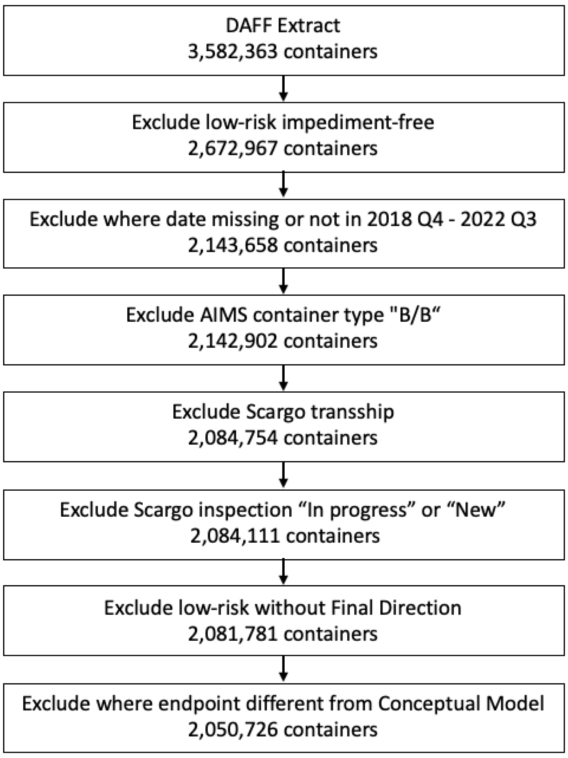

Figure 6.1.: Steps involved in cleaning the raw containers dataset. The number of containers remaining after each filtering step is given.

For most nodes, conditional probabilities of their states were estimated from the data as the relative frequencies of their states for each possible combination of their parents' states. An example is provided in Table 6.1, which gives the relative frequencies (as probabilities p1 and p2) of a CAL six-sided inspection occurring (or not), given the states of the *CAL Country* and *SCHS* nodes. Dependence amongst the nodes at the top level (*CAL Country*, *Rural Destination*, *Type of Goods*, and *SCHS*[2]), for which network structure was unknown, was estimated by applying a structure-learning algorithm that accounted for the relative frequencies of alternative combinations of these nodes observed in the container dataset. Specifically, a Hill-Climbing algorithm (implemented in the `hc` function provided by the `bnlearn` R package) was used to learn the correlation structure of these nodes based on a sample, with replacement and weighted by empirical frequencies, of 10 million containers from the full containers dataset. Subsequently, the conditional probabilities corresponding to that structure were estimated through maximum likelihood using the function `bn.fit`, again from `bnlearn`.

For each node representing the outcome of a particular inspection method, the relative probabilities of the states (i.e., “no contamination”, “low level of contamination”, or “high level of contamination”) were estimated from the data through proportional ordered logistic regression modelling. Internal and external contamination approach rates (i.e., the true proportions of containers that are contaminated) were subsequently estimated from these detection frequencies by assuming that all inspection methods have an inspection effectiveness (or “sensitivity”, the probability of correctly labelling a contaminated container as being contaminated) of 90%.

---

[2] The SCHS reliability was not used in this version of the model, so SCHS became the fourth key container property.

Table 6.1.: Conditional distribution of `I_CAL6sided`, 6-sided inspections for CAL containers. States of SCHS indicate whether containers did not pass through a SCHS facility (i.e., *no*), or else the SCHS facility's sampling rate at the point in time when the container arrived (*5/20/50/100%*).

CAL Country	SCHS	I CAL6sided=TRUE	I CAL6sided=FALSE
FALSE	no	p1	p2
FALSE	5%	p1	p2
FALSE	20%	p1	p2
FALSE	50%	p1	p2
FALSE	100%	p1	p2
TRUE	no	p1	p2
TRUE	5%	p1	p2
TRUE	20%	p1	p2
TRUE	50%	p1	p2
TRUE	100%	p1	p2

Inferring the approach rate from inspection outcomes and effectiveness:

$$\text{Approach rate} = \frac{\text{Detection rate}}{\text{Inspection effectiveness}}$$

Ideally, inspection effectiveness would be elicited from experts; however, for a proof-of-concept, this rather arbitrarily defined setting is adequate, and is readily adjusted to reflect further consultation with the department as well as any findings regarding the impact of uncertainty about this parameter on model outcomes (e.g., modelled leakage rates).

In all cases, frequencies were calculated at an annual time scale, supporting integration with the Value model (see Chapter 10). Further, while container data comprised inspection outcomes for each of seven contaminant types (snails, ants, insects, seeds, soil, other plant material, and other animal material), we presently consider these in aggregate.

Conditional probability tables such as that in Table 6.1 were used to construct a BN mirroring the probabilistic portion of the BN-like model described in Section 5.1, using the package `bnlearn` (version 4.7.1, Scutari, 2010) in R (version 4.2.1, R Core Team, 2022)[3]. This BN can be queried to identify the joint probability of any combination of states across multiple nodes. If required, quantities defined by functional relationships between nodes can be calculated outside of the BN, by querying the BN and passing the outcomes to separate R functions. We demonstrate some functionality in Chapter 7.

6.2. The cut flowers pathway

The development of the simulation model of the cut flowers and foliage pathway followed the same overall steps that we established for the container pathway, by parametrising the model using the cut flowers data provided by the department.

---

[3] R code provided at `https://gitlab.unimelb.edu.au/jbau/pathways`.

The department provided one large pre-combined data set with 101,022 total unique cut flowers and foliage lines, that was extracted from AIMS, for the period of November 2017 to February 2022 (inclusive). Note that a consignment is imported by a single importer, and it can have several lines. Since different lines in the same consignment can have different origins, flower types, etc.—and hence different levels of risk—we consider lines rather than consignments. Data cleaning involved removing the entries marked as 'Withdrawn Entry', 'Administration Withdrawn', or 'Admin Duplicated Entry' (266 lines in total). A total of 100,756 lines were used for parametrisation. Note that these lines include dried or preserved flowers imported during the same period.

Cleaning and categorising the remaining data required making various assumptions due to ambiguous results. Many lines had multiple directives of the same type, and/or had missing directives such as documentation processing – even though we expected that all consignments would have documentation checks according to the conceptual diagram. We made assumptions regarding their movement through the risk control pathway (such as assuming consignment release if no other information is present) (see Section 9.4 and Appendix G). There was also incorrectly filed data, which we could sometimes rectify by doing basic analysis on the free-text fields (which contained more information). However, it was infeasible for us to check every individual line of data.

As part of the data processing, country risk levels were needed, akin to the CAL for the container pathway. In absence of accessible BICON data tables, the country risk levels were determined using information available in Table XVII from Department of Agriculture (2019) and Table VIII from Department of Agriculture, Water and the Environment (2021), where countries were deemed more risky if they appeared more often as being potential locations/sources of pests with potential for establishment and spread within Australia, and with potential for economic consequences. Based on appearance frequency, countries were divided into roughly equally sized groups: "High", "Medium" and "Low".

It is worth noting that an important aspect of the cut flowers and foliage inspection is missing from the data (and indeed unrecorded overall): the proportion of stems inspected within any given filed inspection. Only a maximum of 600 stems are inspected in any consignment, but consignment sizes can vary widely. This variation was not taken into account in our modelling, i.e., we assumed that there was no contamination in the entire line if an inspection of part of a line was found uncontaminated. Note that this "clean" status is part of the data-cleaning step only. Given that we have a non-100% inspection sensitivity, the line could still be unknowingly contaminated and constitute a leakage. Leakages are estimated later, with the BN-like model.

Additionally, the results from real lines affect other lines in the same consignment, with treatment applied at the consignment level. However, we focused on the line level and so noted treatment information at the line level. However, it should be noted that "clean" lines and un-inspected lines can be treated, which we captured in the processed data.

We also note that while contamination results and actionable status in our dataset were recorded at the line level, it *might not be attributable to that line*, but may be a false attribution generated by the results from another line in the same consignment. When that is the case, such errors will propagate through the model and decrease the reliability of the results, as different lines can have different characteristics (from a different source country, containing different goods).

Once the data was cleaned and processed, we created the conditional probability tables needed for the parametrisation of the simulation model.

For most nodes, the conditional probabilities of their states were estimated from the data just like in the container pathway, from the top-level attribute distribution, to the various inspection results. An example of a conditional probability table used is Table 6.2, which gives the probability of one of three document assessment results (incomplete, not ok, ok) given the source country risk (low, medium, high) and goods type (flowers, foliage).

The external/goods-packaging and internal/goods contamination approach rate was estimated by assuming that inspection had an effectiveness of 90%.

Again, note that we did not include any scaling factor to account for the fact that at most 600 stems are inspected, regardless of the consignment size. The inspection effectiveness could be reduced to account for this, but we do not have any data on the distribution of consignment sizes.

Table 6.2.: Conditional probability distribution for the Initial Documentation Assessment Result node, `InitialDocAssResult` = incomplete, not ok, ok.

SourceCountry	TypeOfGoods	incomplete	not ok	ok
low risk	flowers	0.146153846	0.076923077	0.776923077
low risk	foliage	0.602040816	0.102040816	0.295918367
medium risk	flowers	0.016101546	0.018951637	0.964946818
medium risk	foliage	0.138543517	0.079928952	0.781527531
high risk	flowers	0.028601029	0.049474346	0.921924625
high risk	foliage	0.081678956	0.103658537	0.814662507

7. Demonstrating the functionality of the simulation models

This chapter demonstrates how the simulation models work, in particular how to construct (Section 7.1) and query (Section 7.2) the network, and how to visualise the flow of containers (Section 7.3). These sections largely focus on the BN for the container pathway, but similar functionality is possible for the cut flowers pathway.

To construct the BN for the two pathways, we developed the `pathways` R package[1]. Included in the package are pre-defined container and cut flowers pathway models, as well as the raw conditional probability tables and code used to create them. In this chapter, we show how to use the package to customise these BNs by varying the historical and future inspection effectiveness, and/or the treatment effectiveness, and then demonstrate how nodes of the network can be queried to obtain joint and conditional probabilities of multiple attributes and inspection activities. Finally, we demonstrate how the models can be used to simulate scenarios of altered biosecurity controls. Note that the values presented here are influenced by assumptions made about treatment effectiveness when parametrising the model, and should not be used in production.

7.1. Constructing the network

The `pathways` package includes the in-built objects `containers` and `cutflowers`, each of which is a BN object (class `grain`) that has been pre-constructed using the raw data and scripts included in the "data-raw" folder of the package source. Briefly, these networks were created by firstly learning the top-level node correlation structure, and parsing and combining conditional probability tables (CPTs) structured as .csv files. In this preliminary process, we used the included `define_bn()` function to generate a BN from these combined CPTs with class `grain` or `bn.fit`. Users can refer to the package documentation and source files for further detail. It is also possible to call `define_bn` manually, passing in a set of custom CPTs, if required.

These initial, prebuilt pathway models are complete and can be interrogated, but they assume an inspection effectiveness of 100% across all inspection methods. This is likely inappropriate, and so the function `make_pathway()` can be used to modify the existing networks according to scenarios of inspection and/or treatment effectiveness. The assumption about inspection effectiveness can have profound consequences on leakage estimates, since, together with the modelled (data-driven) detection probabilities embedded in the CPTs underlying the prebuilt models, it determines the internal and external approach rates for each combination of top-level nodes.

To create a container pathway BN that assumes 90% inspection effectiveness and 100% treatment effectiveness, we can do:

```
net <- make_pathway('containers', insp_eff=0.9, treat_eff=1)
```

---

[1] https://gitlab.unimelb.edu.au/jbau/pathways

The default BN class is `grain` (i.e. compatible with the `gRain` R package); to instead request a `bnlearn` network, add the argument `type='bnlearn'` to the above command. Each has its advantages, but here we use the grain implementation for its ease and efficiency in querying the network.

7.2. Querying the network

We can query the network to extract quantities of interest from the parametrised model. For example, given fitted model `net`, we can use the `pEvidence` function provided by the `gRain` package (version 1.3.12, Højsgaard, 2012) to extract the proportion of containers that are destined for rural locations (`RuralDest = 'TRUE'`) and have a high level of external contamination (`Contamination_Ext = 'HLC'`):

```
pEvidence(net, c(RuralDest='TRUE', Contamination_Ext='HLC'))
## [1] 0.004345568
```

This result indicates that the joint probability of the two events (i.e., the probability of both events occurring together) is roughly 0.4%. We can also identify the proportion of containers that have a rural destination, have a high level of contamination, and are detected as contaminated by the biosecurity inspections:

```
pEvidence(
  net,
  c(RuralDest='TRUE', Contamination_Ext='HLC', Outcome_Ext_HLC='TRUE')
)
## [1] 0.003786325
```

We see here that around 87.1% (0.003786325/0.004345568) of the containers with high external contamination, which are heading to rural destinations, are captured and treated by the biosecurity system[2]. It follows, then, that the remaining 12.9% (= 0.06% of all containers) represent leakage, and we can confirm this by querying the model:

```
pEvidence(
  net,
  c(RuralDest='TRUE', Contamination_Ext='HLC', Outcome_Ext_HLC='FALSE')
)
## [1] 0.0005212736
```

This can be easier to see by querying conditional, rather than joint, probabilities using the `qgrain` function[3]. The following shows the probability of one or more inspection methods detecting a high level of external contamination, given that the container does indeed have external HLC, and is destined for a rural area.

```
library(gRain)
qgrain(
  net,
  nodes=c('Outcome_Ext_HLC'),
  evidence=list(Contamination_Ext='HLC', RuralDest='TRUE'),
  type='conditional'
)
## Outcome_Ext_HLC
##     FALSE      TRUE
## 0.1286927 0.8713073
```

---

[2] Once again, this fraction depends heavily on assumptions made about the effectiveness of inspection methods.

[3] For calculating leakage rates, the `leakage()` function contained in the repository is convenient.

Queries of arbitrary complexity can be made in this way, by specifying the nodes of interest for observation, and optionally fixing the states of other nodes to make such conditional queries.

7.3. Visualising the flow of containers throughout the system

By querying the probability of interesting combinations of nodes' states, we can obtain the quantities necessary for generating a diagram showing the volumes of containers passing through the various controls to various endpoints. Sankey diagrams are a useful tool for inspecting such flows, and the repository includes a function (`flows`) that simplifies generation of these figures.

Figure 7.1 shows the volume of CAL containers moving through the system, and Figure 7.2 shows the same for non-CAL. Finally, Figure 7.3 combines CAL and non-CAL containers in order to see their relative volumes. These images are ordinarily interactive, and hovering the cursor over a location reveals the volume of containers at that point. Note that non-referred ("no-risk") containers are not included here.

These three diagrams are produced by the following code:

```
flows(net, N=550000, CAL='yes')
flows(net, N=550000, CAL='no')
flows(net, N=550000, CAL='all')
```

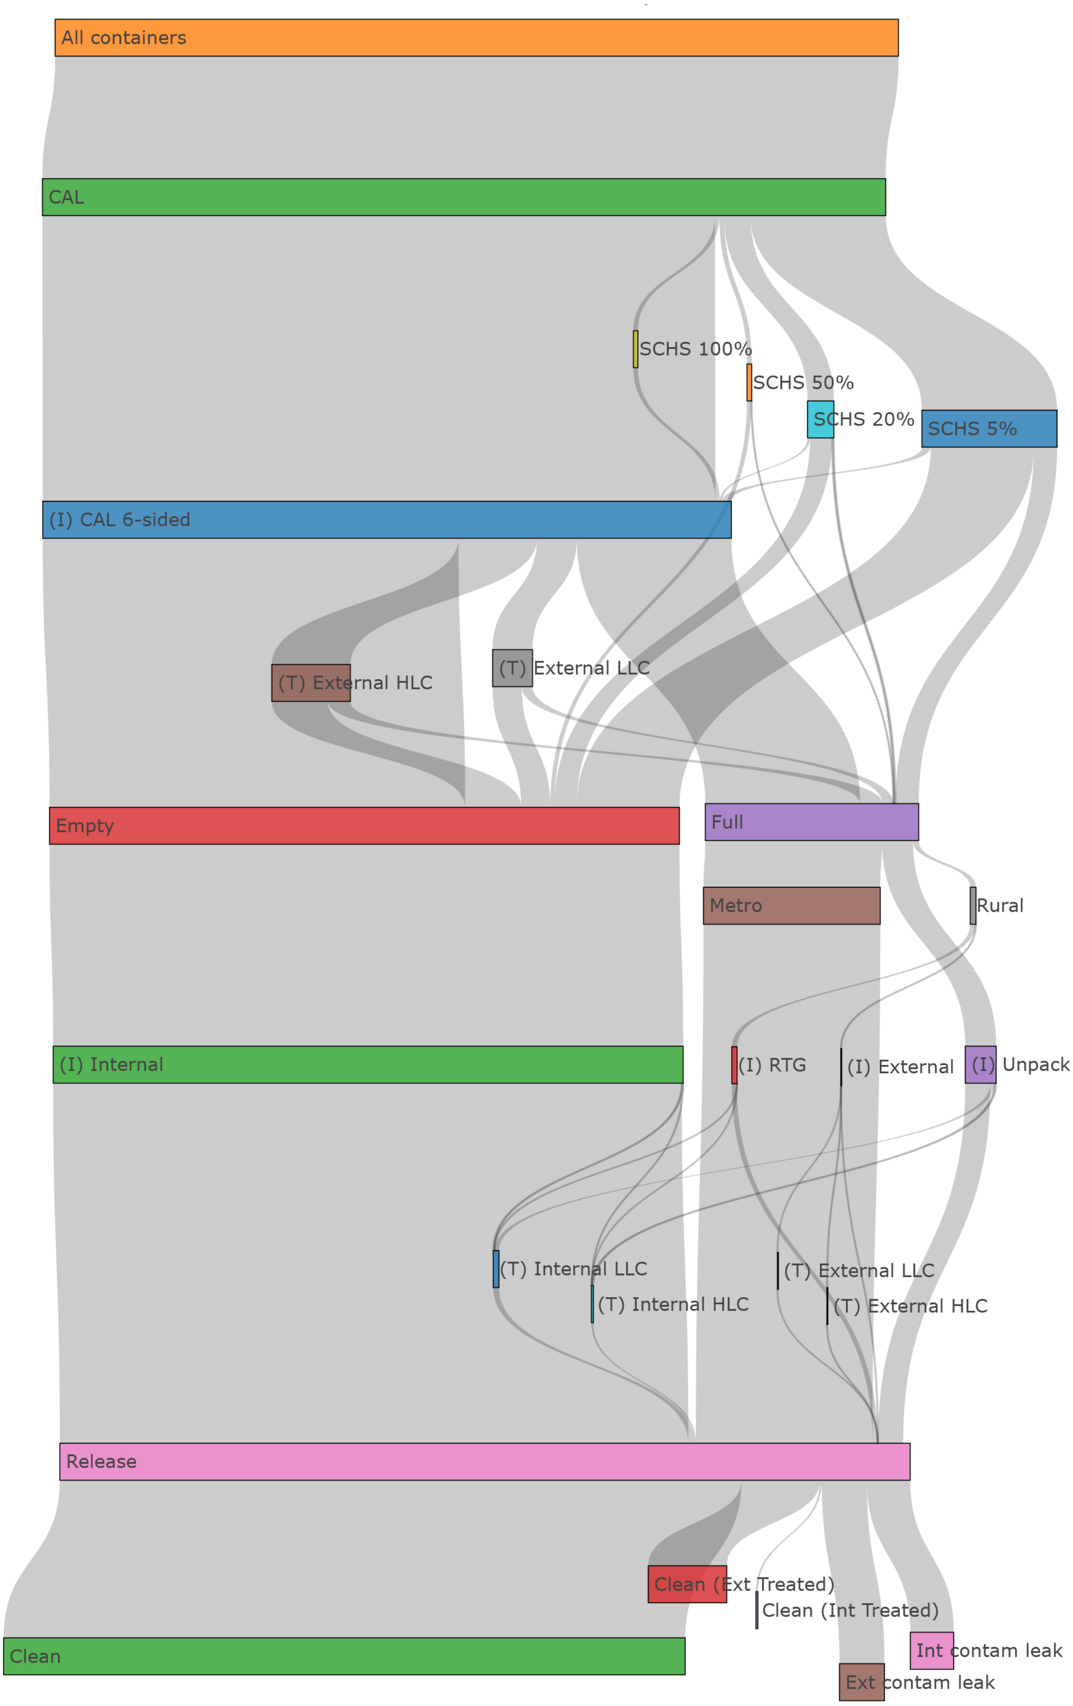

Figure 7.1.: Flows of containers from CAL origins/ports. “(I)” indicates an inspection, and “(T)” indicates a treatment.

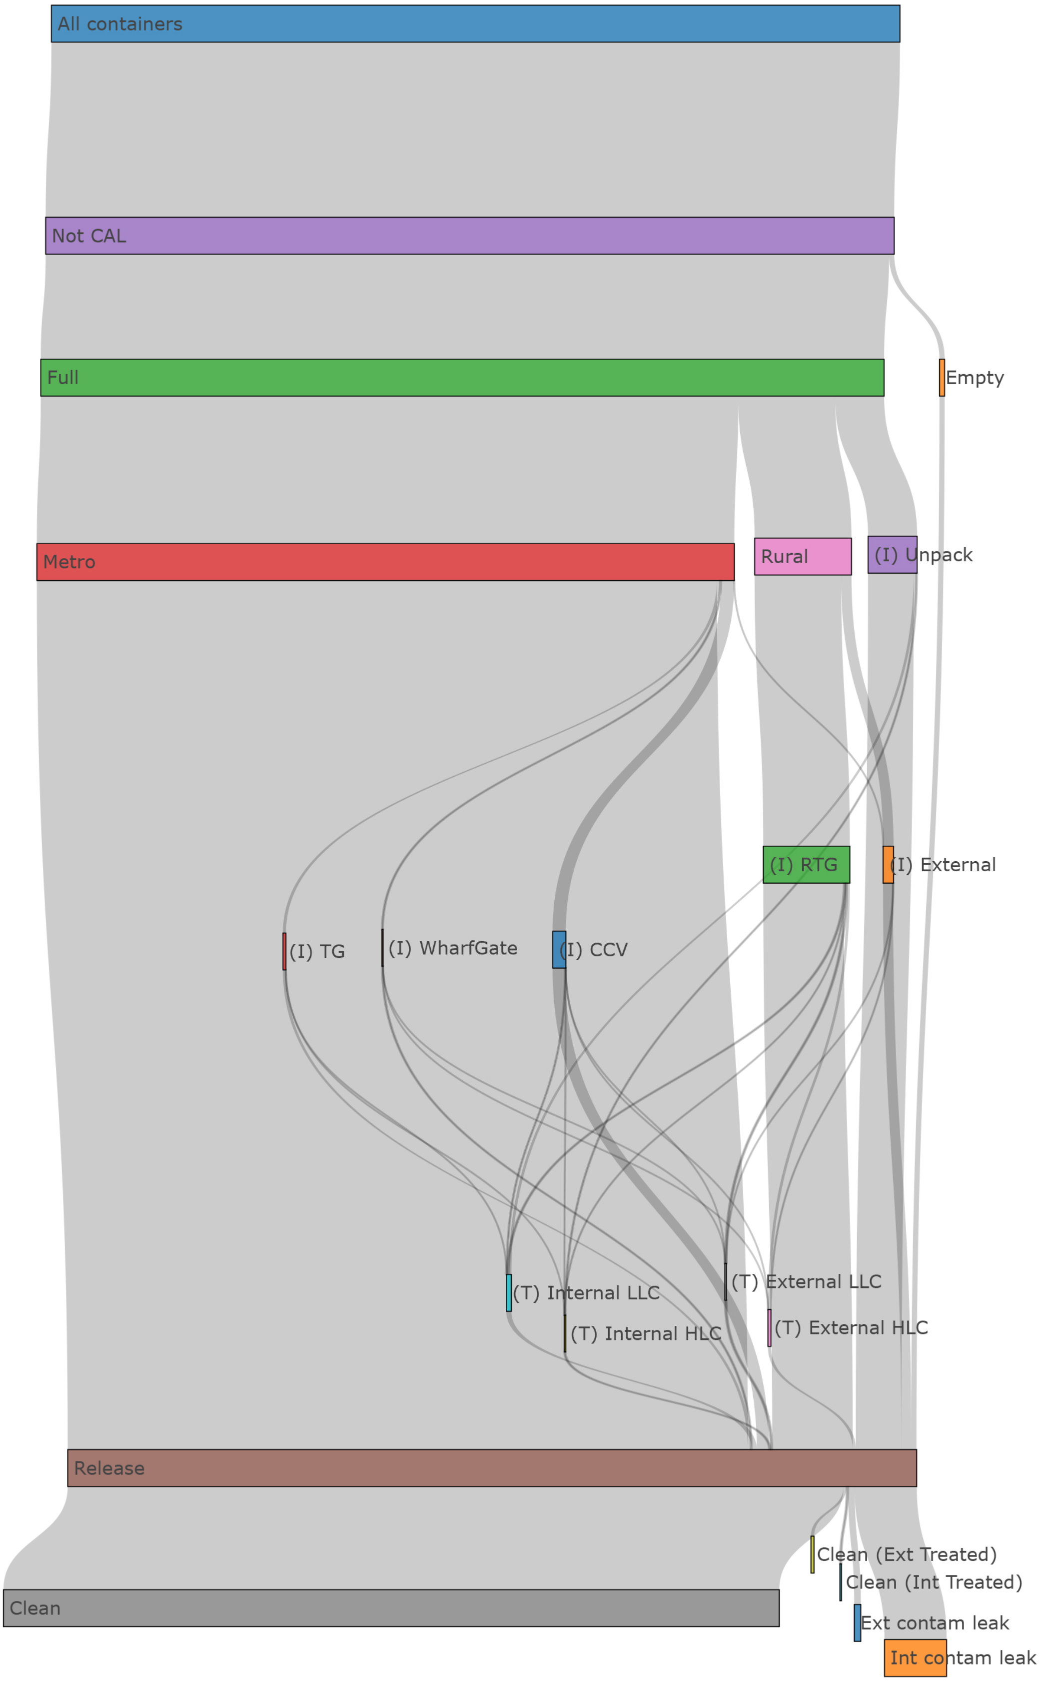

Figure 7.2.: Flows of containers from non-CAL origins/ports. "(I)" indicates an inspection, and "(T)" indicates a treatment.

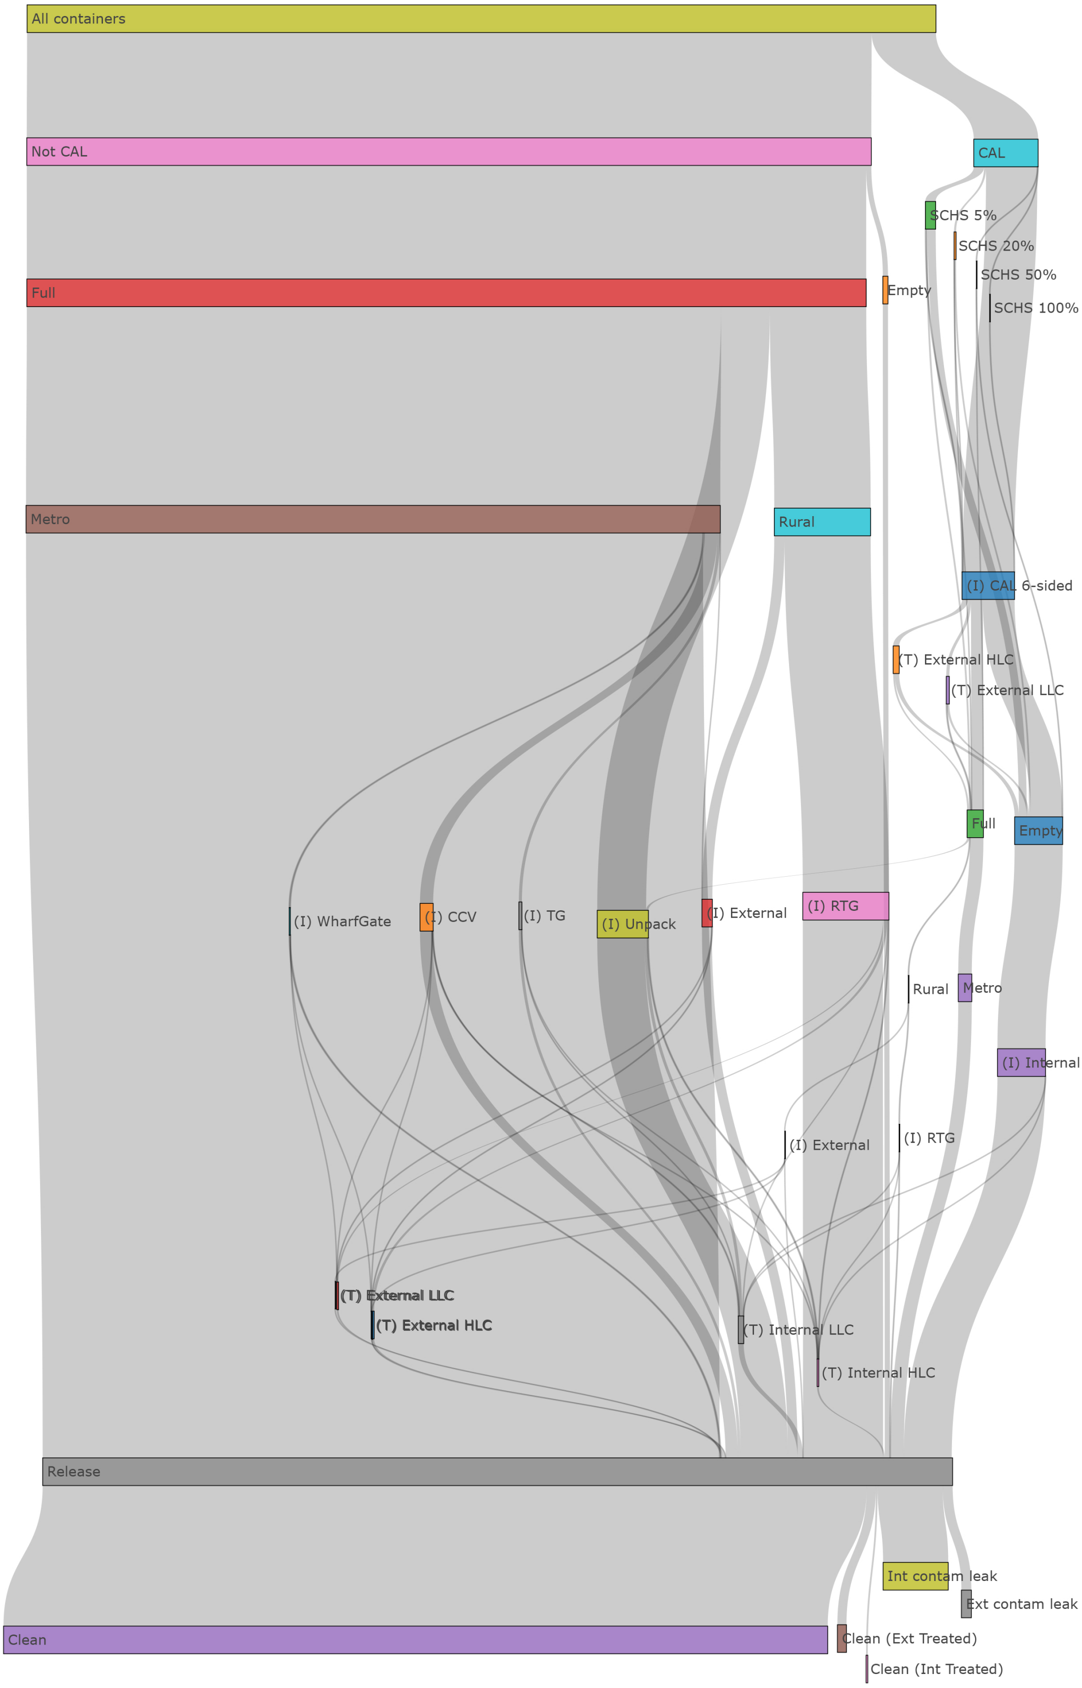

Figure 7.3.: Flows of containers from CAL and non-CAL origins/ports. “(I)” indicates an inspection, and “(T)” indicates a treatment.

8. Analysis of the container pathway

The purpose of the simulation model is to describe the complex relationships between container properties and risk controls, and to reveal the consequences, for contaminant leakage, of varying these controls. We demonstrate use of the model in the following sections, where we first examine the uncertainty in outputs (8.1) and then the impacts on the leakage rate of a range of counterfactual scenarios (8.2) and changes in inspection and treatment effectiveness (8.3).

Note that the leakage rates stated in this chapter are modelling estimates, involving various simplifications and caveats, and are underpinned by several assumptions especially where data gaps exist.

8.1. Uncertainty in outputs

Uncertainty was calculated by using simulations of 550,000 containers, which is the approximate referred container population per year. Figure 8.1 shows the uncertainty around the model's point estimates of leakage rate corresponding to each combination of top-level node states. Red curves show the binomial sampling distributions centred on modelled leakage rates; the close match to simulated outcomes indicates that the nodes between the top-level and the leakage rate are largely either deterministic given their parents, or explain little from the remaining variability in leakage rate. Many of the combinations of top-level node states corresponding to the higher leakage rates were infrequent and thus contribute relatively little to overall leakage (e.g., the combination CAL, Rural, Exceptions, SCHS 100% has a leakage rate of approx. 31%, but the combination only occurs with frequency < 0.001%).

8.2. Counterfactual analysis

In this section, we will assess the influence on the leakage rate of six scenarios: changes to the proportion of countries on the CAL (Section 8.2.1), changes to the proportion of CAL containers that are empty (Section 8.2.2), changes to the rate of full CAL containers going to rural vs. metro destinations (Section 8.2.3), changes to the proportion of containers participating in the SCHS (Section 8.2.4), changes to the inspection rate of SCHS containers (Section 8.2.5), and changes to the proportion of containers requiring a full unpack (Section 8.2.6).

Many of the following figures are presented in side-by-side panels to contrast various scenarios. Please note that the two panels (left hand side and right hand side) of the figures often have different scales for their y-axis, and care should be taken in comparing effects to avoid misinterpretation.

8.2.1. Increasing the proportion of CAL-like countries

This counterfactual analysis changed the proportion of countries on the CAL, i.e. simulating the addition or removal of countries from the list, and investigated the impact on the leakage rate. In Chapter 3.6, we emphasised the importance of reparametrising the model when the relationship between top-level node states and contamination rates is modified. However, for simplification (and since reparametrisation is beyond the scope of this report), we assume here that when a country is added to the CAL, the containers originating from that country inherit the attributes of a typical CAL container in terms of goods type, destination, SCHS scheme, and contamination approach rate.

Figure 8.2 shows the changes in the external (low contamination, shown in or-

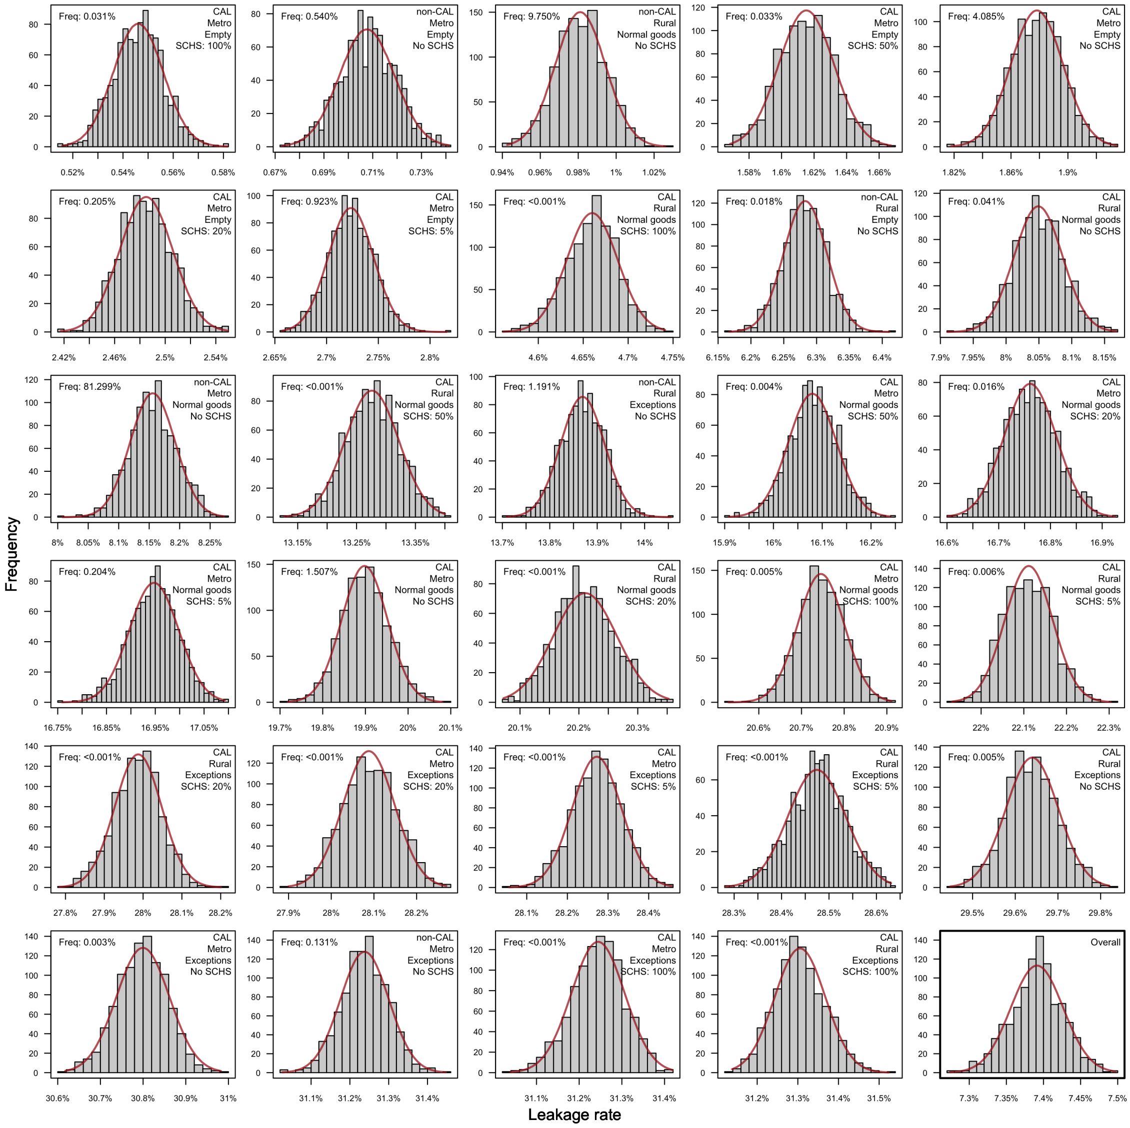

Figure 8.1.: Leakage rate distributions for different types of containers. Each panel shows the result, as a distribution of 1,000 model runs, using a specific combination of the states of the top level nodes. Other than the lower-right panel, which gives the unconditional leakage rate (i.e., across all combinations of top level node states), panels are presented in order of mean leakage rate.

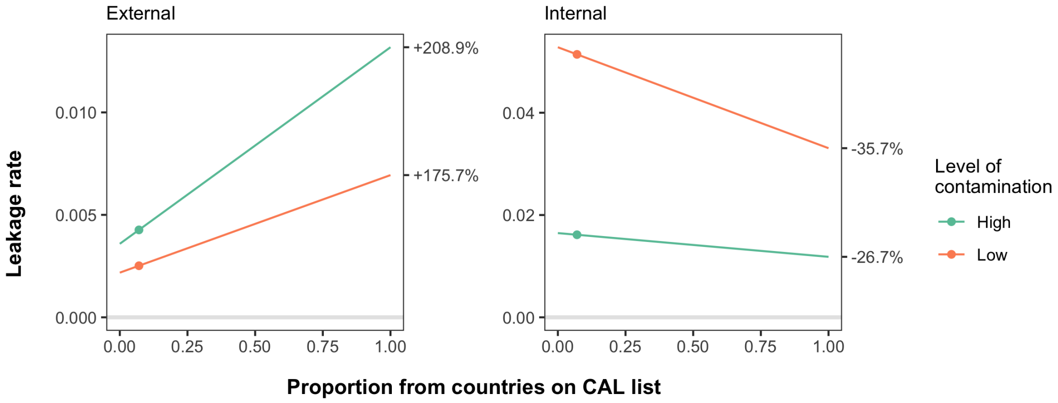

Figure 8.2.: Increasing the proportion of CAL-like countries. Points on lines indicate the leakage rates corresponding to the status quo of the proportion of containers on the CAL (i.e., 7.07% of containers are on the CAL). High levels of external and internal contamination are represented by green lines and low levels of contamination by orange lines. Percentages indicated in the right margins are changes in leakage relative to the status quo.

ange and high contamination, shown in green) contamination on the lhs panel and the changes in the internal (low/orange and high/green) contamination on the rhs. While leakage of internal contamination decreases with an increase in the proportion of countries on the CAL, the external contamination leaked presents an increasing pattern, suggesting more countries with CAL country properties will present a larger threat in terms of external contamination. This is not a surprising result given that CAL containers have a much higher external approach rate than non-CAL containers (approx. ratio of 13:1).

8.2.2. Changes in the proportion of empty CAL containers

This analysis increased or decreased the proportion of empty CAL containers and calculated changes to the leakage rate. Figure 8.3 shows no change in the leakage rate in terms of external contamination, as the number of empty CAL containers increases. One reason may be that in this situation, less containers are directed to CAL inspection (empty CAL containers: 80%, full containers: 88%) because empty containers are more likely to be on the SCHS and therefore exempt from inspection. An indication of the same can be observed if we query the network (as in Section 7.2) and investigate what the external leakage is more sensitive to when we condition on CAL countries alone (`pEvidence(net, c(CALSourceCountry='TRUE'))`). Under this condition the external leakage rate does not change for normal goods (16.7%) and empty containers (15.0%), as the high approach rate of the CAL containers is driven by the "exception" goods (48.9%). The internal contamination leakage rate presents a slight decrease, as more containers are directed to internal inspections.

8.2.3. Changes in the rate of full CAL containers going to rural vs. metro

This analysis changed the proportion of full CAL containers going to rural destinations. This change may reflect a change in the classification of some postcodes from ru-

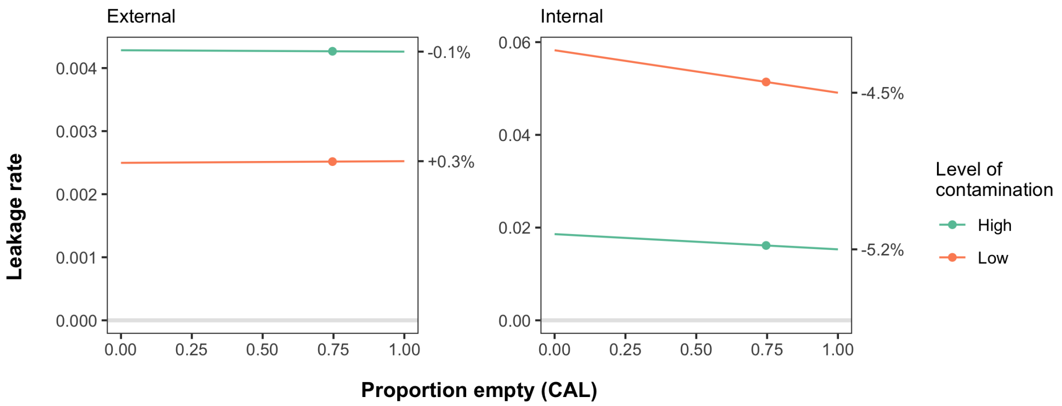

Figure 8.3.: Increasing or decreasing the proportion of empty CAL containers. Points on lines indicate the leakage rates corresponding to the status quo of the proportion of empty CAL containers. High levels of external and internal contamination are represented by green lines and low levels of contamination by orange lines. Percentages indicated in the right margins are changes in leakage relative to the status quo.

ral to metro (or the other way around). Exactly matching the proportion of postcodes which were reclassified through the CEBRA Rural postcode classification project with the resulting proportions for our scenarios would require extra data analysis and is therefore out of scope. However, those proportions will be captured, but not exactly identified on the plots presented in this subsection.

Figure 8.4 shows how the leaked external (left figure) and internal (right figure) levels of contamination change with changing number of CAL containers, carrying normal goods, destined for rural postcodes. A slight increase in the external contamination is still observed (similar to the increase in Figure 8.2) most likely driven by the high external approach rate for CAL containers in general. An increase in containers carrying goods to rural destinations leads to an increase in rural tailgate inspections. However, for CAL containers, these inspections are internal-only, so the present counterfactual scenario does not result in an increase in the number of external inspections, but does explain the decrease in leakage of internal contamination.

8.2.4. Changes in the proportion of containers participating in the SCHS

Not all containers from CAL countries participate in the SCHS. This analysis explored the effects of adding more countries to the SCHS. Changing the proportion of CAL containers carrying normal goods and empty containers participating in the SCHS (i.e., reallocating an increasing proportion of containers not participating in SCHS into the 100% sampling rate category) had a negligible effect on the leakage rate (Figures E.1 and E.2).

8.2.5. Changes in the inspection rate of SCHS containers

This counterfactual scenario investigated the relationship between the inspection rate of CAL containers participating in the SCHS, and the overall leakage rate. The inspection rate of the whole cohort of containers was varied. As expected, the effect on the

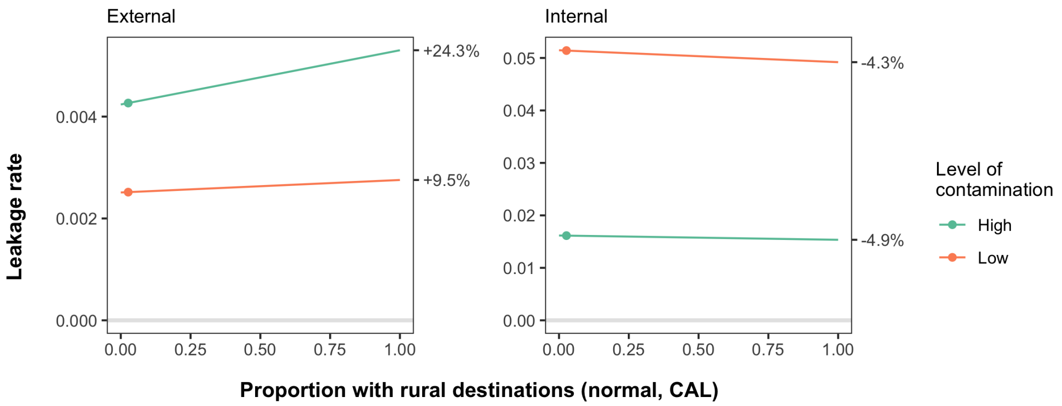

Figure 8.4.: Increasing the proportion of full CAL containers going to a rural destination. Points on lines indicate the leakage rates corresponding to the status quo of the proportion of full CAL containers going to a rural destination. High levels of external and internal contamination are represented by green lines and low levels of contamination by orange lines. Percentages indicated in the right margins are changes in leakage relative to the status quo.

leakage rate was insignificant because of the small sample size of SCHS containers (less than 2%, Figure E.3).

8.2.6. Changes in the proportion of containers requiring a full unpack

This analysis examined how changing the rate of inspections of goods (full unpack) affects the leakage rate. The proportion of containers requiring full unpack at an AA was changed for different combinations of the top level attributes, i.e., containers going to a metro destination originating from a CAL/non-CAL country and carrying normal goods/goods exempt from internal inspection.

Subjecting CAL or non-CAL containers heading to a metro destination with goods that are exempt from internal inspection to a full unpack did not affect the leakage rate (Figures E.4 and E.5). The leakage from CAL containers with normal goods going to a metro destination decreased slightly with increasing full unpack rates (Figure E.6). It is worth mentioning that these changes are for the entire population of referred containers, and the chosen cohort represents a very small fraction of this population. What seems like an insignificant system-wide impact (in terms of the reduction of overall leakage) may (and does) come from a significant impact on the selected cohort.

In contrast, increasing the proportion of non-CAL metro containers with normal goods undergoing a full unpack significantly decreased the leakage of internal contamination, but did not affect leakage of external contamination (Figure 8.5).

This is unsurprising given that the cohort of non-CAL metro containers with normal goods represent the vast majority of all containers. We conjecture that the slight increase in the external contamination leaked under these circumstances is driven by the lack of external inspections when a full unpack is undertaken.

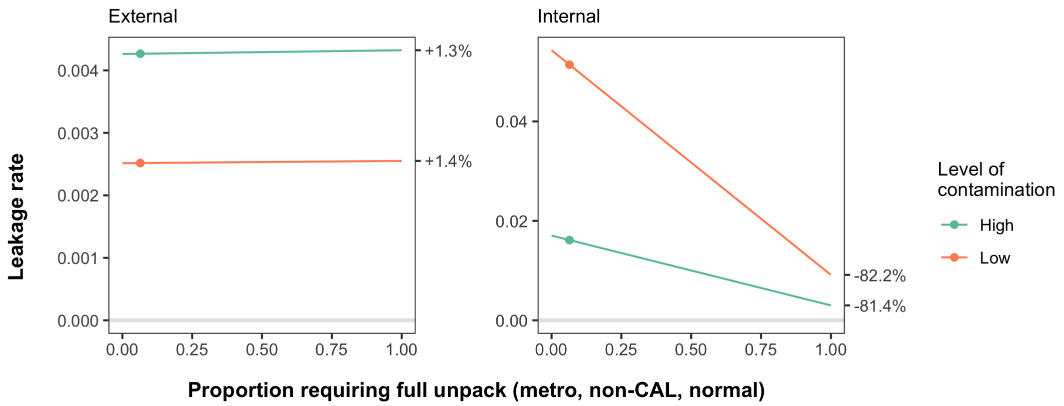

Figure 8.5.: Changing the proportion of non-CAL metro containers with normal goods undergoing a full unpack. Points on lines indicate the leakage rates corresponding to the status quo of the proportion of non-CAL metro containers with normal goods that are being subjected to a full unpack. High levels of external and internal contamination are represented by green lines and low levels of contamination by orange lines. Percentages indicated in the right margins are changes in leakage relative to the status quo.

8.3. Sensitivity analysis

The sensitivity analysis for the container pathway investigated how changes in inspection effectiveness (Section 8.3.1) and in treatment effectiveness (Section 8.3.2) affect the leakage rate.

8.3.1. Changes in inspection effectiveness

Here we examine the implications of modifying our assumption about the historical inspection effectiveness inherent in the container dataset used to parametrise the network. Using the `sensitivity` function, we can examine the change in leakage (or the values of any other nodes we specify) attributable to changes in inspection effectiveness and/or treatment effectiveness. Below, we vary inspection effectiveness from 70% through 100%, for each of three levels of treatment effectiveness (50%, 75%, and 100%).

```
sens_insp_eff <- sensitivity(
  pathway = 'containers',
  nodes = c('Leakage', 'Leakage_Int_HLC', 'Leakage_Int_LLC',
            'Leakage_Ext_HLC', 'Leakage_Ext_LLC'),
  insp_eff = seq(0.7, 1, 0.01), treat_eff = c(0.5, 0.75, 1)
)
```

The resulting table (below) shows the probability of each of the specified nodes' states, for each combination of inspection and treatment effectiveness, and the corresponding curves reveal the changes of interest. Figures 8.6 and 8.7 show that inspection effectiveness strongly influences the leakage rate. External inspections that are more effective in detecting high levels of contamination reduce the leakage rate significantly. The use of more effective treatment methods further reduces the leakage rate.

```
> sens_insp_eff
# A tibble: 930 × 5
   insp_eff treat_eff node    state  prob
      <dbl>     <dbl> <chr>   <chr> <dbl>
 1     0.7       0.5  Leakage FALSE 0.887
 2     0.7       0.75 Leakage FALSE 0.893
 3     0.7       1    Leakage FALSE 0.899
 4     0.71      0.5  Leakage FALSE 0.889
 5     0.71      0.75 Leakage FALSE 0.895
 6     0.71      1    Leakage FALSE 0.900
 7     0.72      0.5  Leakage FALSE 0.891
 8     0.72      0.75 Leakage FALSE 0.896
 9     0.72      1    Leakage FALSE 0.902
10     0.73      0.5  Leakage FALSE 0.892
#      with 920 more rows
#      Use `print(n = ...)` to see more rows
```

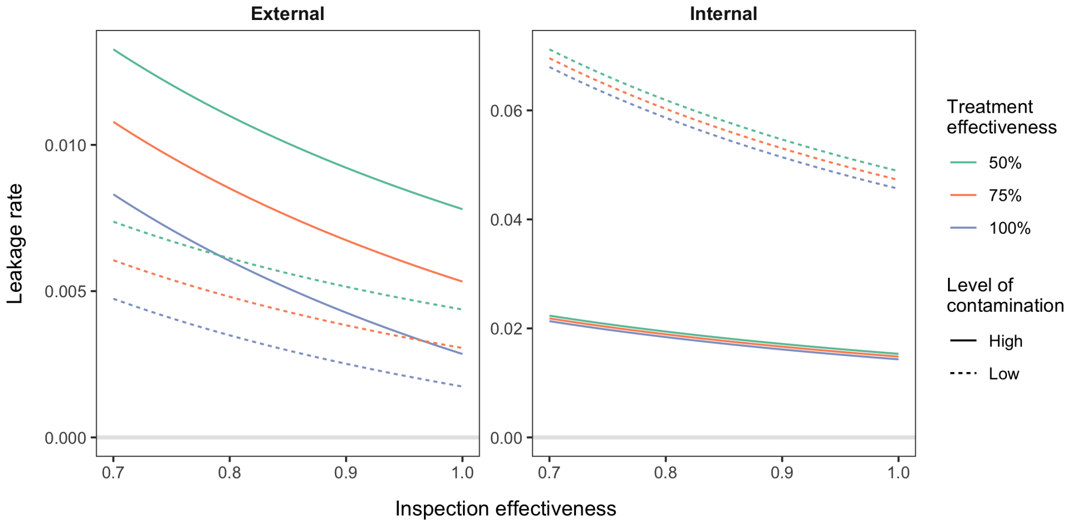

Figure 8.6.: Leakage rate as a function of external and internal inspection effectiveness at three levels of treatment effectiveness (50% - green line, 75% - orange line, and 100% - blue line). Solid lines represent high levels of external contamination and dashed lines represent low levels of contamination.

8.3.2. Changes in treatment effectiveness

In the container model, treatment effectiveness is assumed to be 100%, i.e, the applied treatment removes all of the detected high level contamination[1]. For the sensitivity analysis we have varied treatment effectiveness, as the proportion of contaminated containers after treatment, to examine how this change affects the leakage rate. Different levels of inspection effectiveness were also included in the analysis, to vary the likelihood of detecting contamination after treatment. Code is similar to that for assessing sensitivity to inspection effectiveness, with the key modification being that we examine smaller intervals of treatment effectiveness and fewer levels of inspection ef-

---

[1] Low level contamination, if detected, is assumed to be cleaned on the spot with 100% effectiveness.

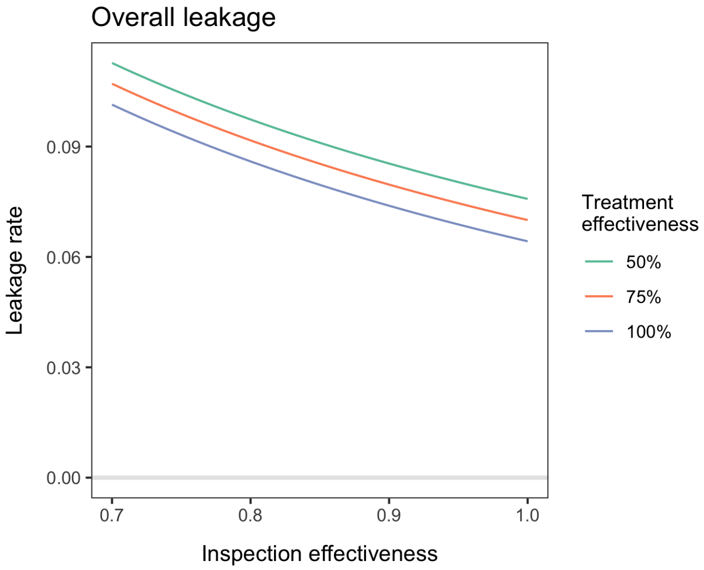

Figure 8.7.: Overall leakage rate as a function of inspection effectiveness at three levels of treatment effectiveness (50% - green line, 75% - orange line, and 100% - blue line).

fectiveness[2]. Figures 8.8 and 8.9 show that the leakage rate decreases with increasing treatment and inspection effectiveness. This effect is less pronounced for internally contaminated containers, since these are subjected to internal inspection only 12.4% of the time, whereas 76.8% of externally contaminated containers undergo external inspection.

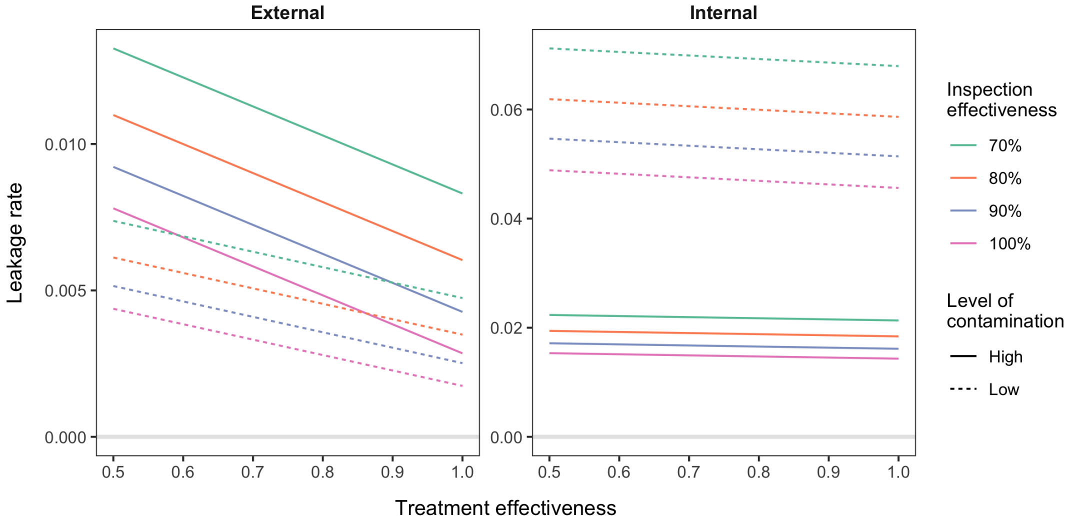

Figure 8.8.: Leakage rate as a function of treatment effectiveness at four levels of inspection effectiveness (70% - green line, 80% - orange line, 90% - blue line, and 100% - pink line). Solid lines represent high levels of external contamination and dashed lines represent low levels of contamination.

---

[2] We could, of course, complete this sensitivity analysis just once with a large number of steps along both dimensions, subsetting the resulting output appropriately for interpretation and visualisation.

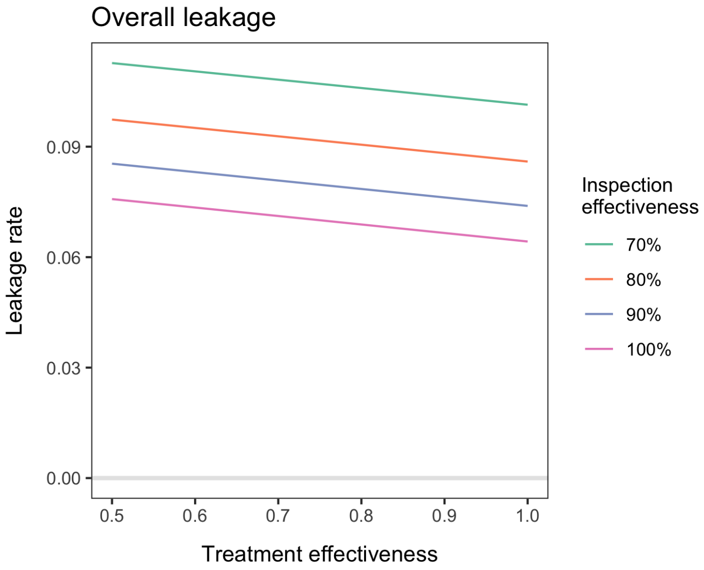

Figure 8.9.: Overall leakage rate as a function of treatment effectiveness at four levels of treatment effectiveness (70% - green line, 80% - orange line, 90% - blue line, and 100% - pink line).

8.3.3. Run time of the sensitivity analyses

The sensitivity analysis was parallelised across 16 cores. The run times were calculated as average run times across 10 runs.

- Varying inspection sensitivity in 0.01 increments from 0.7 to 1 (i.e., 31 values) took 1.4 seconds.
- Varying the inspection sensitivity with 31 steps (as above) together with three different values of treatment effectiveness (0.5, 0.75, 1) took 3.7 seconds.
- Varying the inspection sensitivity with 31 steps together with eleven different values of treatment effectiveness (from 0.5 to 1, with 0.05 increments) took 13.4 seconds.

The run times are affected minimally by the number of queried nodes.

8.4. Caveats and considerations

Key considerations relevant when interpreting this model include:

1. The sensitivity of an inspection (i.e., the probability of detecting a contaminant given that it is present) strongly influences the inferred leakage and approach rates, yet true sensitivities are unknown. For a production model, we would recommend estimating these parameters through a formal expert elicitation process. However, given time constraints and the project's nature as a proof-of-concept, we fabricate sensitivities for the purpose of demonstrating the model.
2. In the current model, we assume that the historical inspection sensitivity is the same regardless of inspection type (e.g. CAL 6-sided, tailgate and CCV inspections are assumed to have the same past sensitivity). We are aware that this is

not true, as different inspection types have different procedures and can involve inspecting different sides of a container; however, we do not have data-informed sensitives as noted above. Additionally, the current model structure does not support multiple historical inspection sensitivity values.

3. In the present model, we assume that treatment is 100% effective. This is unrealistic, and subsequent versions of the model may allow this effectiveness to be varied by adding another variable to the model to represent the probability of an inspection falsely classifying a clean good as not clean.

4. We were unable to include all containers due to their missing data, which can affect the subsequent estimated parameters.

5. In the present model, the "internal" contamination includes both internal container contamination and goods contamination, and goods-inspections (such as full-unpack) were included in the model, that is, this model includes inspection and contamination from both sea containers and containerised cargo.

9. Analysis of the cut flowers pathway

In this chapter, we demonstrate use of the model for the cut flowers pathway. In the following sections we first examine the uncertainty in outputs (Section 9.1) and then the impacts on the leakage rate of a range of counterfactual scenarios (Section 9.2) and changes in inspection and treatment effectiveness (Section 9.3).

Note that the leakage rates stated in this chapter are modelling estimates, involving various simplifications and caveats, and are underpinned by several assumptions especially where data gaps exist.

9.1. Uncertainty in outputs

The uncertainty presented in this section is the sampling uncertainty corresponding to 20,000 lines of cut flowers and foliage in any given (random) year. We chose this number because a total number of approximately 100,000 lines were imported over a period of five years. In this analysis we assumed a historical inspection effectiveness of 90% and a treatment effectiveness of 100%. Figure 9.1 shows the range of leakage of actionable contamination by group (SourceCountry and TypeOfGoods). The uncertainty around the mean leakage is based on the binomial sampling with a given probability parameter.

The overall leakage of "actionable" material on the cut flowers and foliage pathway was 2.77%, at the line level, assuming 90% line-inspection sensitivity. Foliage imported from low risk countries had the highest leakage rate. However, this cohort makes up only 0.10% of the total import volume (i.e., the total number of lines), as can be seen in Figure 9.2 which is a different representation of the data. The biggest contribution to overall leakage comes from flowers imported from medium and high risk countries (21.11% and 73.85%). These cohorts also represented most of the volume of imported lines.

Including another variable in the simulation run, initial document assessment results with outcomes of 'incomplete', 'not ok', and 'ok', provides an even more detailed view on the leakage rate by cohort (see Appendix F; Figure F.1).

9.2. Counterfactual analysis

For the cut flowers pathway, we assessed the influence on the leakage rate of three scenarios: changes in the proportion of lines from high risk countries (Section 9.2.1), changes in the proportion of lines carrying flowers from high risk countries (Section 9.2.2), and changes in the proportion of lines (flowers from high risk countries with documentation that is okay) that are inspected internally (Section 9.2.3). Counterfactual analyses were divided into leakage of actionable and not actionable contamination. The following figures are presented in side-by-side panels to contrast these two scenarios. Please

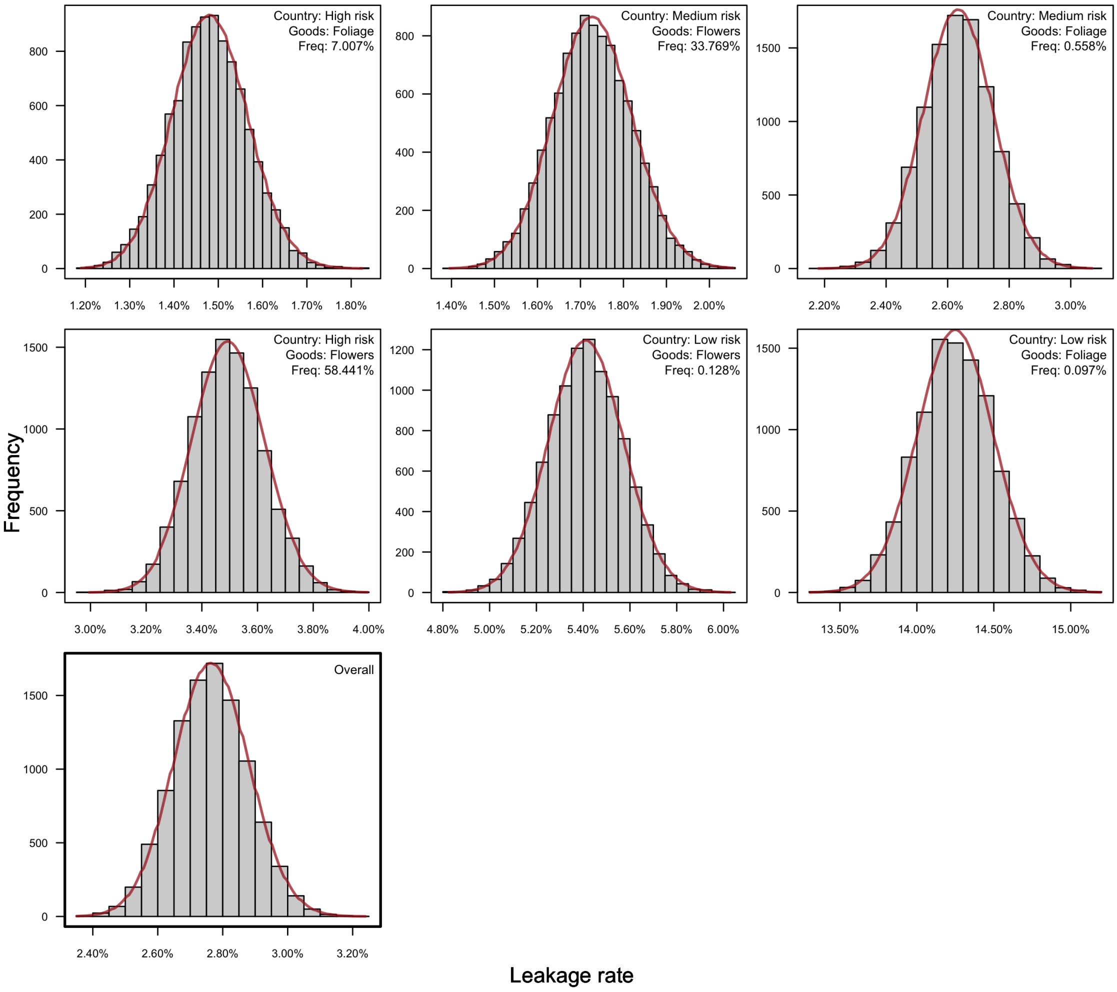

Figure 9.1.: Leakage rate distributions. Each panel shows the result, as a distribution of 1,000 model runs, using a specific combination of the states of the top level nodes. Other than the lower-left panel, which gives the unconditional leakage rate (i.e., across all combinations of top level node states), panels are presented in order of mean leakage rate.

note that the two panels (lhs and rhs) of the figures have different scales for their y-axis, consequently care should be taken in comparing effects to avoid interpretative mistakes.

9.2.1. Changes in the proportion of lines from high risk countries

This counterfactual scenario determined the consequences of changing the proportion of lines of cut flowers and foliage being imported from high risk countries. Currently, around 65% of lines originate from high risk countries. Increasing that proportion to 100% (with the concomitant decrease in lines originating from low and medium risk countries) leads to an expected increase of the leakage of actionable contamination (Figure 9.3). The decrease in leakage of not-actionable contamination is explained by the fact that while lines from medium risk countries have a lower rate of actionable contamination than those from high risk countries, they have a higher rate of not-actionable contamination; hence in this counterfactual we are in essence moving lines

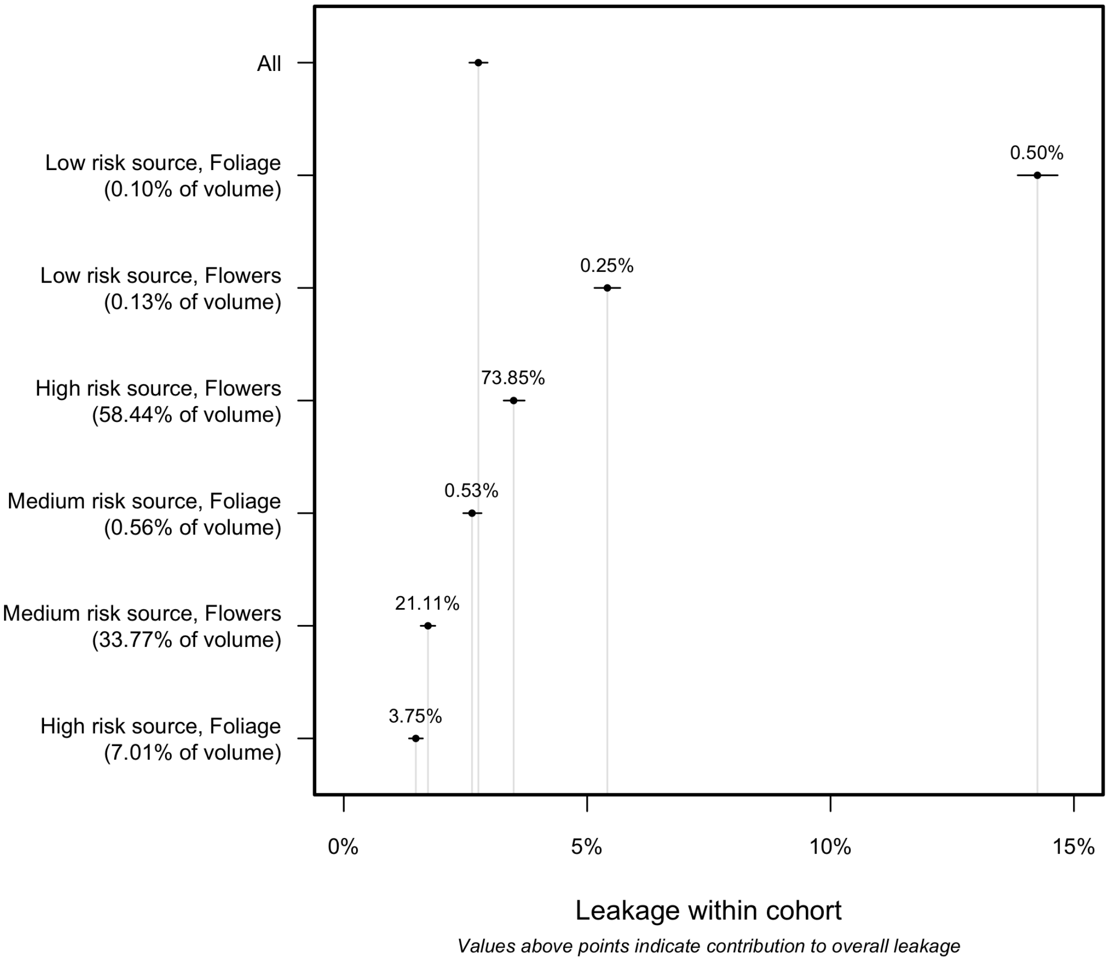

Figure 9.2.: Leakage within cohorts of group (SourceCountry and TypeOfGoods). The y-axis shows the different cohorts and their respective proportion of the total import volume, while the x-axis presents the leakage rate. Dots represent the mean leakage rate and bars the 95% confidence interval. The values above the dots indicate the contribution (in percent) of the cohort to the overall leakage rate.

from a cohort with a relatively high not-actionable contamination rate and low actionable contamination rate, to a cohort with relatively low not-actionable contamination rate and high actionable contamination rate (Figure 9.4).

9.2.2. Changes in the proportion of lines carrying flowers from high risk countries

In this scenario, we changed the proportion of lines carrying flowers that are being imported from high risk countries. Currently, around 89.3% of lines on the cut flowers pathway import flowers rather than foliage. When we increase the number of lines carrying flowers and thereby decrease the number of lines carrying foliage we can observe an increase in the leakage of actionable and not actionable contamination (Figure 9.5). This is due to the higher leakage rate that we expect from the import of flowers from high risk countries compared to the import of foliage, as can be seen in Figure 9.2.

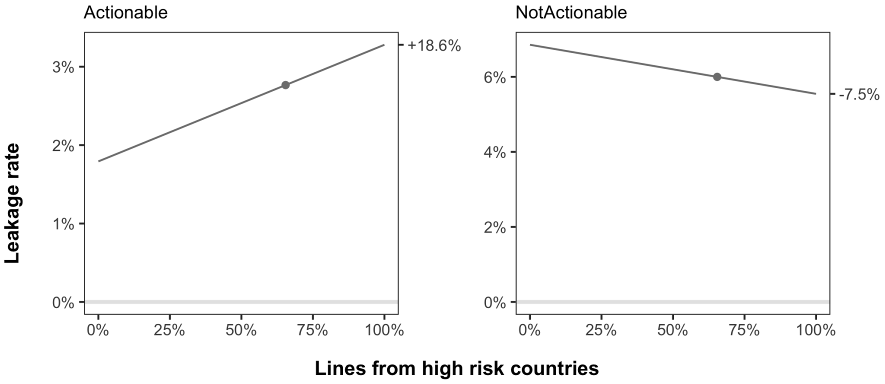

Figure 9.3.: Changing the proportion of lines of cut flowers and foliage imported from high risk countries. Dots on lines indicate the leakage rates corresponding to the status quo of the proportion of lines coming from high risk countries (i.e., 65.4%). The left panel presents the leakage of actionable contamination and the right panel the leakage of not-actionable contamination. The percent values at the right margin of each panel indicate the change in leakage rate relative to the status quo.

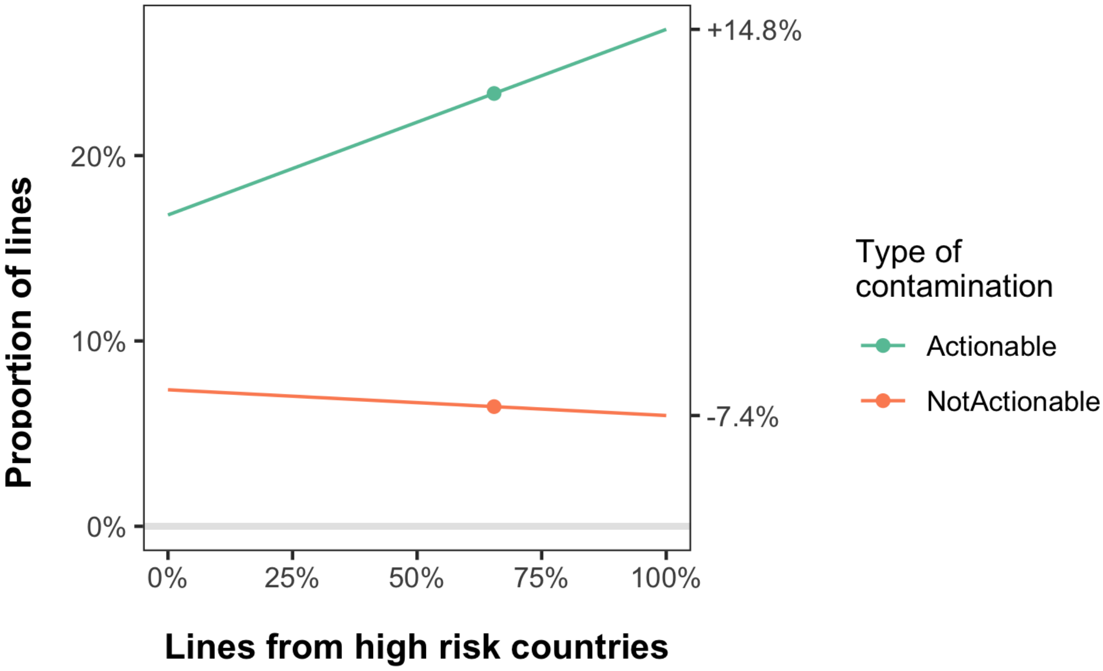

Figure 9.4.: Changes in the relative proportions of actionable and non-actionable contamination as the proportion of lines of cut flowers and foliage imported from high risk countries is varied. Dots on lines indicate the proportions corresponding to the status quo (i.e., 65.4% of lines come from high risk countries). The percent values at the right margin of each panel indicate changes relative to the status quo.

9.2.3. Changes in the proportion of lines that are inspected internally

This counterfactual scenario investigated the influence of internal inspections on leakage rate. Given that border biosecurity staff inspect a very high proportion of lines

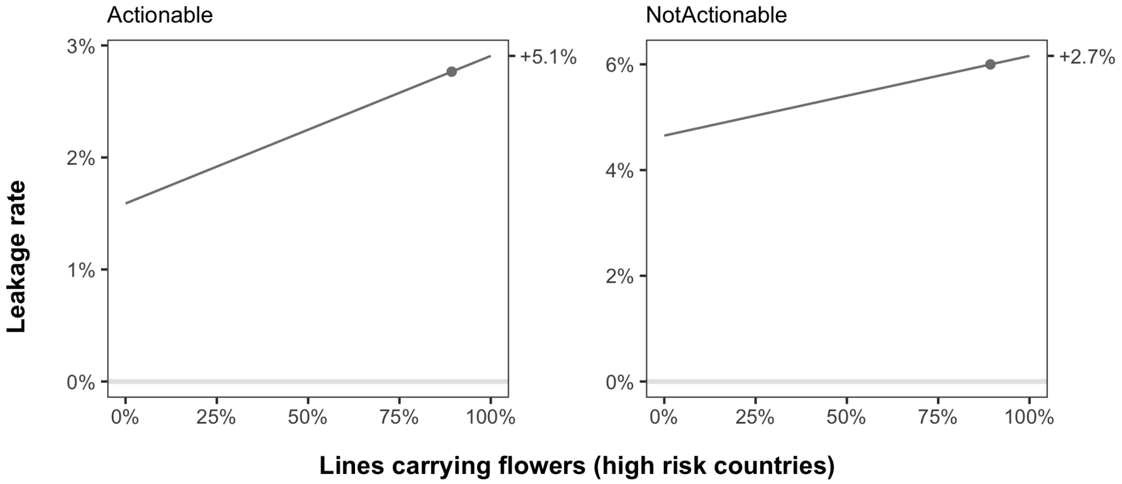

Figure 9.5.: Changing the proportion of lines carrying only cut flowers that are imported from high risk countries. Dots on lines indicate the leakage rates corresponding to the status quo of the proportion of cut flower lines coming from high risk countries (i.e., 89.3%). The left panel presents the leakage of actionable contamination and the right panel the leakage of not actionable contamination. The percent values at the right margin of each panel indicate the change in leakage rate relative to the status quo.

each year (97.6%) we can look at what would happen if inspection rates were reduced. In the example here, we reduce inspection rates for a particular cohort - lines of flowers from high risk countries with compliant import documentation. As expected, the leakage rate increased steeply with reduced inspection effort for actionable contamination (Figure 9.6). If no lines of cut flowers were inspected internally, then the leakage of actionable contamination would be more than 14%. The slight positive slope for not-actionable contamination is counterintuitive; it arises due to treatment rates being higher, in some cases, for compliant lines (i.e., lines found to be clean at inspection or lines presumed to be clean if they have not been inspected) than for inspected lines found to have not-actionable contamination (see Section 9.4).

9.3. Sensitivity analysis

In this section, we investigate how changes in inspection effectiveness (Section 9.3.1) and in treatment effectiveness (Section 9.3.2) affect the leakage rate.

9.3.1. Changes in inspection effectiveness

Here, we examine the impacts of modifying our assumption about the historical inspection effectiveness inherent in the cut flowers dataset. To test the sensitivity of the model outputs (i.e., the leakage of actionable contamination in this context) to changes in inspection effectiveness, we increased the inspection effectiveness to 100% (i.e., all contaminated lines that were inspected were labelled as being contaminated). Increasing inspection effectiveness leads to a decrease in the leakage rate (Figure 9.7) because of (1) the associated reduction in inferred contamination approach rates (i.e., if we assume inspections underlying the data are 100% effective, then the detection rate ap-

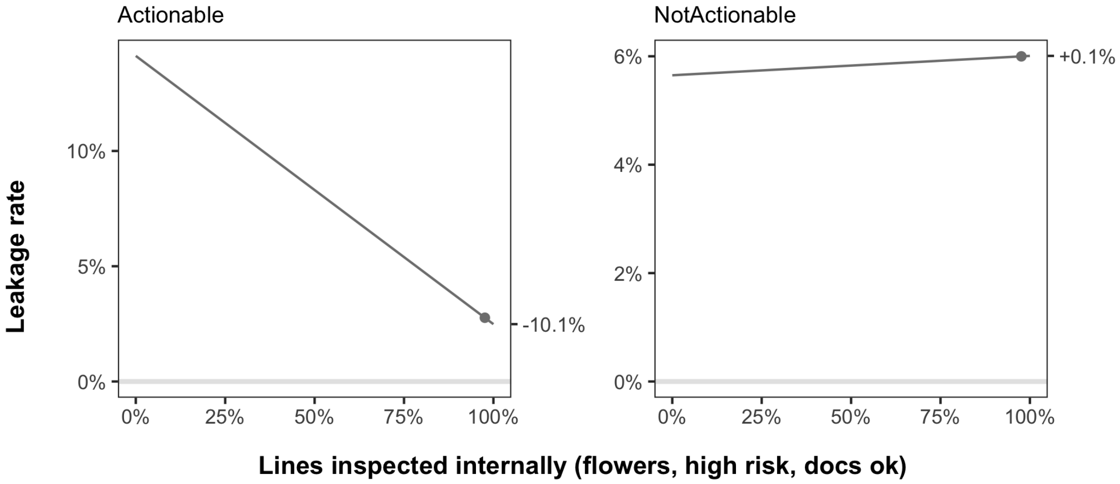

Figure 9.6.: Changing the internal inspection effort for imported flowers from high risk countries with compliant documentation. Dots on lines indicate the leakage rates corresponding to the status quo of the proportion of lines from this cohort that are being inspected internally (i.e., 97.6%). The left panel presents the leakage of actionable contamination and the right panel the leakage of not actionable contamination. The percent values at the right margin of each panel indicate the change in leakage rate relative to the status quo.

proximates the true rate of contamination; however, if we assume 90% inspection effectiveness, then the detection rate represents only 90% of the true contamination rate), and (2) the causal relationship between inspection effectiveness and leakage which means that a greater proportion of actionable contamination present on inspected lines is detected. Leakage of not-actionable contamination was not particularly sensitive to inspection effectiveness, because only a small fraction (7.1%) of those lines with not-actionable contamination are treated, representing just 0.46% of all lines. This is in contrast to the lines with actionable contamination, 88.2% of which undergo treatment (20.6% of all lines). As expected, the leakage rate decreased with increasing effectiveness of treatments (0%, 50%, and 100%).

9.3.2. Changes in treatment effectiveness

Figure 9.8 shows the impacts on the rate of leakage of actionable contamination when the effectiveness of treatments is decreased from 100% to 50%. A reduction in treatment effectiveness to 50% results in a steep increase in the leakage rate because only half of the detected actionable contamination is removed effectively. As discussed in Section 9.3.1, the leakage rate decreases with increasing inspection effectiveness.

9.4. Caveats and considerations

Key considerations relevant when interpreting this model include:

1. Inspection of a line or consignment does not usually mean inspection of *all* flowers or foliage in that line or consignment. While we would expect that the contamination of stems within one consignment to be interlinked, it is highly possible that uninspected stems are actually contaminated when the inspected stems

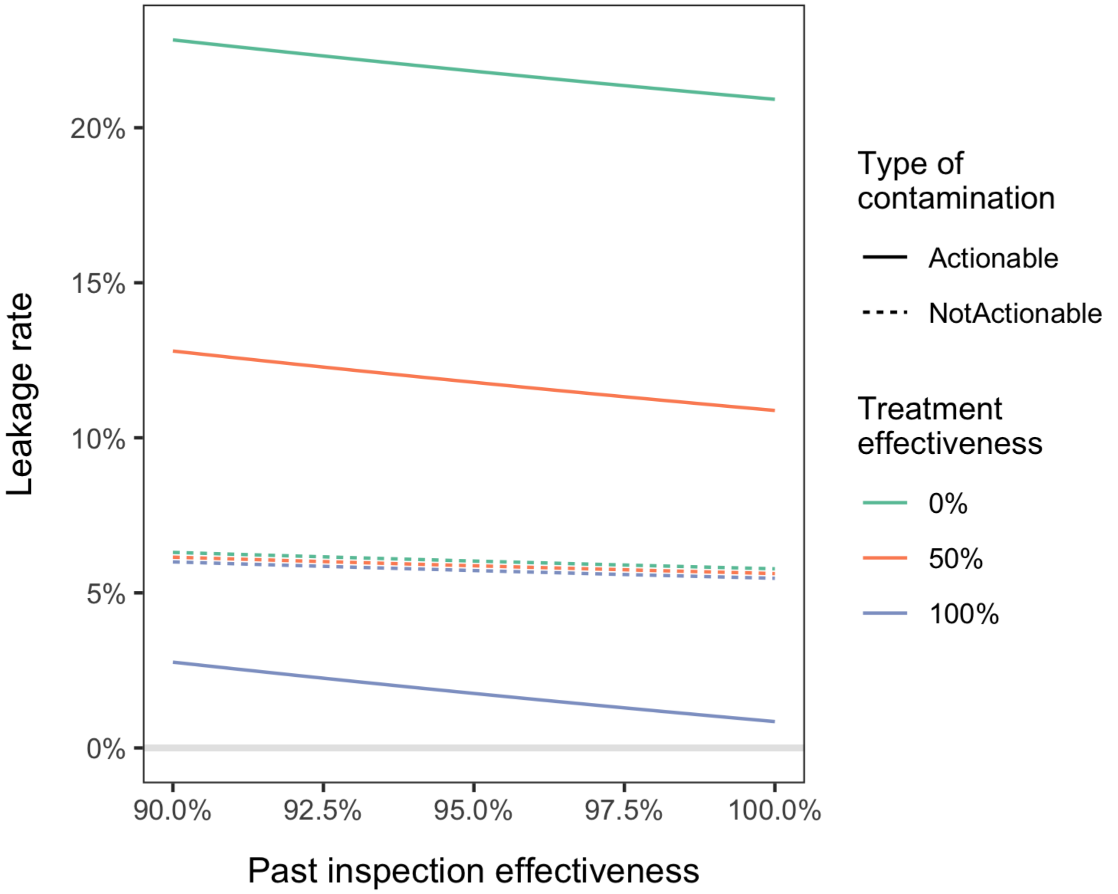

Figure 9.7.: Overall leakage rate of actionable and not-actionable contamination as a function of inspection effectiveness at three levels of treatment effectiveness (0% - green line, 50% - orange line, and 100% - blue line). Actionable contamination is represented by solid lines and not actionable contamination by dashed lines.

are clean. However, we did not include a node into our BN-like network describing the proportion of stems inspected, which would then modify the leakage rate. This is possible, but it would require data which may not exist.

2. Treatment rates, required for each combination of internal inspection outcome class and external inspection outcome class, were calculated as empirical frequencies from the data where possible. However, for particular combinations no instances existed and the corresponding rates were imputed by assigning medians of relevant groups (e.g., median treatment rate for lines with actionable contamination was assigned to unrepresented combinations that involved actionable contamination). These imputed rates may not accurately reflect actual treatment rates should such combinations be encountered.

3. Note that we did not explicitly link lines within the same consignment together. Because of this, "clean" lines, or lines without inspection, can still be treated within our model (e.g., because other lines in the consignment were found to be contaminated). Additionally, this could affect counterfactual results as we modify treatment efforts. Linking the actions upon lines within the same consignment

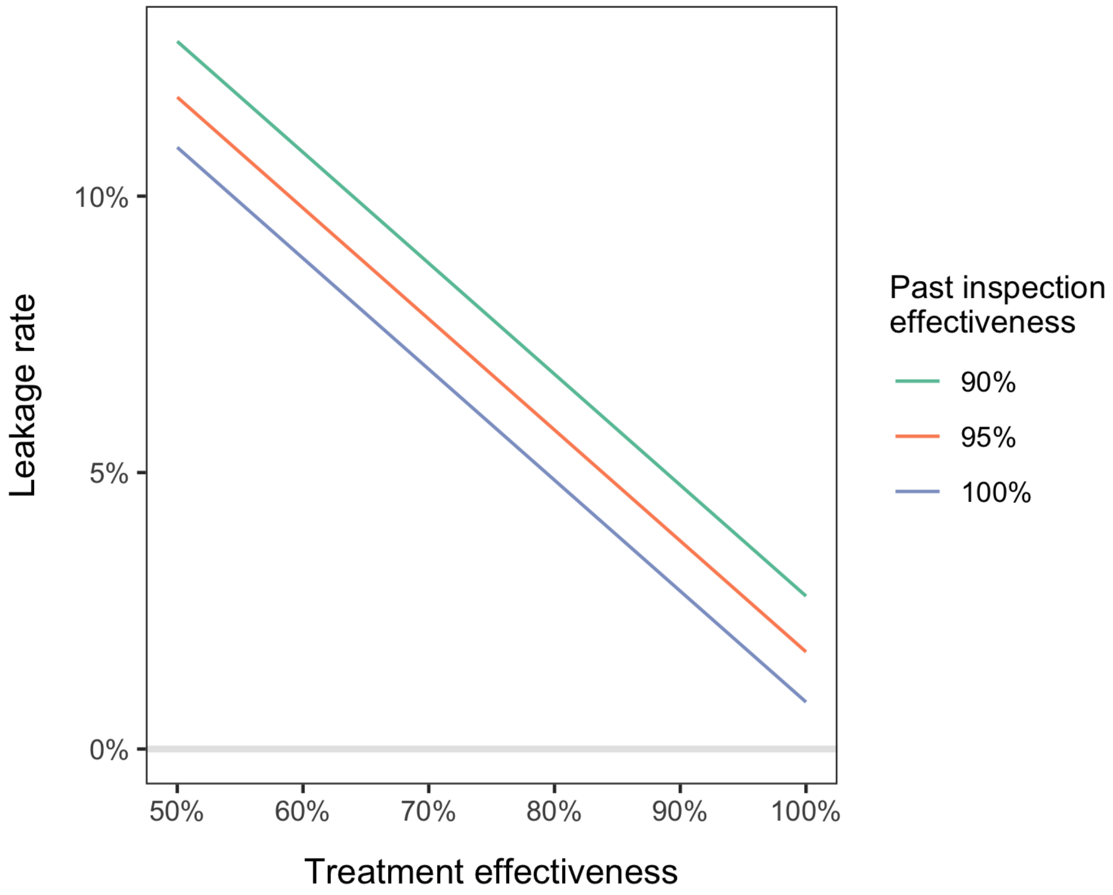

Figure 9.8.: Overall leakage rate of actionable contamination as a function of treatment effectiveness at three levels of past inspection effectiveness (90% - green line, 95% - orange line, and 100% - blue line).

would most likely involving the additional nodes to the BN-like network.

4. In some cases, treatment frequencies were higher for lines that were found to be clean at inspection and/or lines that were presumed clean without inspection than for inspected lines with not-actionable contamination. This will in part be due to small sample sizes for some groups. Given our assumption that uninspected lines are deemed to be clean and compliant, the above can result in the counterintuitive situation where reduced inspection leads to a higher rate of treatment, and hence a lower rate of leakage. While the general structure of the network is logical, the estimated treatment rates could benefit from refinement based on additional data or insights from Departmental personnel.

5. The sensitivity of an inspection strongly influences the inferred leakage. We do not have a formal estimate of the sensitivity of inspection.

6. It is possible to lower the sensitivity of inspection by some factor to account for the fact that not all stems are inspected, by formulating some kind of average inspection sensitivity given data about the proportion of stems inspected, and information or assumptions about the distribution of contamination in lines and consignments.

7. We have mentioned two methods that could be used to adjust for the fact that a maximum of 600 stems are inspected in any given consignment: adding another node in the BN-like network, or lowering inspection sensitivity. Note that we do not have any preference, and we expect both methods to provide similar results if calibrated correctly. Regardless, both methods would require data.

8. Our leakage estimate should be an underestimate, because we did not incorporate information about how not all stems were inspected in a line. Given how the maximum number of stems examined is based on finding lines with over 0.5% contamination, our leakage rate may be better posed as the leakage rate for 'highly contaminated' lines, rather than for any contamination at all. However, quantifying what is meant by 'highly contaminated' would require further work.

9. Many lines and consignments had missing information (i.e. missing at least one of documentation, verification, inspection, or final directives), which we complemented with default values. It is possible that these lines should have been excluded from the analysis. Alternatively, if different default values were chosen the subsequent results may have been different.

10. We also note that some directions were incorrectly filed and that often (but not always), crucial clarifying information was located in the free text fields. While we did use the free text field as a last resort to see if it contained any missing information, we did not use it to its fullest potential, which may have required natural language processing.

11. Note that the data we used included dried or preserved flowers, which have different biosecurity risk from fresh cut flowers and foliage. As such, leakage results should be considered as being for both preserved and fresh flowers and foliage.

12. Note that contamination results might not be attributable to the line it was recorded against, but could be due to the results from another line in the same entry. As different lines in the same entry can have different attributes (from a different source country, containing cut flowers and foliage), this will give reduced reliability to our estimated contamination rates given different sets of attributes.

10. Resource Allocation Tool

10.1. Overview of the Tool

The resource allocation tool ("the Tool") is delivered as a Shiny web app (Chang *et al.*, 2022) integrating the two pathway models (containerised goods, and cut flowers and foliage) with CEBRA's "Value Model", a spatially explicit bio-economic model developed to estimate the net present value of Australia's biosecurity system by simulating the arrival, spread and impact of biosecurity hazards (Dodd *et al.*, 2020). By linking the pathway models with the Value Model, rates of risk material leakage along the pathways are captured in arrival rates driving simulation behaviour, thereby revealing the economic impact – in terms of differential damage to assets – of pathway-specific control measures (Figure 10.1).

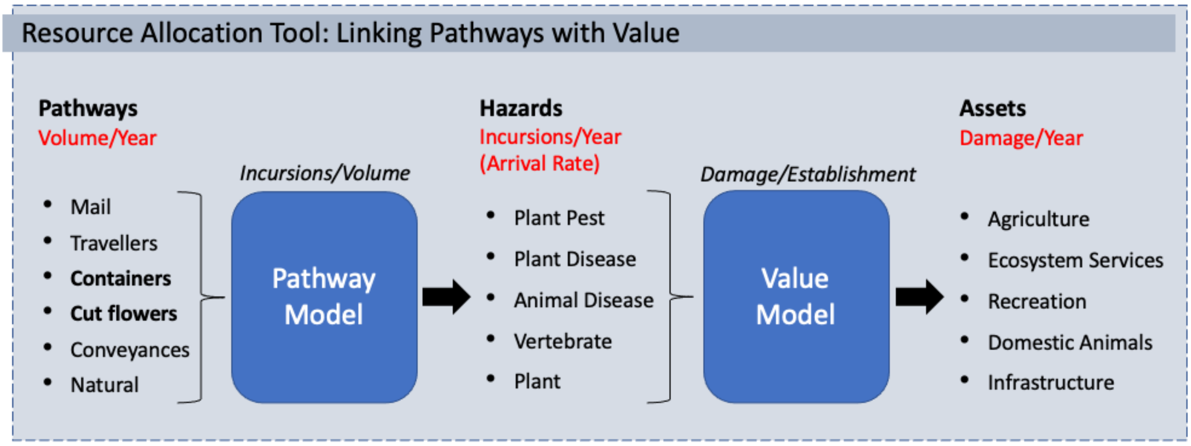

Figure 10.1.: Relationship between the pathway models and the Value Model (Figure adapted from Dodd *et al.*, 2022).

The Tool's key purpose is to allow users to assess the relative benefits of competing strategies of resource allocation into biosecurity measures. To this end, the user interface permits specification of alternative levels of resource allocation into controls on one or both focal pathways. As a proof-of-concept, the Tool limits this functionality to varying the effectiveness of individual inspection methods, i.e., the probability of detecting contamination when present on an inspected item (a container, or line of cut flowers/foliage). For each alternative level of effectiveness, the following process occurs:

1. The predefined pathway models are updated and the probability that an item is released with contamination is recalculated;
2. The probabilities from (1) are multiplied by the user-defined total annual volume

of items on each pathway to determine the overall leakage of biosecurity risk material (i.e., the count of contaminated items leaked);

3. Using sets of probabilities provided by the Department, leakage of risk material from (2) is then separated into the 40 functional groups considered by the Value Model, and these functional group-specific leakage rates are translated to rates of arrival and successful establishment attributable to the focal pathways;
4. These pathway-specific establishment rates are added to the background rate of establishment through other pathways, and the resulting rates are used to parameterise the Value Model;
5. The Value Model simulates establishment and spread consistent with the overall rate of establishment, and calculates resulting damage.

Damage estimates are then compared across alternative control scenarios and displayed by the Tool in tabular and graphical form.

10.2. User interface

The Tool's user interface comprises a web app with a series of tabs for adjusting simulation settings, running analyses, and viewing results. At any time, users can click the "Save configuration" button to save all settings; doing so produces a specific URL that when visited, will restore the saved settings.

10.2.1. Global settings

The *Settings* tab (Figure 10.2) provides inputs allowing users to define the spatial extent (Demo = Greater Melbourne; Full = all of Australia), duration, and number of iterations, as well as annual container and flower/foliage volume and the effectiveness of inspection on each pathway during the period represented by data used to parametrise the pathway models. Note that container volume here refers to the number of containers referred through ICS profiling to AIMS. Non-referred containers are not considered by the Tool; instead, their contribution to overall hazard establishment is captured in the constant background establishment attributable to all other pathways[1]. This tab also allows users to specify the alternative scenarios of inspection effectiveness, as proportions ranging from 0 to 1, that should be compared during the analyses – this can optionally be specified per inspection method.

10.2.2. Analyses—Compare scenarios

The "Compare scenarios" analysis tab (Figure 10.3) provides additional settings specific to comparing control strategies across one or more controls (i.e., effectiveness scenarios across one or more inspection methods), involving one or both pathways. After selecting the species of interest, the chosen pathway(s) and the controls to be modified, the Run Analysis button can be used to start the analysis.

---

[1] A consequence is that assessment of resource allocation into CCV does not account for the impacts of CCV on leakage through non-referred containers, despite a small fraction being inspected by this means.

Upon commencing an analysis, the full set of alternative scenarios is computed using a full factorial design. The figures in this section show the outcomes of varying a single control, but in the case where multiple controls are examined simultaneously, all combinations of levels of controls across chosen inspection methods are considered. For example, if a user defines the alternative levels of control for all inspection methods as `0, 0.5, 1` (on the Settings tab), and selects “CAL 6-sided (External)” and “Rural Tailgate (non-CAL, External)” as the controls to examine (on the Compare scenarios tab), then the nine possible combinations (i.e., `0 0`, `0 0.5`, ..., `1 1`) will be considered. All other inspection methods will be kept constant at their 2018–2022 *status quo* values as defined on the Settings tab.

While running, progress is indicated at the bottom of the screen, and upon completion, the “Key Results“ and ‘Graphical Summary“ tabs will be populated with comparison results. These include (1) a tabular summary (Figure 10.4) showing the mean ± SEM (Standard Error of the Mean) and median damage across iterations for each scenario, as well as the benefit (reduction in damage) of each scenario relative to the

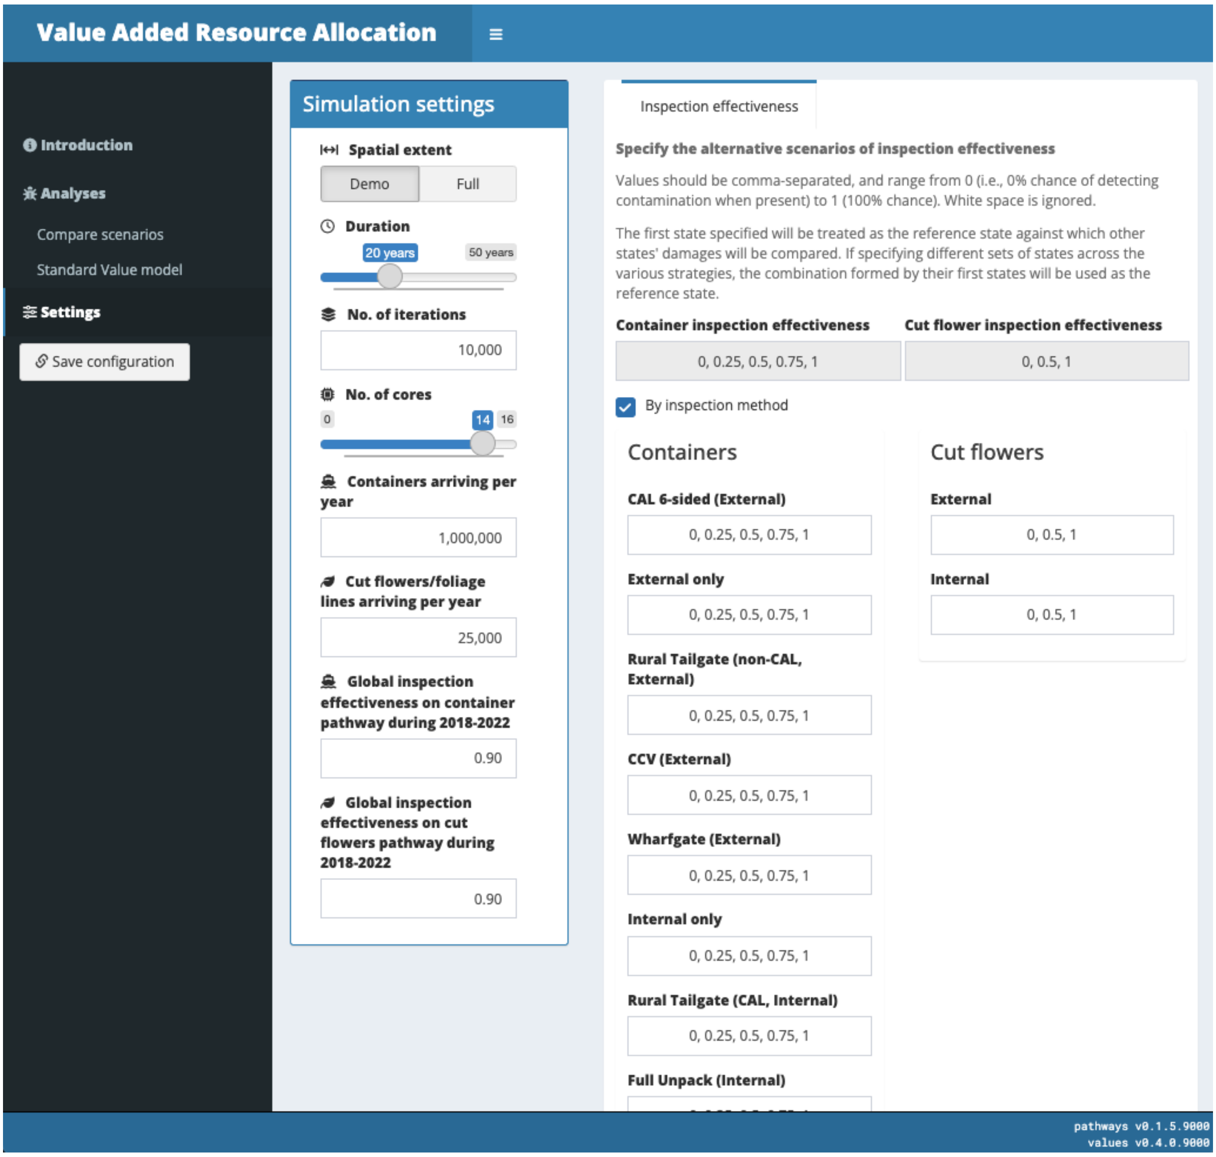

Simulation settings

Spatial extent: Demo

Duration: 20 years (slider to 50 years)

No. of iterations: 10,000

No. of cores: 14

Containers arriving per year: 1,000,000

Cut flowers/foliage lines arriving per year: 25,000

Global inspection effectiveness on container pathway during 2018-2022: 0.90

Global inspection effectiveness on cut flowers pathway during 2018-2022: 0.90

Inspection effectiveness

Specify the alternative scenarios of inspection effectiveness

Values should be comma-separated, and range from 0 (i.e., 0% chance of detecting contamination when present) to 1 (100% chance). White space is ignored.

The first state specified will be treated as the reference state against which other states' damages will be compared. If specifying different sets of states across the various strategies, the combination formed by their first states will be used as the reference state.

Container inspection effectiveness: 0, 0.25, 0.5, 0.75, 1

Cut flower inspection effectiveness: 0, 0.5, 1

☑ By inspection method

Containers

CAL 6-sided (External): 0, 0.25, 0.5, 0.75, 1

External only: 0, 0.25, 0.5, 0.75, 1

Rural Tailgate (non-CAL, External): 0, 0.25, 0.5, 0.75, 1

CCV (External): 0, 0.25, 0.5, 0.75, 1

Wharfgate (External): 0, 0.25, 0.5, 0.75, 1

Internal only: 0, 0.25, 0.5, 0.75, 1

Rural Tailgate (CAL, Internal): 0, 0.25, 0.5, 0.75, 1

Full Unpack (Internal)

Cut flowers

External: 0, 0.5, 1

Internal: 0, 0.5, 1

Figure 10.2.: The *Settings* tab of the Resource Allocation Tool’s user interface.

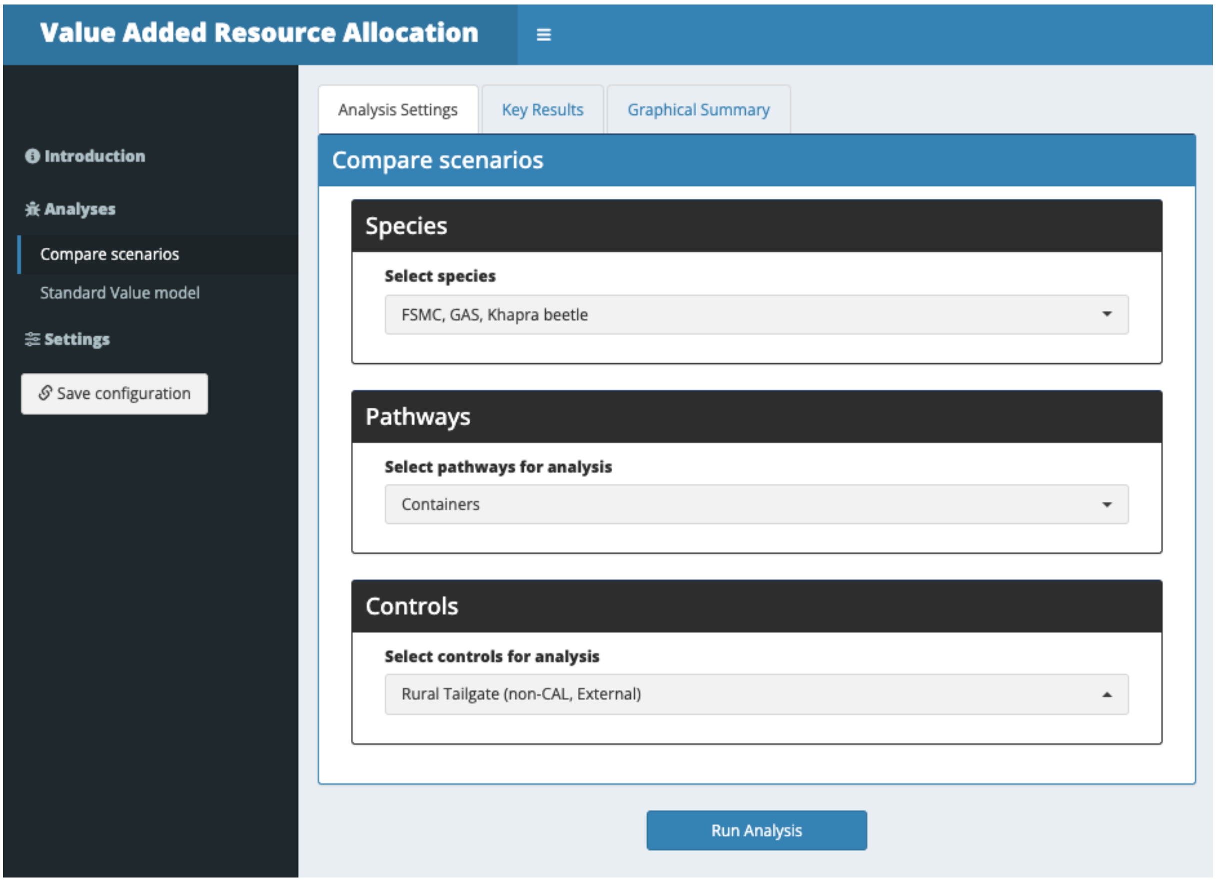

Select species
FSMC, GAS, Khapra beetle

Select pathways for analysis
Containers

Select controls for analysis
Rural Tailgate (non-CAL, External)

Figure 10.3.: The *Compare scenarios* tab of the Resource Allocation Tool's user interface.

reference scenario (defined as the combination of controls produced by taking the first-listed effectiveness value for each considered control); and (2) a graphical summary showing the average damage over time (both per asset and summing across assets), the distribution of cumulative damage at the end of the simulation, and the benefit relative to the reference state, for each control scenario (Figures 10.5 and 10.6). Note that figures are interactive, allowing zooming on areas of interest and showing values on hover.

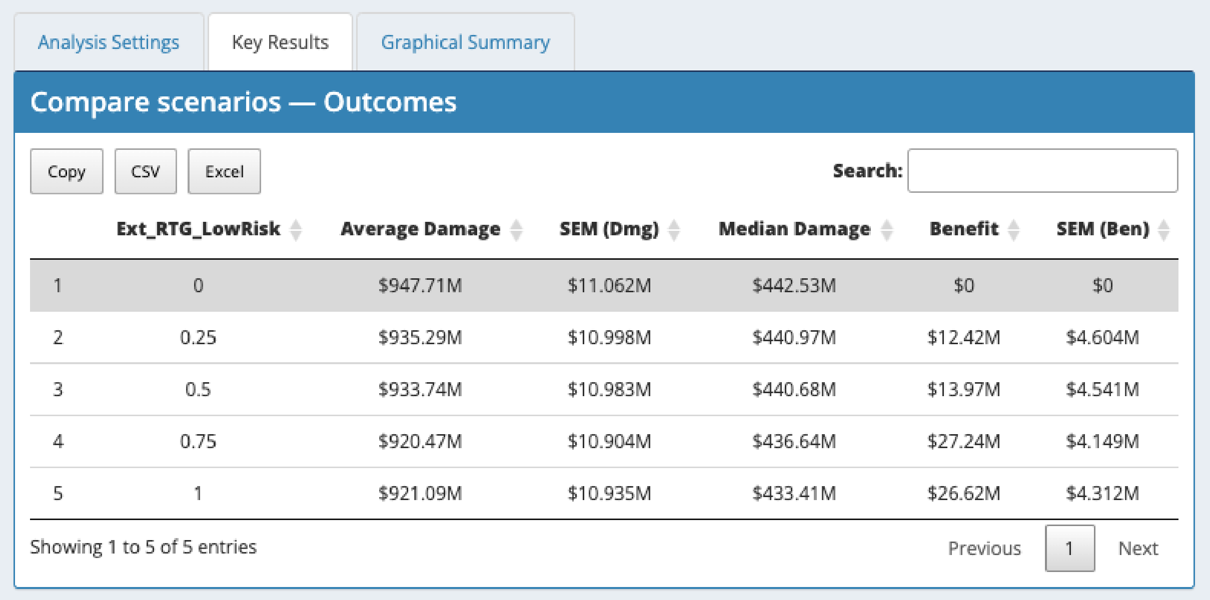

Compare scenarios — Outcomes

	Ext_RTG_LowRisk	Average Damage	SEM (Dmg)	Median Damage	Benefit	SEM (Ben)
1	0	$947.71M	$11.062M	$442.53M	$0	$0
2	0.25	$935.29M	$10.998M	$440.97M	$12.42M	$4.604M
3	0.5	$933.74M	$10.983M	$440.68M	$13.97M	$4.541M
4	0.75	$920.47M	$10.904M	$436.64M	$27.24M	$4.149M
5	1	$921.09M	$10.935M	$433.41M	$26.62M	$4.312M

Showing 1 to 5 of 5 entries

Figure 10.4.: Tabular results summary for the *Compare scenarios* analysis of the Resource Allocation Tool. Buttons permit exporting the table to CSV or Excel format, or to the clipboard. Note that reported values are intended for demonstration only.

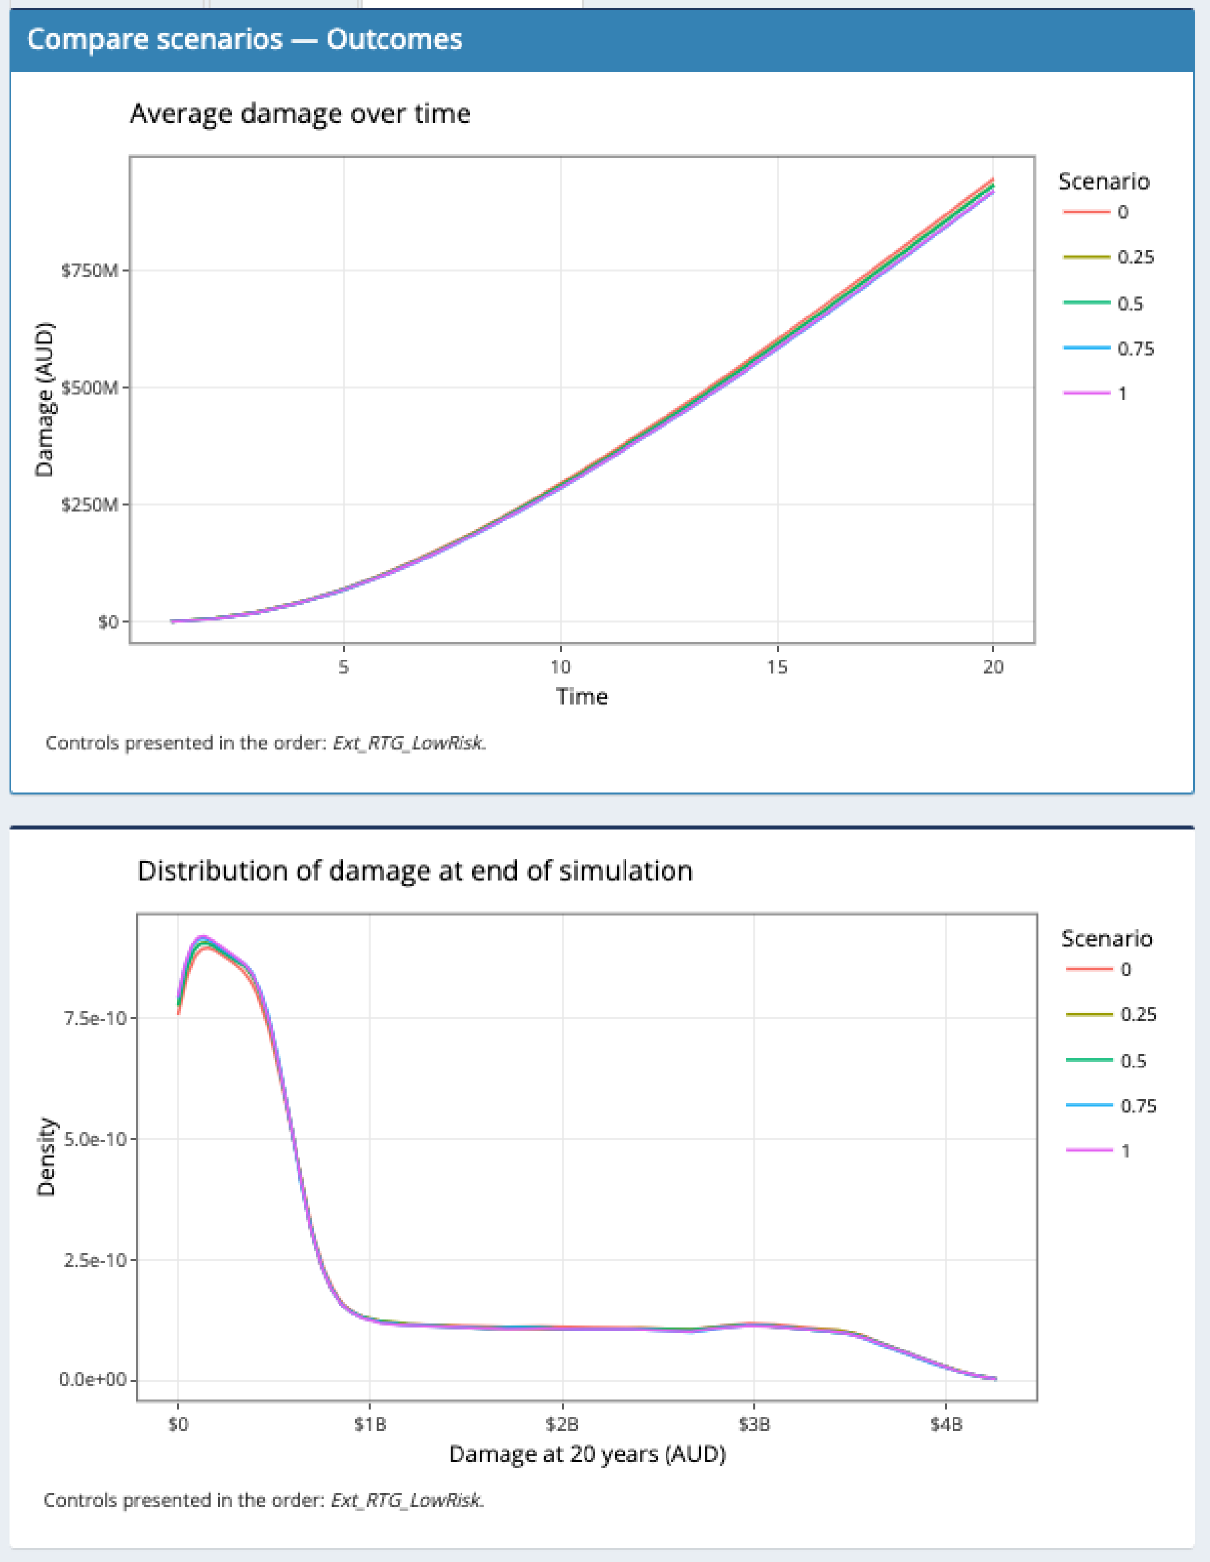

Controls presented in the order: *Ext_RTG_LowRisk*.

Controls presented in the order: *Ext_RTG_LowRisk*.

Figure 10.5.: Graphical summary for the *Compare scenarios* analysis of the Resource Allocation Tool. Figures show the average cumulative damage over time, corresponding to each of the control scenarios (top); and the distribution of damages across iterations at the end of the simulation (year 20; bottom). Note that reported values are intended for demonstration only.

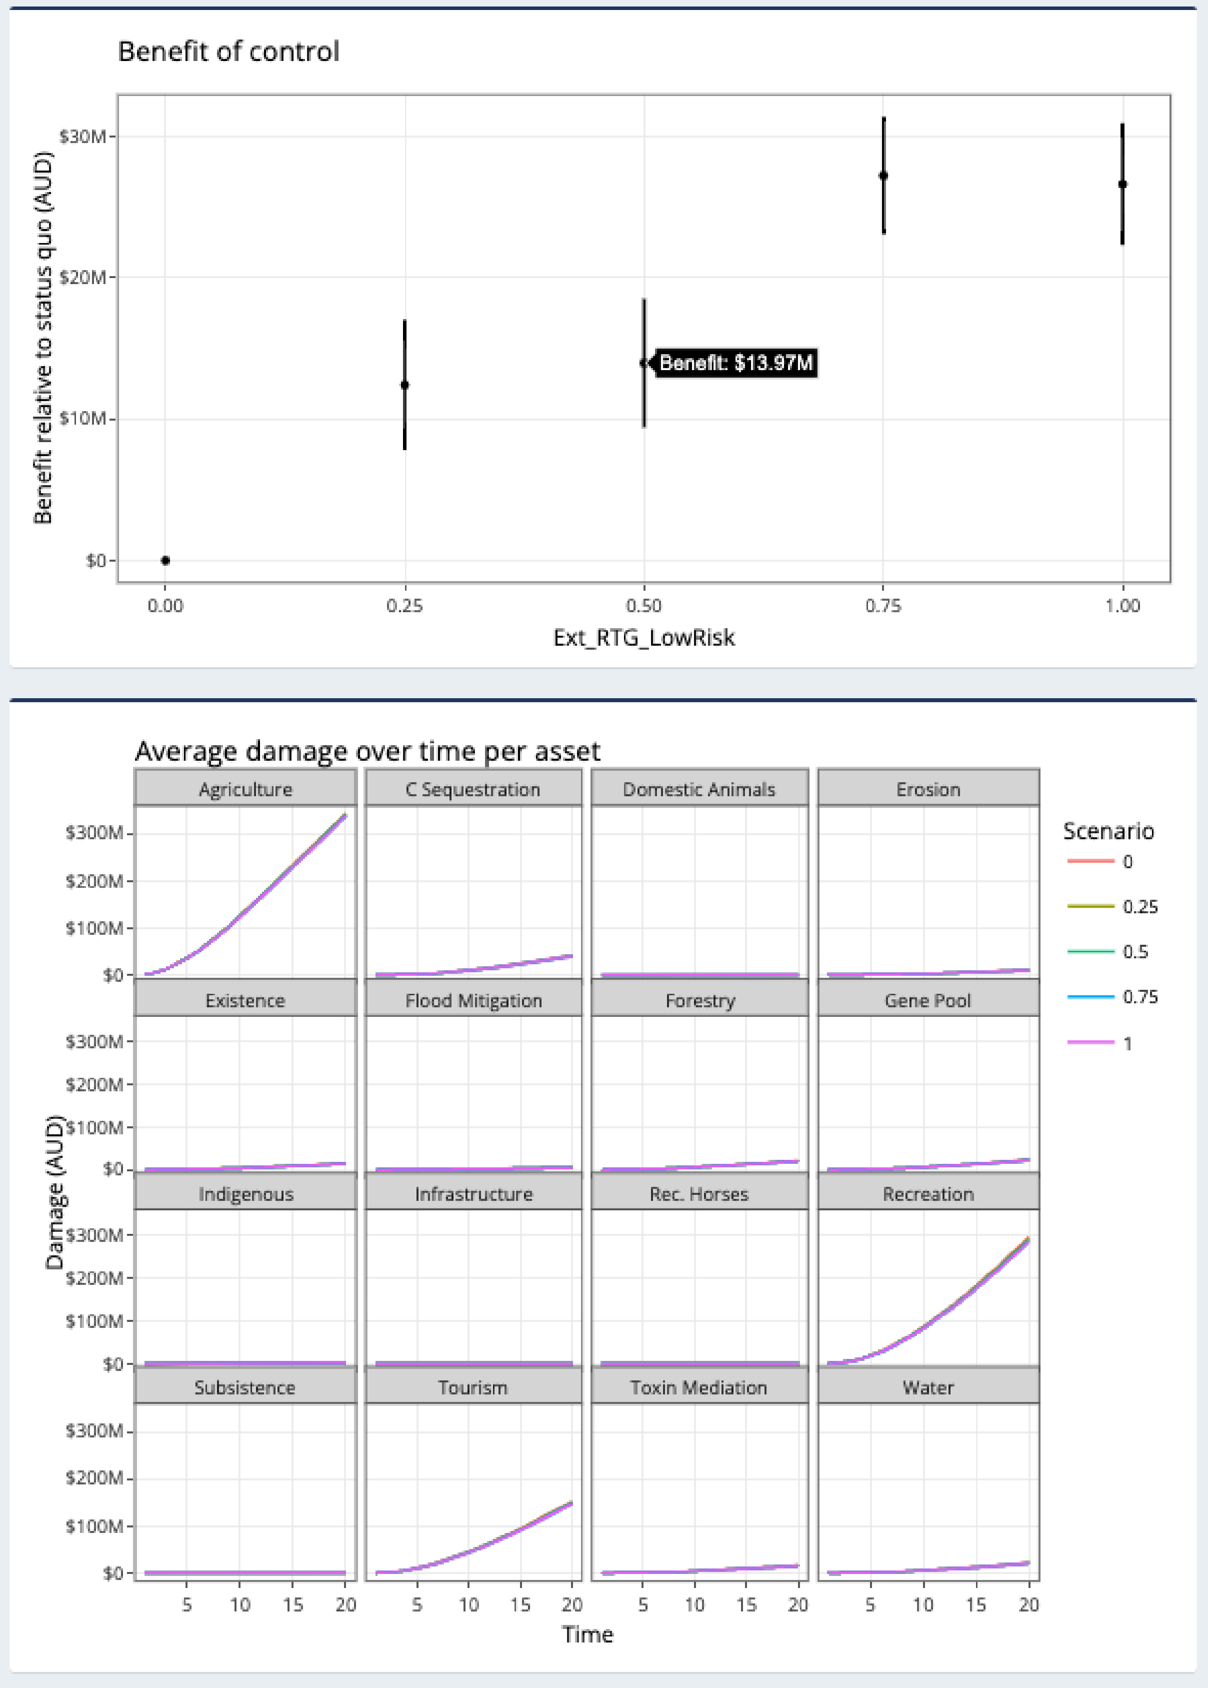

Figure 10.6.: Graphical summary for the *Compare scenarios* analysis of the Resource Allocation Tool. Figures show the mean ± SEM benefit of control scenarios relative to the reference scenario (`0 0` in this case; top), and the average cumulative damage over time per asset, attributable to the selected species, for each of the control scenarios (bottom). Note that reported values are intended for demonstration only.

10.2.3. Analyses—Standard Value Model

The Tool also serves as an interface to the original Value Model comparison of biosecurity system *on* versus *off* as examined in Project 170713 (Dodd *et al.*, 2020). All 40 species, or a subset thereof, can be selected, and the trajectory database sampling approach (Section H.1) is used for efficient simulation. A tabular summary shows the benefit of the system versus having no system (Figure 10.7), and graphical results show damage over time for each scenario, as well as the distribution of damage at the end of the simulation (Figure 10.8).

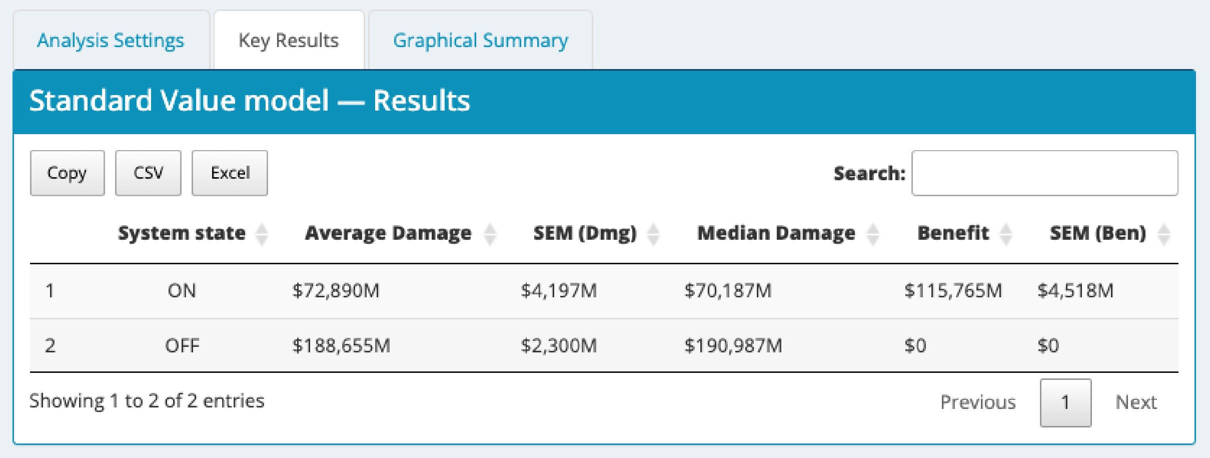

Standard Value model — Results

	System state	Average Damage	SEM (Dmg)	Median Damage	Benefit	SEM (Ben)
1	ON	$72,890M	$4,197M	$70,187M	$115,765M	$4,518M
2	OFF	$188,655M	$2,300M	$190,987M	$0	$0

Showing 1 to 2 of 2 entries

Figure 10.7.: Tabular results summary for the *Standard Value Model* analysis of the Resource Allocation Tool. Buttons permit exporting the table to CSV or Excel format, or to the clipboard. Note that reported values correspond to a subset of species and are intended for demonstration only.

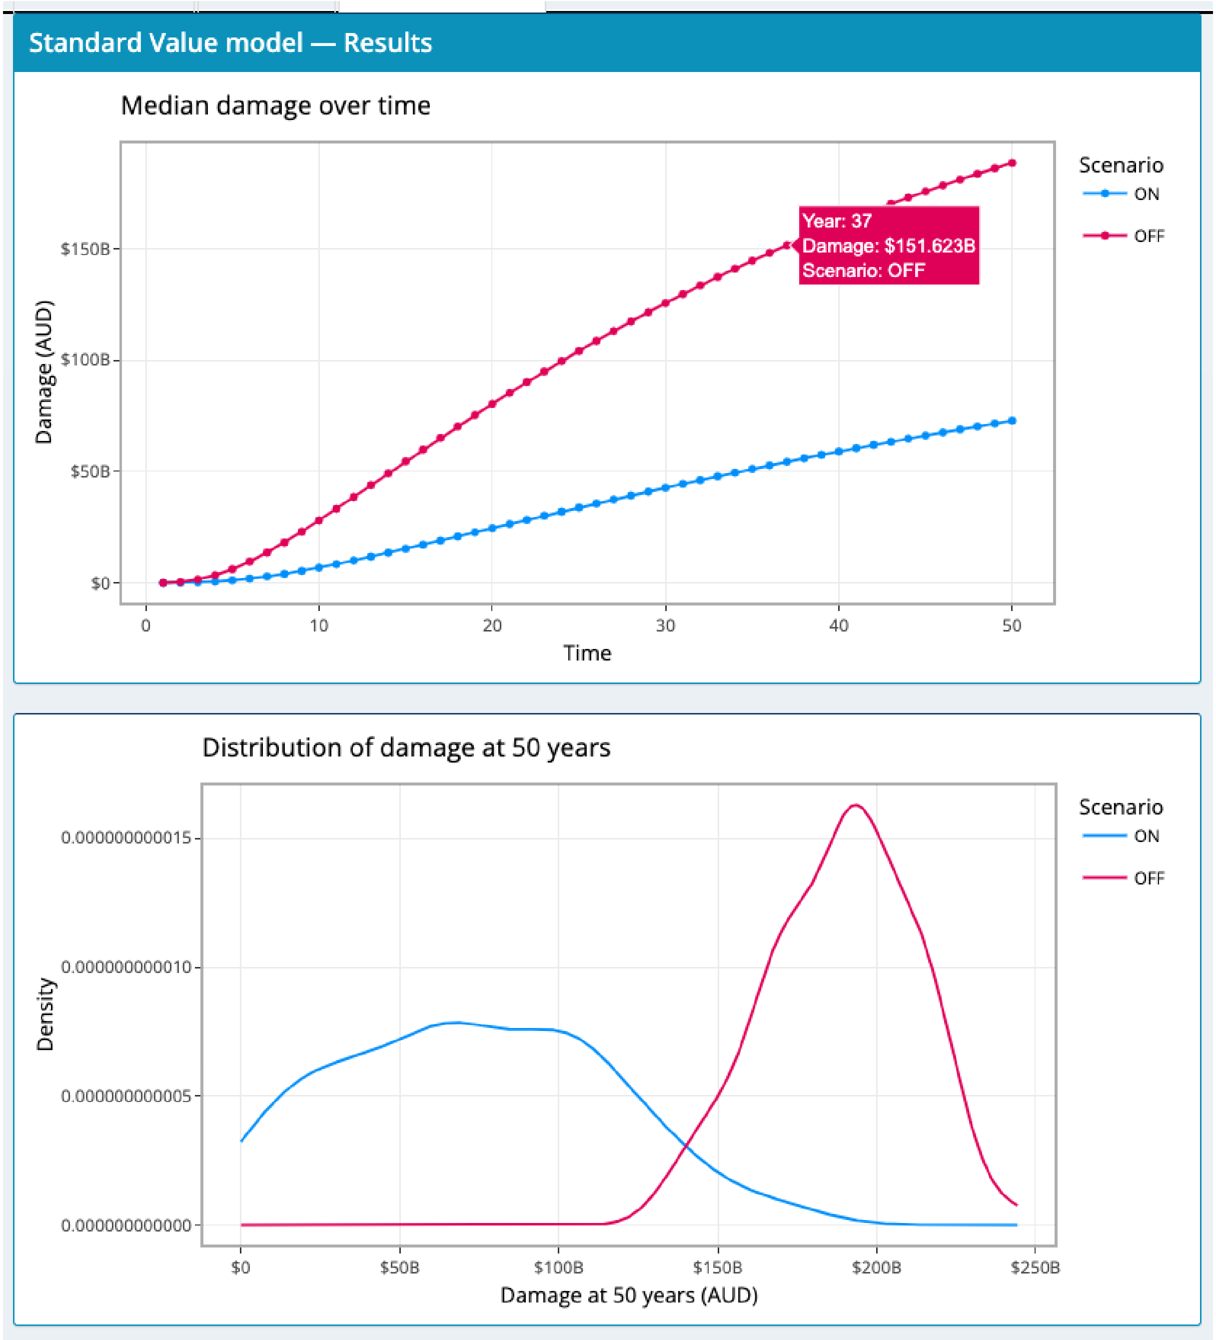

Figure 10.8.: Graphical summary for the *Standard Value Model* analysis of the Resource Allocation Tool. Figures show the median cumulative damage over time, corresponding to each of the control scenarios (top); and the distribution of damages across iterations at the end of the simulation (year 50; bottom). Note that reported values correspond to a subset of species and are intended for demonstration only.

11. Discussion and recommendations

11.1. Questions that can be answered with this research

In this project, we developed simulation models for the container and cut flowers pathways and parametrised them with the department's import management and border inspection data. This approach allowed us to estimate the leakage rate of each pathway and test the effects of changing the values of the top level variables on model outputs. Looking at pathways independent of each other, one question that can be answered with this work is how future changes in a biosecurity risk profile may affect leakage rates. For example, on the container pathway, increasing the proportion of CAL-like countries strongly increased the leakage of high level contamination. Similarly, importing high risk lines of cut flowers (i.e., lines from high risk source countries) increased the leakage of actionable contamination. These results can give the department some indication of what magnitude of change in leakage to expect should more containers and lines from high risk countries arrive in Australia, which is a realistic future scenario. Furthermore, the use of the simulation models developed in this project can help decision makers to estimate what the consequences of changes in investment in risk controls might be. For example, increasing participation of CAL countries in the SCHS may not give a good return on investment, because the simulation results suggested that allocating more containers to the 100% sampling rate category (assuming an increased participation rate of CAL countries) did not affect the leakage rate.

11.2. Optimising investment in risk controls

The development of simulation models for two import pathways was the foundation for the main aim of the project, which was to develop a simple resource allocation tool that allows decision makers to vary investment in border risk controls across two or more pathways and observe the effects on the leakage rate and subsequently on damages to assets.

The Shiny web app developed in this project is a simple and intuitive interface for defining different investment scenarios that end users can use to inform the decision-making process. Investment scenarios compare damages relative to the status quo, so potential changes to border inspection policies can be tested in a safe environment without real life consequences, providing analysts with a deeper understanding of the interactions between border risk controls for the container and cut flowers pathways. The resource allocation tool is a proof-of-concept that has demonstrated the utility of adopting a simulation approach for addressing resource allocation questions. It is flexible to incorporate more import pathways and can also be adapted to integrate information on the cost of risk controls.

11.3. Generalisation of pathways

The development of the simulation models for the two pathways consisted of the same general steps: 1) build a conceptual diagram, 2) use that to inform the model structure, 3) parameterise the model based on data. Each general step followed the same process.

The conceptual diagram aimed to capture the border risk controls and processes currently implemented by the department on a particular pathway. These controls and processes are triggered by various general attributes of the units present on a particular pathway (e.g. containers, cut flower consignments). Different combinations of attributes' values trigger different interventions. While each pathway may have a different set of general attributes, we expect a certain level of commonality across pathways. For example the source country risk and the type of goods were common attributes for the two pathways. The container pathway had extra attributes which determined different interventions, like the destination or the presence of offshore interventions.

The (simulation) model structure mirrors the conceptual diagram and complements it with placeholder nodes that data is collected about, i.e. frequency and outcomes of interventions.

The set of general attributes determined (for both pathways) the type and frequency of interventions (i.e., documents assessment, internal inspections, external inspections). It is the same set of attributes (characterising the population of pathway units) which can explain the true (but unobserved) contamination rate. The observed/detected contamination rate is then simply a function of the true contamination and the detection capability (i.e., the inspection sensitivity). This reasoning can be applied to any pathway; the differences will arise from the different set of interventions undertaken. For example, for the container pathway we modelled multiple different types of internal and external inspections, whereas for the cut flower pathway, there was only one internal and one external inspection.

Assumptions about the inspection(s) sensitivity need to be made for all pathways, and based on these, the leakage rate can be estimated.

Parametrising the model requires the data to be processed and filtered such that alternative combinations of the states of the general attributes, along with their relative frequencies can be combined to produce a set of conditional probability tables that characterise the relative frequencies of interventions given all combinations of the states of the general attributes.

11.4. Operationalising the project outputs

Transferring the findings of this project into an operational context requires technical expertise, and access to the programming software R as well as adequate computing capacity. Foremost, departmental staff need to be comfortable with using and adapting the code base. If DAFF risk analysts wish to incorporate other pathways into the resource allocation tool, these pathways would need to be modelled first as described in this report. The modelling approach that was used in this project is generalisable and can be applied to other pathways as discussed in the previous section, which means that end users do not need to learn or understand other methods for data analysis and development of simulation models.

11.5. Recommendations

Based on our experience with the container and cut flowers data sets we offer the following recommendations to improve the quality of model outputs and to support the uptake of the resource allocation tool:

- **Improve data quality** by recording additional relevant information (e.g., proportion of stems inspected in a line or consignment of cut flowers — volume data is needed for this) and by curating database entries (e.g., introducing automatic data entry constraints so that missing data or incorrectly filed directions can be avoided; reviewing direction categories). Better data streams can facilitate the identification and matching of 1) units (e.g., containers, goods) with their attributes (source country, offshore/on shore treatment etc.), 2) applied actions/directions (inspection, inspection type, treatment etc.) with their outcomes (contamination found, release or export etc.).

- **Estimate inspection sensitivities**. To demonstrate the utility of the proof-of-concept resource allocation tool we assumed inspection sensitivities for the different types of border inspections undertaken on both pathways. However, when operationalising the project outputs, sensitivity parameters should be estimated through a formal expert elicitation process because inspection sensitivity directly influences the inferred approach rate and thus the leakage rate. Reducing uncertainty around inspection sensitivity parameters would improve the confidence in model outputs.

- **Hold workshop(s)** for use of the R code base. Successful integration of our research findings into departmental processes requires a solid understanding of the R packages developed in this project. CEBRA proposes to support the transition with tailored workshops for relevant departmental staff.

Acknowledgements

This report is a product of the Centre of Excellence for Biosecurity Risk Analysis (CEBRA). In preparing this report, the authors acknowledge the financial and other support provided by the Australian Department of Agriculture, Fisheries and Forestry, the New Zealand Ministry for Primary Industries and the University of Melbourne. In particular, we thank the advice and ongoing support provided by our department colleagues Blaine Wentworth, Pip Griffin, Richard Gao, Tristan Kearns, Vivian Xing, Ellen Miech, Margaret Yang, and Adam Usher. We would also like to thank two reviewers on CEBRA's independent Scientific Review Panel for their useful feedback on the draft report.

Bibliography

Chang W, Cheng J, Allaire J, *et al.* (2022) *shiny: Web Application Framework for R*. URL `https://CRAN.R-project.org/package=shiny`. R package version 1.7.2.

Craik W, Palmer D, Sheldrake D (2017) Priorities for Australia's biosecurity system. An independent review of the capacity of the national biosecurity system and its underpinning Intergovernmental Agreement. Tech. rep., Department of Agriculture and Water Resources, Canberra, Australia.

DAFF (2022) Country Action List (CAL). URL `https://www.agriculture.gov.au/biosecurity-trade/import/arrival/pests/cal#country-action-list`.

Department of Agriculture (2019) Final Pest Risk Analysis for Cut Flower and Foliage Imports—Part 1. Tech. rep., Department of Agriculture, Canberra. URL `https://www.agriculture.gov.au/sites/default/files/sitecollectiondocuments/biosecurity/risk-analysis/plant-reviews/final-report-cut-flowers.pdf`.

Department of Agriculture, Water and the Environment (2021) Final Pest Risk Analysis for Cut Flower and Foliage Imports—Part 2. Tech. rep., Department of Agriculture, Water and the Environment, Canberra. URL `https://www.agriculture.gov.au/sites/default/files/documents/final-report-cut-flowers-foliage-imports-part-2.pdf`.

Dodd A, Arndt E, Baumgartner J (2021) User requirements specification. CEBRA Project 21D: Value Added. Modelling the marginal return on investment within and across pathways. Tech. rep., Centre of Excellence for Biosecurity Risk Analysis, The University of Melbourne.

Dodd A, Arndt E, Hanea A (2022) Methods for determining the optimal level of investment in border-biosecurity risk controls with dependent benefits. CEBRA Project 21D: Value Added. Modelling the marginal return on investment within and across pathways. Tech. rep., Centre of Excellence for Biosecurity Risk Analysis, The University of Melbourne.

Dodd AJ, Stoeckl N, Baumgartner JB, Kompas TF (2020) Key Result Summary: Valuing Australia's Biosecurity System. Tech. rep., Centre of Excellence for Biosecurity Risk Analysis (CEBRA), The University of Melbourne, Melbourne. URL `https://cebra.unimelb.edu.au/__data/assets/pdf_file/0020/3535013/CEBRA_Value_Docs_KeyResultSummary_v0.6_Endorsed.pdf`.

Hijmans RJ (2023a) *raster: Geographic Data Analysis and Modeling*. URL `https://CRAN.R-project.org/package=raster`. R package version 3.6-13.

Hijmans RJ (2023b) *terra: Spatial Data Analysis*. URL `https://CRAN.R-project.org/package=terra`. R package version 1.7-39.

Højsgaard S (2012) Graphical independence networks with the gRain package for R. *Journal of Statistical Software*, **46**, 1–26. doi:10.18637/jss.v046.i10. URL `https://www.jstatsoft.org/v46/i10/`.

Holt J, Leach AW, Johnson S, *et al.* (2018) Bayesian Networks to Compare Pest Control

Interventions on Commodities Along Agricultural Production Chains. *Risk Analysis*, **38**, 297–310. doi:10.1111/risa.12852.

IGB (2018) Hitchhiker pest and contaminant biosecurity risk management in Australia. Tech. rep., Inspector-General of Biosecurity, Department of Agriculture and Water Resources, Canberra, Australia.

IGB (2019) Pest and disease interceptions and incursions in Australia. Tech. rep., Inspector-General of Biosecurity, Department of Agriculture and Water Resources, Canberra, Australia.

International Standards for Phytosanitary Measures (2008) ISPM No. 31 Methodologies for Sampling of Consignments. Tech. rep., IPPC. URL `https://www.ippc.int/en/publications/83473/`.

Jamieson L, Woodberry O, Mascaro S, Meurisse N, Jaksons R, Brown S, Ormsby M (2021) An Integrated Biosecurity Risk Assessment Model (IBRAM) For Evaluating the Risk of Import Pathways for the Establishment of Invasive Species. *Risk Analysis*, **42**, 1325 – 1345. doi:https://doi.org/10.1111/risa.13861. URL `https://doi.org/10.1111/risa.13861`.

Korb K, Nicholson A (2010) *Bayesian Artificial Intelligence*. CRC Press, Inc., 2 edn., 491 pp.

Landau WM (2021) The targets r package: a dynamic make-like function-oriented pipeline toolkit for reproducibility and high-performance computing. *Journal of Open Source Software*, **6**, 2959. URL `https://doi.org/10.21105/joss.02959`.

Pearl J (1988) *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers, San Mateo, California.

R Core Team (2022) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL `https://www.R-project.org/`.

Sardain A, Sardain E, Leung B (2019) Global forecasts of shipping traffic and biological invasions to 2050. *Nature Sustainability*, **2**, 274–282. doi:10.1038/s41893-019-0245-y.

Scutari M (2010) Learning Bayesian Networks with the bnlearn R package. *Journal of Statistical Software*, **35**, 1–22. doi:10.18637/jss.v035.i03.

A. Provided reading material

The department provided the following reading material on the container and cut flowers pathways, including work instructions, policy documents and e-mail content.

A.1. Container pathway

Reading material included:

- Inspecting sea containers at the wharf exit gate (Biosecurity Act 2015).pdf
- Managing high risk pathway sea cargo.pdf - this document covers only CAL countries
- SCHS Quality Manual PNG & SOL ver 3.0 (Sept 2014).pdf
- SCHS COP LAE PORT SERVICES v 4.0.pdf
- Performing tailgate and rural tailgate inspections.pdf
- Undertaking cargo compliance verification inspections.pdf
- Wharfgate Week -2 Stats July BNE.xlsm
- Assessing Cargo Risk Analysis profiles in the Agriculture Import Management System.pdf
- OSN 2022-107 Changes to the assessment of Cargo Compliance Verification (CCV) entries.pdf

The DAFF project lead sent the following e-mail content on 11 April 2022. It included further explanations from the DAFF operational area on the container inspection process.

Container inspection processes including how containers are sampled (chosen for inspection):

CAL (managed in S Cargo and AIMS)

- All CAL full or empty containers are held in the ICS via profiles, the profile targets all containers or break bulk which have either a Country of origin or a Port of Loading which is on the Country Action List (CAL). The profile applies a GAS Hold in the ICS which is then communicated to S Cargo. If any containers on the vessel are GAS Held then the vessel and all its containers and break bulk are referred to S Cargo. The Container NCC (CNCC) then issue CAL inspection directions and a list of cargo that is GAS Held to the principle shipping agent and cc the stevedore. These receive a full six-sided external inspection on the wharf at the first port of discharge by a biosecurity officer.

- When cargo passes external inspection the inspecting officer lifts the GAS Hold in S Cargo. S Cargo communicates a GAS Release message to ICS which lifts the hold. The inspection results are sent to the CNCC who enter them directly into S Cargo (it is intended that inspecting officers will perform this task in the future). If concerns are found and the wharf does not have a wash pad a very manual conditional release process is actioned in both S Cargo and ICS by the CNCC to allow the container to move to an AA for treatment.

- All empty CAL containers go to an AA for an internal inspection once all external issues are resolved. Containers arriving with goods (fulls) are not subject to this process.

- The only exceptions made are for containers that have been processed through the Sea Container Hygiene System (SCHS) or those that are transhipping through Australia and are not destined for domestic use. These containers will be GAS Exempted upon application to CNCC.

- The CAL policy and profiles do not target goods inside containers. AIMS profiles manage any goods concerns. The GAS Hold is the highest hierarchy in the ICS system and no other intervention can occur until this is lifted.

- When looking at S Cargo reports keep in mind there are 5 GAS Hold lift types to consider:
    1. GAS RELEASED = inspected and released
    2. GAS EXEMPTED = not inspected (tranship approved)
    3. GAS CANCELLED = not inspected usually withdrawn cargo that did not arrive
    4. SCHS CANCELLED = not inspected due to reduced intervention SCHS processing
    5. CONDITIONAL RELEASE = containers that had to move off wharf for further treatment or inspection (S Cargo currently does not change this status to released, Cond Rel is manually lifted directly in the ICS).

Rural tailgate (RTG) (managed in AIMS)

- Containers destined for rural areas rely on brokers to declare this in the CRA profile questions. If declared for rural destination the container is referred to AIMS for RTG inspection in the metro area of discharge prior to delivery to the rural destination.

- These are controlled via Conditional Releases (ICS to AIMS) which allow the cargo to exit the wharf.

- Rural tailgate and Tailgate WI (attached).

Tailgate (managed in AIMS)

- Tailgate inspections vary. They are directed where concerns are identified due to not meeting documentary requirements such as packing declarations etc. If concerns are identified when the broker answers the CRA profile questions or for some commodities such as plant or biological stock feed that require sampling usually imported under a permit.
- These are controlled via Conditional Releases (ICS to AIMS) which allow the cargo to exit the wharf.

Wharf-gate inspection

- Containers randomly inspected at the wharf exit gate. Very few regions continue to man the exit gates due to staffing. The CNCC collect and store any data received example attached.
- WI for wharf gate can be found attached.
- This is random so no hold type is used (could be considered a leakage survey).

External container inspection (managed in AIMS)

- These are generally performed when a rural tailgate or tailgate inspection cannot be performed due to dangerous goods or other WHS/biosecurity concerns.

Policy and operational details of the CCV scheme

- CCV policy is here.
- WI is attached.
- My understanding of CCV is that there are ICS profiles which randomly select cargo (2% was a figure I was given) for CCV to verify goods are as declared and/or to target goods we don't normally inspect. I have requested the details of the profile from Targeting and will provide when I have them.

Policy and operational details of the SCHS scheme and associated verification/assurance.
SCHS (managed in S Cargo and AIMS)

- Currently SCHS facilities are located in ports that are on the CAL.
- The SCHS facility send the CNCC a list of containers that have been treated (one list per vessel), this is uploaded to S Cargo. S Cargo then changes the risk type from HIGH (CAL) to SCHS and applies the intervention rate (5, 20, 50 or 100%). Containers not selected for inspection are SCHS Cancelled and do not receive an inspection.
- The SCHS management policy, a copy of a Quality Manual covering audits etc and a sample of a COP are attached (internal documents). An information sheet can be found here.

A.2. Cut flowers pathway

The department provided the following reading material:

- `https://www.agriculture.gov.au/biosecurity-trade/import/goods/plant-products/cut-flowers-foliage` - webpage contains an introduction to overall process,
- *1. Sampling and inspecting imported cut flowers.pdf* - document containing work instructions,
- *2. Horticultural produce and cut-flower inspection techniques.pdf* - document containing work instructions,
- *3. Policy for managing imported fresh cut flowers and foliage that do not comply with import conditions.pdf* - document containing policy for cut flowers,
- *4. 20210622_Cut-Flower-pathway-modelling-leakage (O).docx* - a draft report on modelling cut flower biosecurity risk,
- *5. Sampling_CF_units_v2.pdf* - an ABARES draft report, *Within consignment sampling of cut flowers: choosing what to Sample (2017)*,
- *6. Sampling_discussion_final.pdf* - an ABARES draft report, *Within consignment sampling of cut flowers: issues to consider (2017)*,
- *7. BICON - Import Conditions - Overview.pdf* - overview of the different subpathways within the cut flowers pathway,
- *8. BICON - Import conditions detailed.pdf* - detailed conditions for different flowers and foliage types and different source countries.

Additionally, the department also provided a *Response to CEBRA data request.docx* document accompanying the cut flowers dataset, describing policy changes in the last 5 years.

B. Consultation with DAFF experts

During the development of the conceptual diagrams for the container and cut flowers pathways the CEBRA project team consulted with a range of DAFF experts who are listed in Table B.1.

Table B.1.: DAFF experts consulted in the development of the conceptual diagrams.

Name	Role
Jose Arias	Assistant Director, Profiling and Targeting
Richard Gao	Assistant Director, Biosecurity System Modelling
Greg Hankins	Assistant Director, Profiling and Targeting
Rama Karri	Director, Hitchhiker Working Group
Tim Killesteyn	Pathway Policy (Cargo and Conveyances)
George Peyiotou	Data Analyst, Biosecurity Analytics Centre
Cindy Pretty	Assistant Director, Hitchhiker Working Group
Sarah Rake	Assistant Director, Hitchhiker Working Group
Rachel Slatyer	Assistant Director, Biosecurity Analytics Centre
Tamara Thorn	Assistant Director, CCV Program, Profiling and Targeting
Jon Webber	Director, Approved Arrangements
Craig Scheibel	Assistant Director, Horticulture Imports Program
Douglas Kerruish	Director, Horticulture Imports Program
Adam Usher	Assistant Director, Biosecurity System Modelling
Tristan Kearns	Business Analyst, Biosecurity System Modelling
Oscar Yang	Applied Statistician, Biosecurity System Modelling
Qiang Li	Data Scientist, Biosecurity System Modelling

C. Deriving the structure of the model (technical details)

In a BN, all nodes represent variables which are probabilistic and have either discrete or continuous distributions. Arcs represent dependence between parent and child nodes (an arc points from a parent to a child) and this dependence is captured by the conditional distribution of a child given its set of parents. However, some of the inspection nodes are deterministic given the parents, these are: I Internal, I Unpack, I RTG Low Risk, and I RTG High Risk.

The variables represented as nodes in the container pathway BN are listed and defined in Table C.1.

Table C.1.: The BN nodes for the container model

Node name	Definition	Node type / States	Estimated from
CAL Country	The source country of a random container, classified as being on the current list of countries on the action list or not. CAL stands for country action list	binary (true/false)	data
Rural Destination	The destination of a random container being rural or not, informed by the current rural/metro classification	binary (true/false)	data
Type Of Goods	The type of goods a random container is filled with	discrete/categorical (normal/exceptions/none)	data
SCHS	A random container being on the SCHS or not; when on the scheme, different sampling rates for inspection are covered by the "yes" states	discrete/categorical (no, yes5, yes20, yes50, yes100)	data
SCHS Reliability	The reliability of the SCHS which was intended to inform quarterly sampling rates for the containers on the SCHS; not used, hence not properly defined	discrete or continuous	data
Approach Rate Internal	The true internal approach/contamination rate inferred from Output Internal Inspection and Sensitivity Internal Insp	discrete/ordinal (high level contamination - HLC, low level contamination - LLC, no contamination - none)	function

Continued on next page

Node name	Definition	Node type / States	Estimated from
Approach Rate External	The true external approach/contamination rate inferred from Output External Inspection and Sensitivity External Insp	discrete/ordinal (high level contamination - HLC, low level contamination - LLC, no contamination - none)	function
Requires Full Unpack	A random container requiring a full unpack inspection	binary (true/false)	data
I Unpack	A random container/goods being inspected with this inspection	binary (true/false)	data
I CAL6sided	A random container/goods being inspected with this inspection	binary (true/false)	data
I External	A random container/goods being inspected with this inspection	binary (true/false)	data
I Internal	A random container/goods being inspected with this inspection	binary (true/false)	data
I RTG High Risk	A random container/goods being inspected with this inspection	binary (true/false)	data
I RTG Low Risk	A random container/goods being inspected with this inspection	binary (true/false)	data
I TG	A random container/goods being inspected with this inspection	binary (true/false)	data
I CCV	A random container/goods being inspected with this inspection	binary (true/false)	data
I Wharfgate	A random container/goods being inspected with this inspection	binary (true/false)	data

Continued on next page

Node name	Definition	Node type / States	Estimated from
Output Internal Insp.	Observed internal contamination level of a random container	discrete/ordinal (high level contamination - HLC, low level contamination - LLC, no contamination - none)	data
Output External Insp.	Observed external contamination level of a random container	discrete/ordinal (high level contamination - HLC, low level contamination - LLC, no contamination - none)	data
Sensitivity Internal Insp	the probability of internal contamination to be found when present in a random container	constant value	guess; ideally replaced by structured expert judgement
Sensitivity External Insp	the probability of external contamination to be found when present on a random container	constant value	guess; ideally replaced by structured expert judgement
Treatment Internal	the probability of a random container to be sent to internal treatment	binary (true/false)	data
Treatment External	the probability of a random container to be sent to external treatment	binary (true/false)	data

Continued on next page

Node name	Definition	Node type / States	Estimated from
Leakage Rate Internal	the probability of a random container to be internally contaminated or containerised goods being contaminated and undetected as such	binary (true/false)	function
Leakage Rate External	the probability of a random container to be externally contaminated or containerised goods being contaminated and undetected as such	binary (true/false)	function

From the 86 arcs in the container pathway BN (see Figure 5.1) only 41 correspond to probabilistic relationships. This means we need to estimate 13 conditional distributions (i.e., conditional probability tables, accounting for 1808 conditional probabilities). The lack of arcs assume conditional independencies which were not yet validated, but rather implied by the conceptual diagram. Dependencies represented by un-directed links were estimated as a four dimensional joint distribution. To do that we used a local structure learning algorithm available in gRain, which assigned a direction to undirected arcs such that the joint dependence is well approximated with a DAG. The direction of these arcs however is not discussed since it is simply a (not-unique) graphical representation of a correlation matrix.

Similarly for the cut flower pathway BN in Table C.2.

From the 38 arcs in the cut flowers pathway BN only 23 correspond to probabilistic relationships. This means we need to estimate 10 conditional distributions (i.e., conditional probability tables, accounting for 368 conditional probabilities). As above, lack of arcs corresponds to the assumption of conditional independencies which were not yet validated, but rather implied by the conceptual diagram. The non-probabilistic relationships are functional.

Table C.2.: The BN nodes for the cut flowers model

Node name	**Definition**	**Node type / States**	**Estimated from**
Source Country	The risk level of the source country of a random consignment informed by the current list of exporting countries	discrete/ordinal(high risk, medium risk, low risk)	data
Type Of Goods	The type of goods a random consigment if filled with	discrete/categorical (flowers, foliage)	data
Initial Document Assessment Result	The type of documents a random consignment is accompanied by; informed by the existence and completeness of import documentation (assuming all units have an initial document assessment)	discrete/categorical (OK, not OK, incomplete)	data
Inspection Internal	A random container/goods being inspected with this inspection (sample inspection)	binary (true/false)	data
Inspection External	A random container/goods being inspected with this inspection (carton & package inspection)	binary (true/false)	data
Outcome Internal Inspection	Observed internal contamination level of a random consignment	discrete/categorical (clean compliant, clean non-compliant, infested actionable, infested non-actionable)	data
Outcome External Inspection	Observed external contamination level of a random consignment	discrete/categorical (clean compliant, clean non-compliant, infested)	data
Contamination Internal	The true internal contamination (actionable and not actionable) of a random consignment	discrete/categorical (clean, infested)	function
Contamination External	The true external contamination (actionable and not actionable) of a random consignment	discrete/categorical (clean, infested)	function
Treatment	the probability of a random consignment to be sent to treatment	binary (true/false)	data
Destroy and Sendback	the probability of a random consignment to be sent to be destroyed or sent back	binary (true/false)	data
Contaminated External Not Actionable	the probability of a random consignment to be classified as this	binary (true/false)	function
Contaminated Not Actionable	the probability of a random consignment to be classified as this	binary (true/false)	function
Cleaned	the probability of a random consignment to be cleaned	binary (true/false)	function
Contaminated Internal Not Actionable	the probability of a random consignment to be classified as this	binary (true/false)	function
Contaminated Internal Actionable	the probability of a random consignment to be classified as this	binary (true/false)	function
Contaminated Actionable	the probability of a random consignment to be classified as this	binary (true/false)	function
Leakage Actionable	the probability of a random consignment to be contaminated and undetected as such	binary (true/false)	function
Leakage Not Actionable	the probability of a random consignment to be contaminated and undetected as such	binary (true/false)	function

D. Annotated structure of the container pathway BN

Figure D.1 provides an annotated version of the BN for the container pathway to improve readability. It shows the top level correlated attributes and groups interventions (i.e., internal external inspections) below. The true contamination on the pathway consists of the internal and external approach rate and the detected/treated contamination is represented by the output and treatment nodes in the net. The leaked contamination is a combination of internal and external leakage.

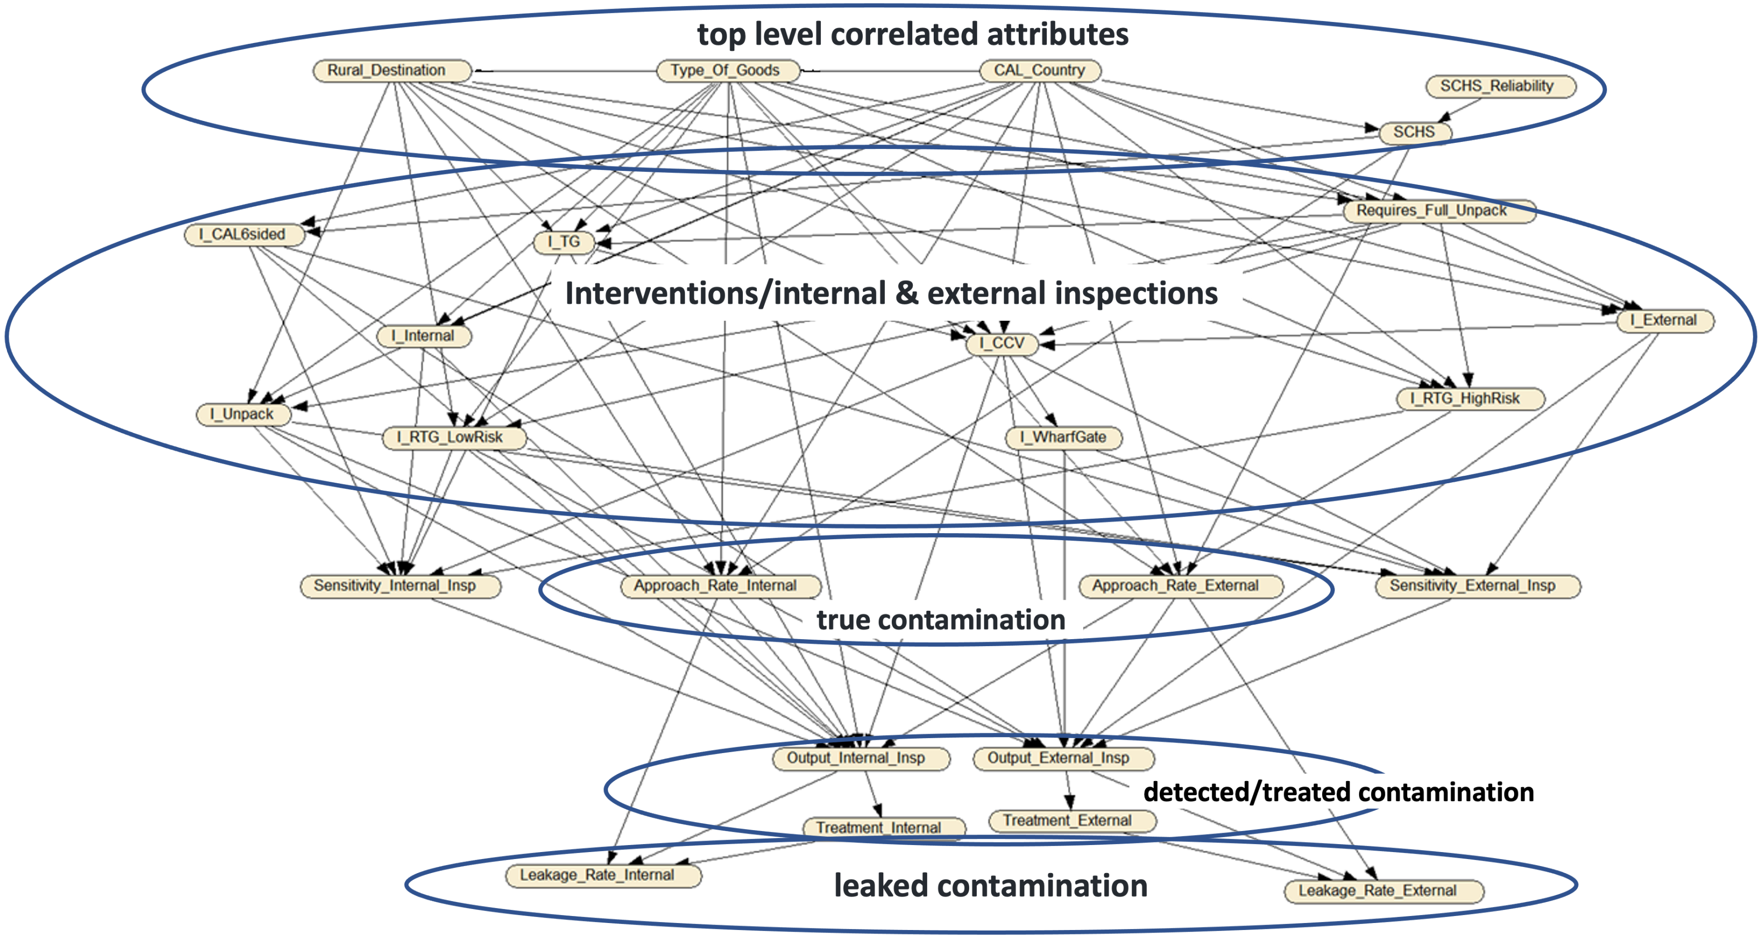

Figure D.1.: Annotated version of the container pathway BN.

E. Additional results on the counterfactual analysis of the container pathway

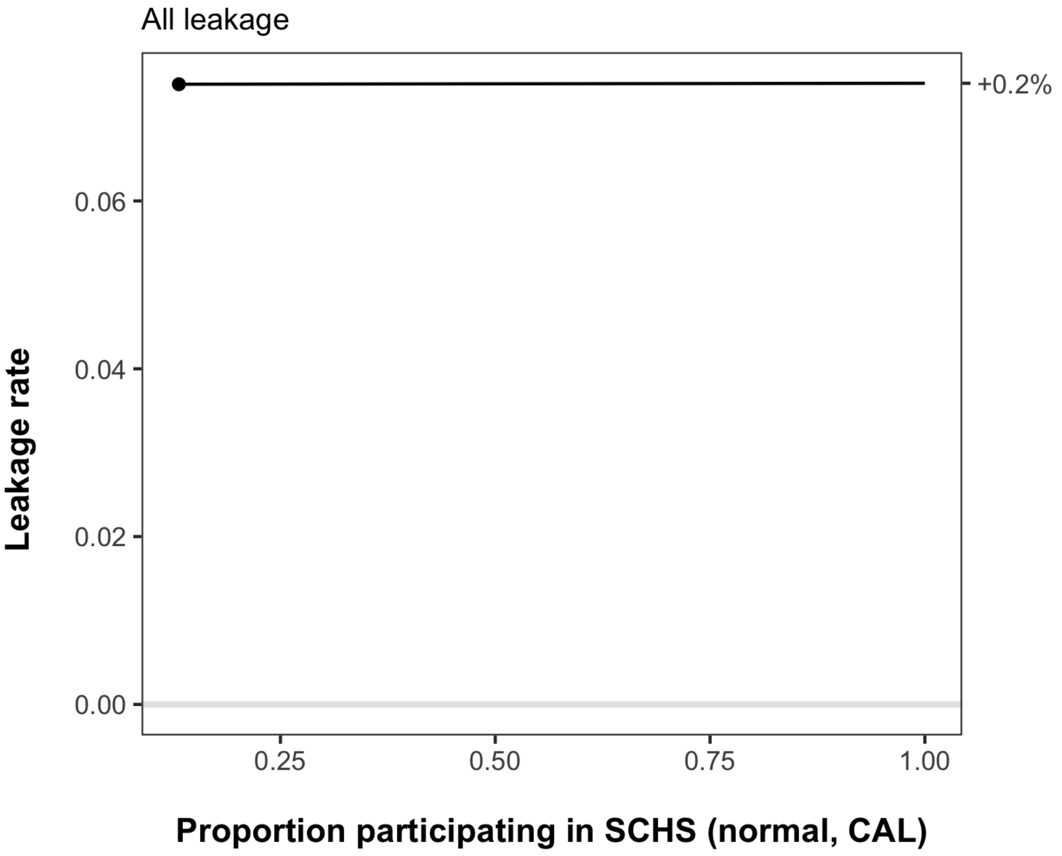

Figure E.1.: Increasing the proportion of CAL containers with normal goods participating in the SCHS. The point on the line represents the status quo of the proportion of CAL containers with normal goods that are participating in the SCHS (i.e., 13.2%). The percentage indicated in the right margin gives the change in leakage relative to the status quo.

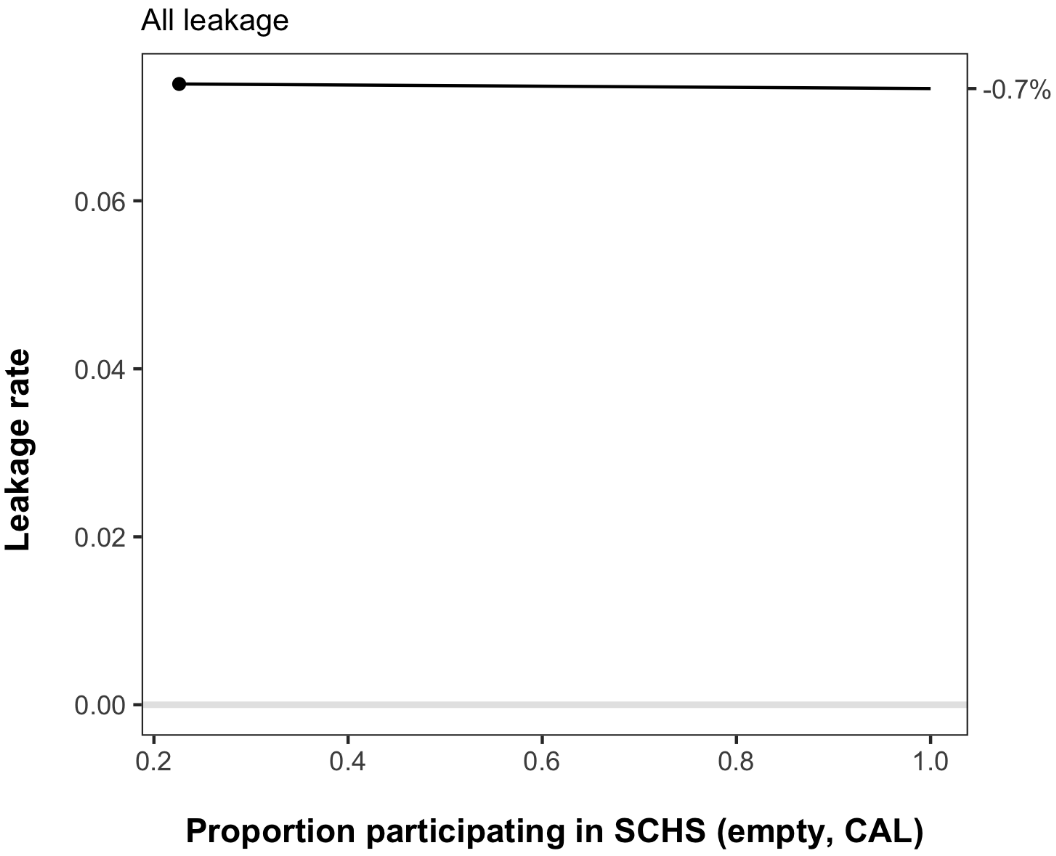

Figure E.2.: Increasing the proportion of empty CAL containers participating in the SCHS. The point on the line represents the status quo of the proportion of CAL containers with normal goods that are participating in the SCHS (i.e., 22.6%). The percentage indicated in the right margin gives the change in leakage relative to the status quo.

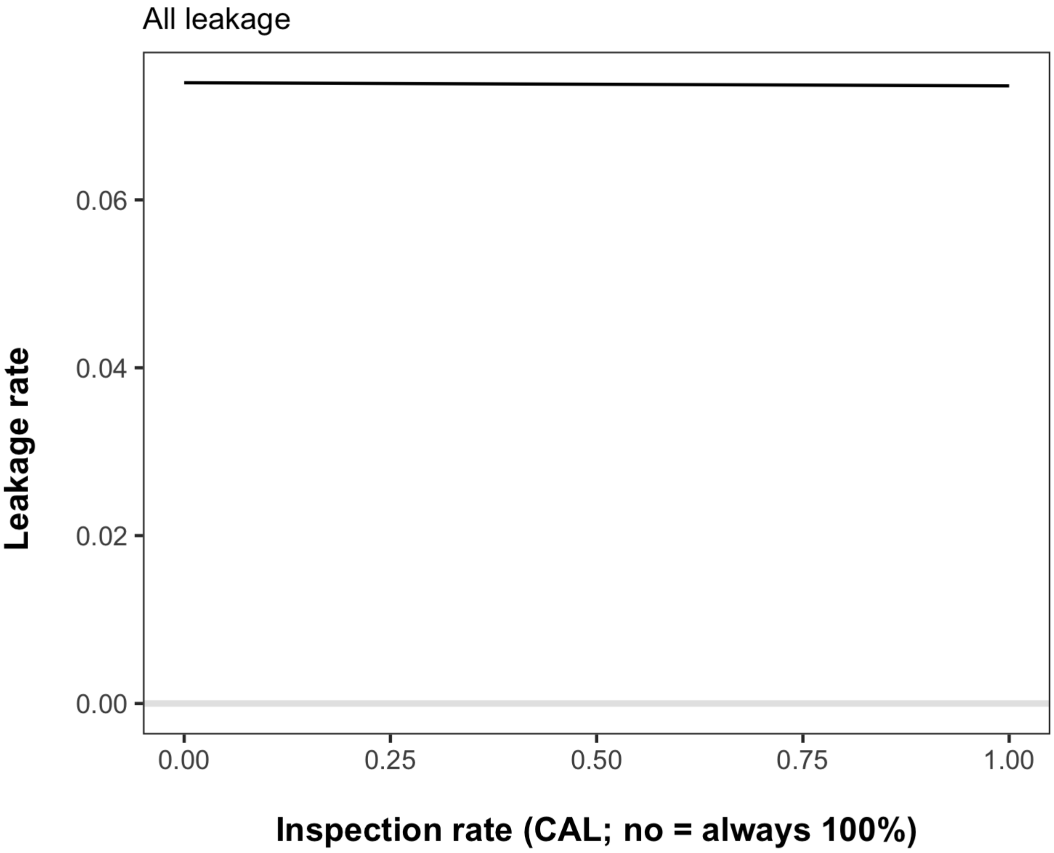

Figure E.3.: Varying the inspection rate of CAL containers participating in the SCHS.

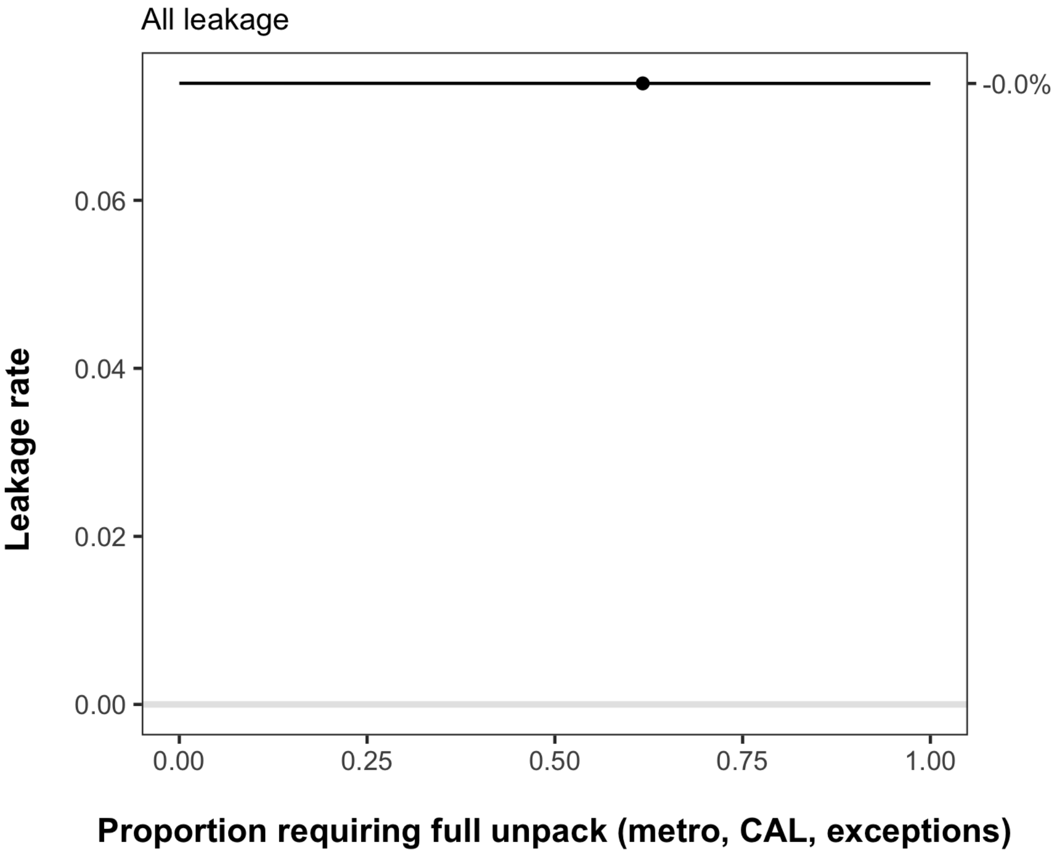

Figure E.4.: Changing the proportion of CAL metro containers with goods exempt from internal inspection undergoing a full unpack. The point on the line represents the status quo of the proportion of CAL metro containers with exempt goods that are being subjected to a full unpack. The percentage indicated in the right margin gives the change in leakage relative to the status quo.

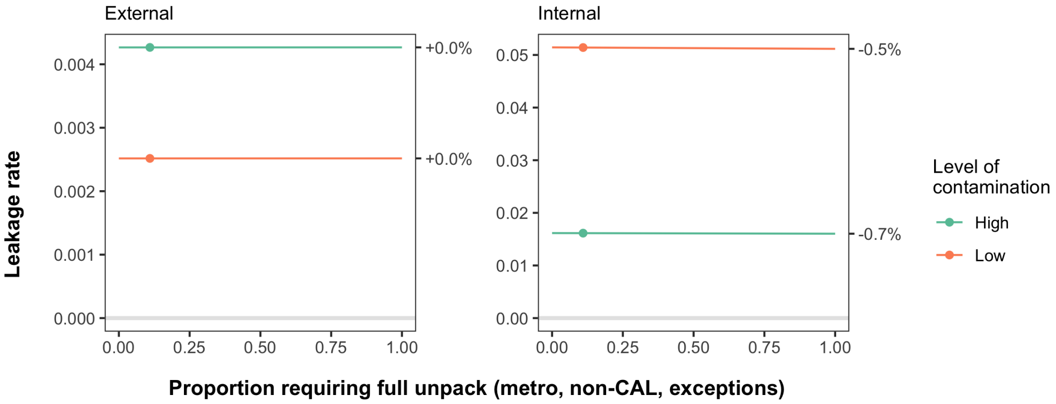

Figure E.5.: Changing the proportion of non-CAL metro containers with goods exempt from internal inspection undergoing a full unpack. Points on lines indicate the leakage rates corresponding to the status quo of the proportion of non-CAL metro containers with exempt goods that are being subjected to a full unpack. High levels of external and internal contamination are represented by green lines and low levels of contamination by orange lines. Percentages indicated in the right margins are changes in leakage relative to the status quo.

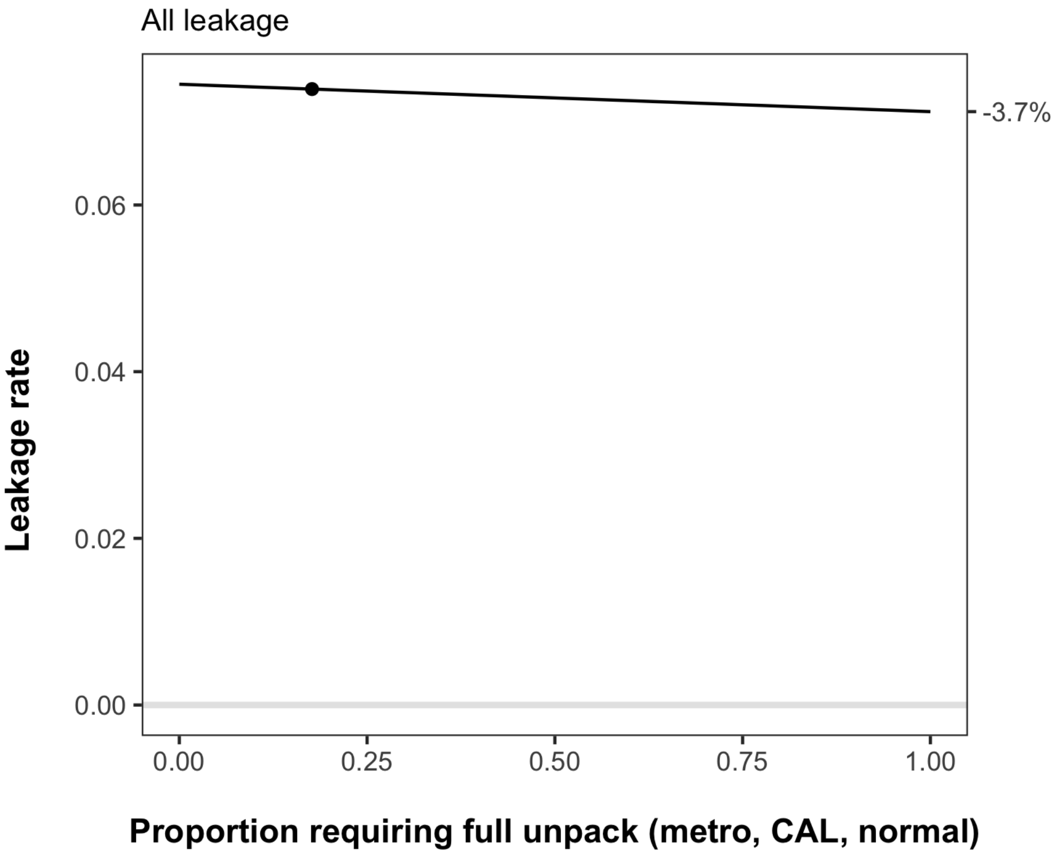

Figure E.6.: Changing the proportion of CAL metro containers with normal goods undergoing a full unpack. The points indicates the leakage rate corresponding to the status quo of the proportion of CAL metro containers with normal goods that are subjected to a full unpack. The percentage indicated in the right margin gives the change in leakage relative to the status quo.

F. Additional results on the uncertainty in outputs of the cut flowers pathway

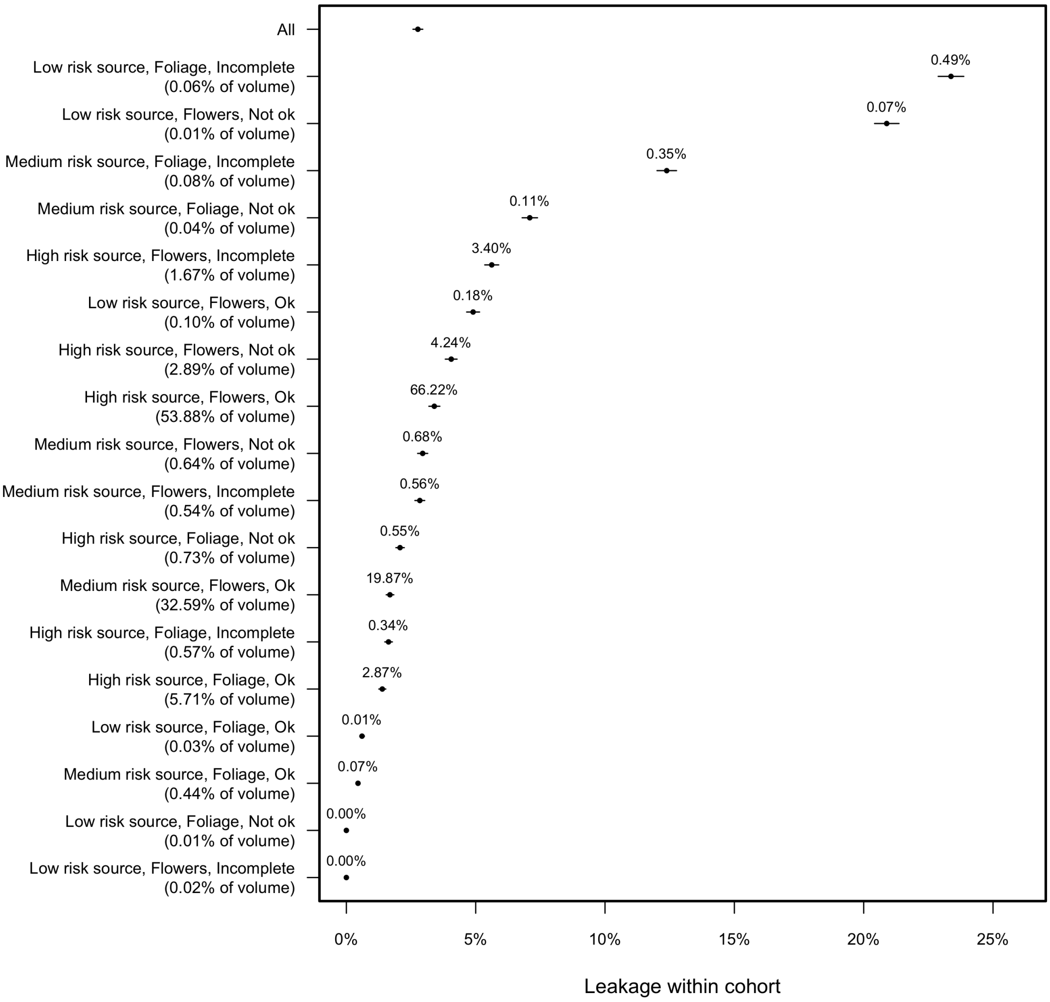

Values above points indicate contribution to overall leakage.
Points and intervals give median and 95% interval assuming 20,000 lines.

Figure F.1.: Leakage within cohorts of group (SourceCountry × TypeOfGoods × InitialDocAssResult). The y-axis shows the different cohorts and their respective proportion of the total import volume, while the x-axis presents the leakage rate. Dots represent the mean leakage and bars the 95% confidence interval. The values above the dots indicate the contribution (in percent) of the cohort to the overall leakage rate.

G. Technical details regarding the cut flowers data processing

The aim of the data cleaning and initial processing is to turn the provided data into the format suitable for the next stage of analysis. For each line, we kept their source country, tariff ID (as this can be relevant for determining flower type), goods description (more information the line contents), in case we need that for future/further processing. But the key attributes we want are: the source country risk level, the type of goods, the initial documentation assessment result, whether the packaging was inspected, and its outcome if inspected, whether the line was inspected and the outcome if inspected, whether there was any treatment, and whether the line was ultimately destroyed/sent back, or released.

However, many of the lines in the data did not have all of this information. As such we made a number of assumptions which we list in the following section.

G.1. Assumptions made in the cut flowers data processing

1. In the case with multiple document directions, the priority order is: *not ok* > *incomplete* > *ok* (i.e. we assume the 'worse' case).
2. In the case of multiple final directions, the priority order is: *Released* > *destroyed/-exported*. Note:
    a) if a line has two directions, *partial commodity (destruction)* and *final directive release*, then we assumed *Released* was the final outcome.
    b) However, if a line only has a *partial commodity disposal/destruction* direction and no final directive, then we have assumed that the entire line is destroyed.
3. In the case of multiple inspection outcomes, the priority order is:
    a) *Infested actionable* > *infested not actionable*,
    b) *Clean not compliant* > *clean compliant* (i.e., we assumed the 'worse case').
4. 'Verification' (usually package inspection) includes 'Compliance Verification'.
5. If there is no documentation directive, we assumed the documentation result to be *incomplete*.
6. If there was no packaging inspection direction, we assumed that no external inspection occurred, with a default *clean compliant* outcome

7. If there was no sample inspection direction, we assumed that no internal inspection occurred, with a default *clean compliant* outcome

8. If there was no treatment direction, we assumed there was no treatment.

9. If there was no final direction or unclear final direction, we assumed the line was *Released*.

10. Additionally, anything direction with 'Not performed' or 'No result' is assumed to not occur.

G.2. Additional caveats with the cut flowers data

In addition to the assumptions made regarding various documentation and inspection outcomes, there are some caveats to note regarding the data:

1. Note that we had not requested time stamps, and did not expect that there would be multiple directions of the same type. This is why we made those 'priority' orderings in the previous section. This also meant that directions that said "See next direction" were particularly ambiguous.

2. There were unclear direction categories such as 'Staff Information See Comments'. The data accompanying these directions was often not used.

3. Some lines appeared to have destruction after treatment.

4. The department was unable to include any offshore measures directly in the data set. This means we did not include attributes about whether the line of cut flowers had off-shore treatment, systems approaches etc.

5. Some directions appear to have been filed incorrectly, such as a line that had inspection "Not performed", but marked that insects had been found in the hazards field and free text field.

6. We must stratify lines by their attributes and properties (such as inspection outcomes) in order to make the conditional probability tables required for the simulation model. However, there were missing combinations. We temporarily fixed this by using averaging across available combinations.

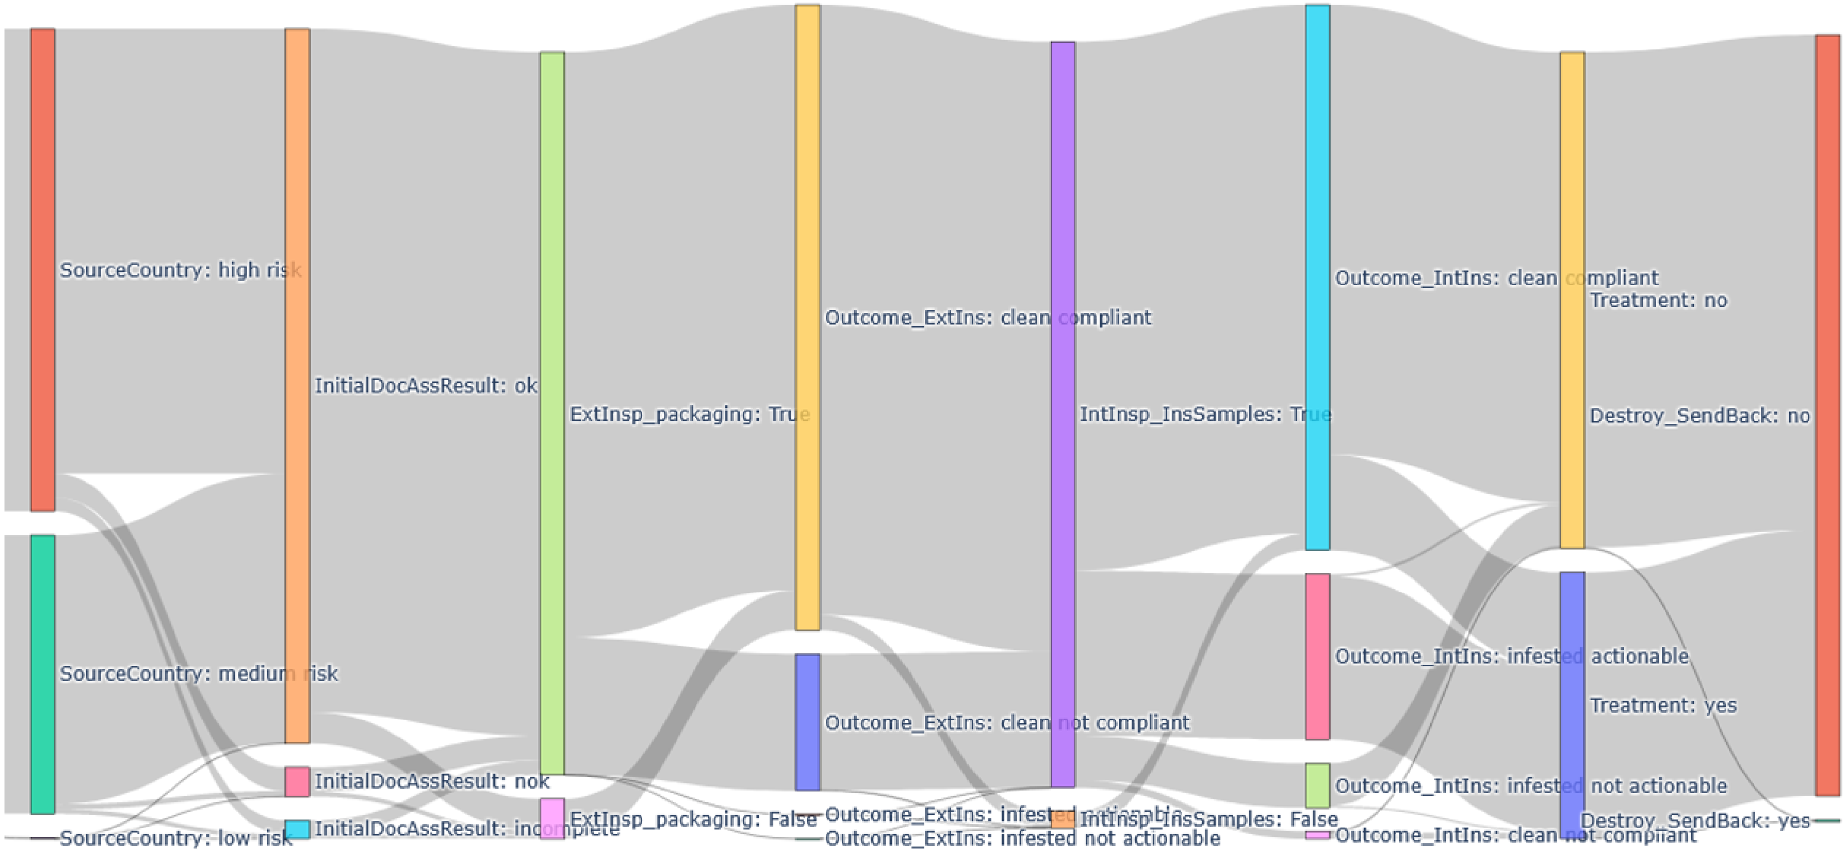

Figure G.1.: Flows of cut flower lines.

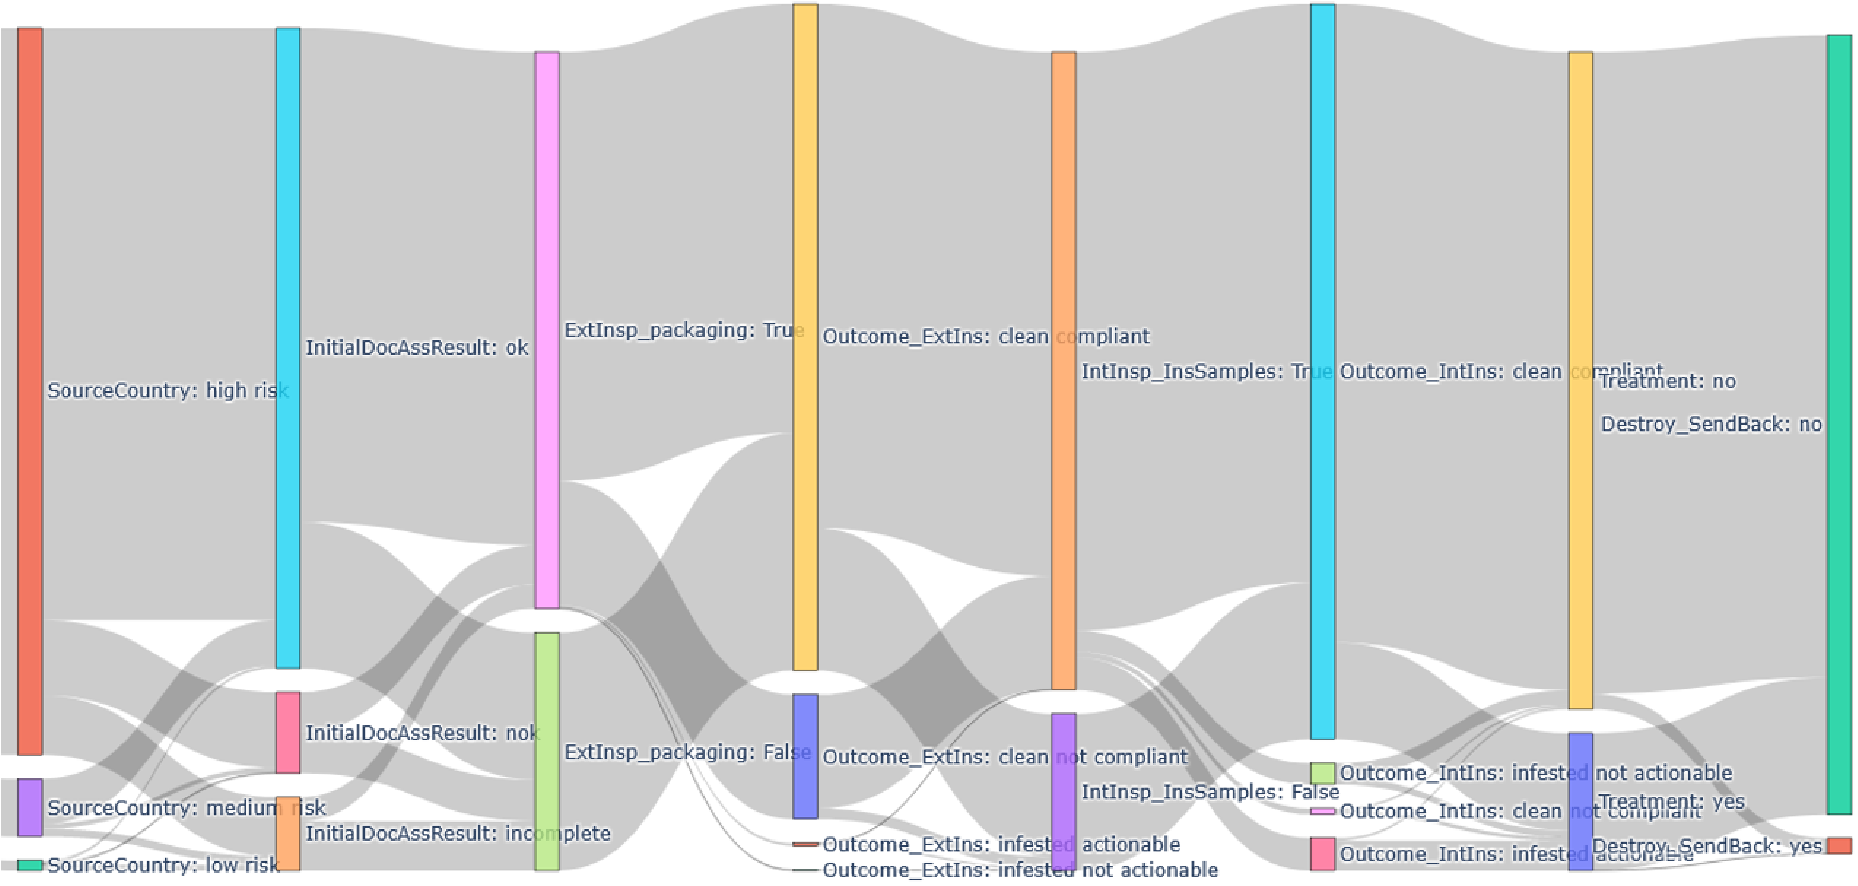

Figure G.2.: Flows of foliage lines.

H. Improvements to the Value Model

To facilitate the integration of the pathway models with the Value Model and permit assessment of a set of alternative scenarios of control via the Tool (Chapter 10), several Value Model enhancements were necessary.

The most significant update to the model involved optimisation of the spatiotemporal simulation algorithm through (1) use of precalculated spread trajectories agnostic to hazards' arrival rates (i.e., the "trajectory database"), and (2) a transition to simulating species independently and subsequently combining the outcomes when calculating damages - rather than simulating species together - thereby vastly improving memory efficiency and enabling enhanced parallelisation. Other key changes included refactoring R code to remove dependency on packages that are soon to be deprecated, in favour of their faster replacements.

H.1. Improved simulation algorithm

The original implementation of the Value Model, as produced for CEBRA Project 170713 (Dodd *et al.*, 2020), constructed time series of multiple hazards' arrival and spread by simulating, for each of many iterations, the stochastic arrival process and the subsequent spread. That is, for each iteration, it looped over hazards, and for each hazard it would: (1) determine the random number of arrivals (if any) given the hazard's arrival rate (sampling from a Poisson distribution with lambda equal to the arrival rate); (2) if one or more arrivals occur, determine the location of those arrivals (sampling randomly from all host cells, weighted by human population density); (3) spread the hazard to cells within the radial dispersal distance of the already infected cells; and (4) simulate the stochastic long-distance jump diffusion process where relevant.

The above steps were necessarily repeated for each iteration, and for each alternative arrival rate. For Project 170713, two alternative arrival rates were examined (i.e., biosecurity system *on* vs *off*), but for Project 21D, several alternative arrival rates must be considered, corresponding to a wider range of alternative scenarios of control. The need to simulate the spread process for each set of arrival rates proved a critical bottleneck impeding the rapid simulation required to deliver Project 21D's resource allocation tool. However, spread dynamics subsequent to arrival are unaffected by arrival rate, and we took advantage of this fact to achieve a vast improvement in run speed. Our redesigned simulation process involves preparing a single database of realisations of arrival and spread (i.e., many individual time series of infection spread, given a single arrival at the first time step), and then sampling from this database according to the arrival rate of interest, offsetting the sampled time series as appropriate given the time step of arrival. Specifically, after the once-off process of preparing the trajectory database is complete, a simulation is constructed for each hazard as follows (see also Figure H.1):

1. sample the number of arrivals from a Poisson distribution for each time step (e.g.,

for a ten year simulation, the number of arrivals for each of the ten years may be 1, 0, 2, 1, 0, 0, 0, 0, 2, 0);

2. draw, with replacement, those numbers of trajectories from the trajectory database (six trajectories in total, for our example); and
3. combine the trajectories by lagging the start of their spread dynamics to commence at the corresponding times of arrival (and ignoring the excess spread beyond the simulation duration) and then calculating the union of the sets of cells infected at each time step.

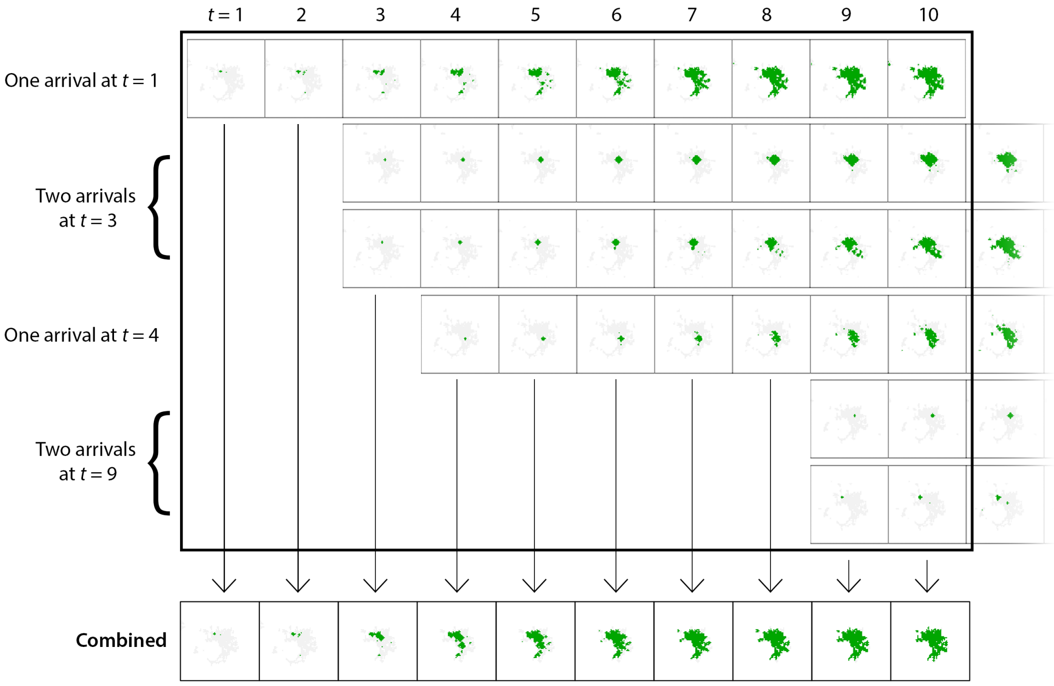

Figure H.1.: A simplified example demonstrating how trajectories are sampled from the trajectory database and combined to construct a single simulation of the arrival and spread of a single hazard. In this example, columns represent time steps (1 through 10) of the simulation, rows represent trajectories drawn from the trajectory database, and within individual panes grey pixels represent uninfected host cells and green pixels infected host cells. Arrivals occur at time steps 1 (n=1), 3 (n=2), 4 (n=1), and 9 (n=2). The corresponding number of trajectories are randomly sampled (with replacement) from the database, offset by their start time (with resulting excess dynamics beyond year 10 ignored), and finally combined into a single time series by calculating the union of infected cells at each time step.

This process is repeated for each hazard, and the resulting damage (yield reduction and/or trade loss) is calculated as in the original Value Model implementation by relating infection dynamics to the values of the underlying assets.

Relative to the original implementation, the trajectory database approach represents a substantial improvement in computational efficiency, with run speed and memory management greatly enhanced, particularly for hazards with many host cells and high arrival rates. In our tests, a trajectory database of 50,000 trajectories per hazard was created in just under 60 hours using 44 cores, suggesting a rate of approximately 19 trajectories (of all hazards) per core per hour. Memory usage peaked at around 3.2GB per core. The rate of trajectory generation varies greatly across hazards, from 112 trajectories per core per hour (screwworm fly, bluetongue) to 9,626 trajectories per core per hour (HPAI), with a median of 2,035 trajectories per core per hour (Table H.1). It is possible that, at least for some hazards, 50,000 trajectories is more than required to adequately represent the distribution of possible scenarios of spread. One could determine hazard-specific trajectory sample sizes based on the number of cells susceptible to the hazards, but this is unlikely to yield a great reduction in run time since the hazards that take the longest are those that have the most host cells (not the ones for which sample size might be reduced). Regardless, once completed, this process does not need to be repeated unless hazard/host attributes affecting arrival and spread dynamics are altered, in which case trajectories need only be reconstructed for the affected hazards.

Sampling and combining trajectories from the database to form the final replicate simulations represents an additional phase of computation. We generated 10,000 such simulations using Project 170713's system on and system off arrival rates. These took around 2.7 hours (92 simulations/core/hour) and 6.8 hours (36/core/hour), respectively, using 40 cores. Thus, together with the time taken for generating the trajectory database, simulations comparable to those produced for Project 170713 were finished in around 70 hours, compared to the original 150 hours (excluding spatial data preparation). As mentioned above, simulating outcomes for further establishment rate scenarios (e.g. corresponding to altered controls) does not incur the cost of preparing the trajectory database.

Importantly, damage estimates arising from the new trajectory database approach described above were congruent with the original estimates from Project 170713, with median discounted damage at 50 years for the system on scenario approximately 0.6% higher than the original estimate. Thus while Project 170713 produced 50,000 simulations per scenario, 10,000 is likely to be sufficient using this approach. Distributions of damage over time for each asset were also generally equivalent to those derived from the original implementation, though there were very slight deviations for Forestry and Tourism (Figure H.2).

The above discrepancy is likely explained by increased jump frequency when using the trajectory database approach. The behaviour simulated by the trajectory database approach is equivalent in all aspects other than jump diffusion, which is permitted to occur more frequently than originally intended. In the original implementation, jump cells are only active until the first time a jump occurs from them. With the new implementation, jump cells are still deactivated after their first jump, but this is now per trajectory, not per simulation. This means that when combining trajectories into a simulation, if multiple sampled trajectories spread to the same jump cell, it is allowed to jump in each of those trajectories. The overlapping infection is handled by calculating the union of infected cells at each time step, but no such post-processing correction is possible to factor out overly frequent jumps. While the impact of this mismatch in logic appears minimal, it could be remedied by reducing the average jump rate.

Table H.1.: Hazard characteristics and number of trajectories generated per core per hour. *Host cells* gives the number of cells susceptible to the hazard; *Dispersal radius (m)* is the distance in metres that the hazard can spread per time step; *Pr(jump)* gives the mean of the Poisson distribution for the number of jump events per jump cell; *Time (min)* gives an indicative time taken to generate 50,000 arrival-spread trajectories assuming 44 cores are in dedicated to the task; and *n/core/hour* gives the corresponding number of trajectories generated per core per hour. Note that for FMD, spread upon arrival is deterministic since it is treated as having infinite dispersal capacity (i.e., all cells are "infected" upon arrival).

Group	Exemplar	Host cells	Dispersal radius (m)	Pr(jump)	Time (min)	n/core/hour
FSMC	Spongy moth	26,220	13,625	3.000	29.9	2,282
Animal other bacteria	Contagious equine metritis	40,201	6,922	0.150	54.0	1,262
Animal other micro other	Equine babesiosis	40,201	6,388	0.150	53.6	1,271
Animal other virus	Equine influenza V	40,201	12,676	0.150	49.6	1,374
Avian virus	hpai rev	81	3,666	25.000	7.1	9,626
Broadacre bacteria	Bacterial blight (angular leaf spot)	75,022	8,876	0.010	27.5	2,480
Broadacre beetle	Large grain borer	75,022	6,816	7.000	47.6	1,432
Broadacre bug thrips mite	Russian wheat aphid	75,022	8,436	11.000	42.3	1,612
Broadacre fungus	wheat stem rust	75,022	8,782	12.000	42.6	1,602
Broadacre mollusc	Golden apple snail	75,022	2,500	0.001	11.7	5,814
Broadacre virus	Cotton leaf curl virus	75,022	7,857	0.010	27.0	2,528
Broadacre weed	Red witchweed	664,257	9,710	0.004	66.5	1,025
FMD	FMD	625,352	Infinite	0.000	0.0	26,223,776
Forestry beetle	alb	21,411	3,053	6.000	39.5	1,728
Forestry fungus	ppc	25,576	12,319	3.000	33.5	2,035
Forestry nematode	Pine wilt nematode	25,576	6,112	2.000	31.0	2,200
Forestry termite	Termites	25,576	2,500	0.050	19.3	3,539
Forestry weed	*Amorpha fruticosa*	165,598	2,500	0.010	19.5	3,491
Fruitfly	Papaya Fruit Fly	768	10,000	5.000	8.1	8,369
GAS	gas	768	10,673	5.000	8.0	8,497
Horticulture bacteria	ccanker	768	2,500	3.000	8.9	7,643
Horticulture beetle	Colorado potato beetle	768	2,500	1.000	9.2	7,419
Horticulture bug thrips mite	thrips	768	3,189	2.500	9.0	7,568
Horticulture fly moth	False codling moth	768	2,500	0.200	10.7	6,358
Horticulture fungus	Citrus powdery mildew	768	2,500	2.000	9.6	7,103
Horticulture nematode	Potato cyst nematode	768	2,500	1.500	9.4	7,288
Horticulture virus	tbrn	768	2,500	2.000	9.9	6,914
Horticulture weed	Generic Cyperus	768	2,500	0.500	9.6	7,087
Khapra beetle	khapra	75,022	22,602	25.000	44.6	1,529
Livestock bacteria	Haemorrhagic septicaemia	625,352	13,016	2.200	236.4	288
Livestock bug thrips mite	Varroa	324,118	11,000	100.000	296.5	230
Livestock fly moth	Screw worm fly	625,352	24,000	5.000	606.2	112
Livestock virus	Bluetongue rev	625,352	24,000	5.000	605.3	113
Nonagricultural bee wasp	Generic Hymenoptera	324,118	2,500	0.050	40.4	1,687
Nonagricultural fly moth	Generic Diptera	4,692	3,000	0.200	15.6	4,370
Nonagricultural fungus	Dutch Elm Disease	1,096,348	10,000	4.500	367.1	186
Nonagricultural micro other	Dutch Elm Disease	1,096,348	10,000	4.500	366.2	186
Nonagricultural vertebrate	Black Spined Toad	1,092,183	10,000	0.010	118.3	576
Nonagricultural weed	MFG	1,092,183	2,500	0.001	50.2	1,358
Tramp ant	RIFA	1,213,717	15,000	0.020	156.8	435

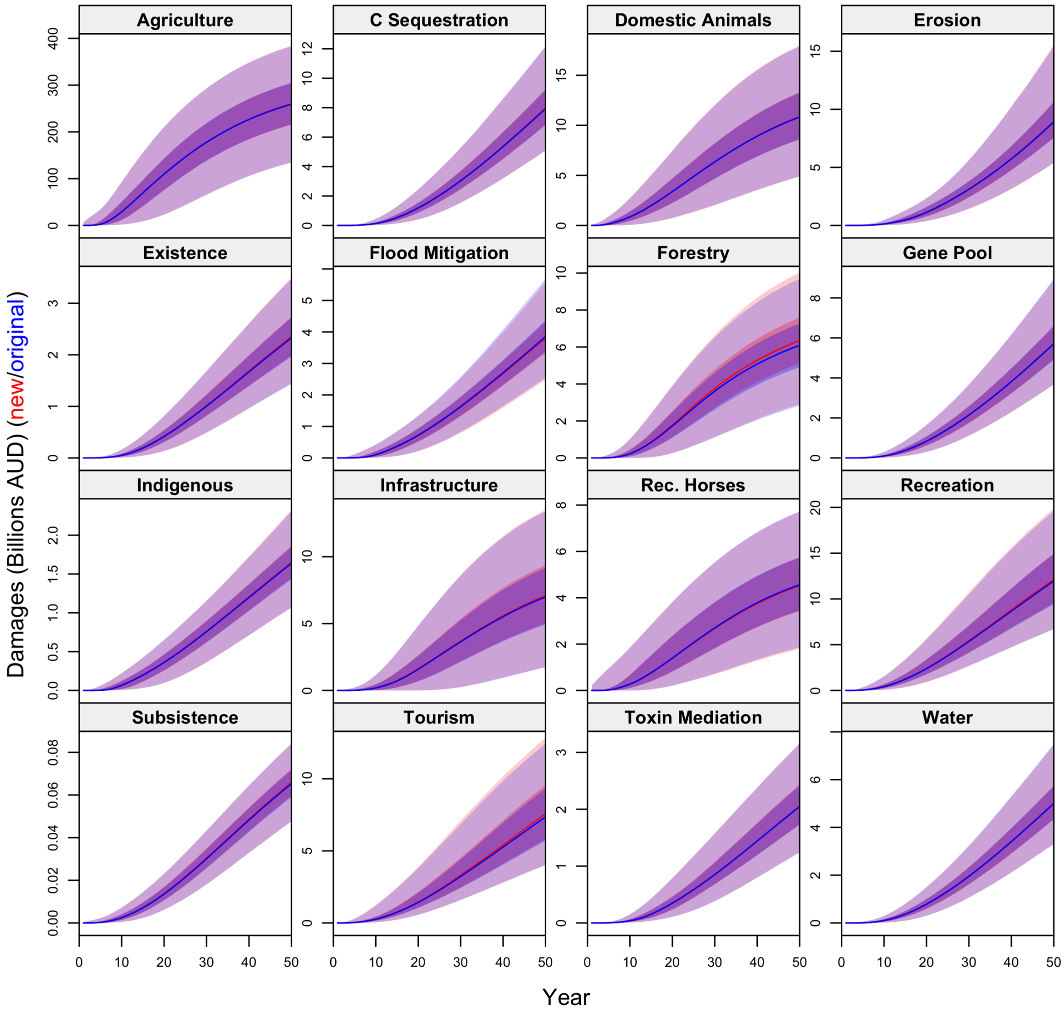

Figure H.2.: Comparison of the distribution of discounted, cumulative damage with the biosecurity system on (Dodd *et al.*, 2020), using the original Value Model simulation implementation (blue) and the new trajectory database implementation (red). Colours are semi-transparent, hence purple indicates areas of congruence. The solid line gives the median across iterations (n=50,000 for original approach; 10,000 for trajectory database approach). Ribbons indicate the 25th and 75th percentiles across iterations (dark ribbon), and the 2.5th and 97.5th percentiles (light ribbon).

H.2. `targets` workflow

The Value Model workflow has now been migrated from a series of lengthy R scripts to `targets` (Landau, 2021), a Make-like pipeline framework that tracks interdependencies of R objects and recreates any that become outdated, and enhances parallelisation of computation.

H.3. Replacing deprecated dependencies

The Value Model made frequent use of the `raster` package (Hijmans, 2023a) for manipulation of raster datasets. This has now been superseded by the considerably faster and more powerful `terra` package (Hijmans, 2023b). Alongside a range of other optimisations to the `value` package (`https://gitlab.unimelb.edu.au/cebra/cebra-r-packages/values`), this has led to a substantial reduction in the time taken to prepare raw spatial datasets – simulation inputs, for all of Australia at 2500 × 2500 m, are now generated from initial raw datasets in approximately 1.4 hours (44 cores), where previously this process took over 24 hours.