Risk Analysis, Vol. 38, No. 4, 2018 DOI: 10.1111/risa.12873

When Does Poor Governance Presage Biosecurity Risk?

Stephen E. Lane ,[1,*] Anthony D. Arthur,[2] Christina Aston,[2] Sam Zhao,[2] and Andrew P. Robinson[2]

Border inspection, and the challenge of deciding which of the tens of millions of consignments that arrive should be inspected, is a perennial problem for regulatory authorities. The objective of these inspections is to minimize the risk of contraband entering the country. As an example, for regulatory authorities in charge of biosecurity material, consignments of goods are classified before arrival according to their economic tariff number. This classification, perhaps along with other information, is used as a screening step to determine whether further biosecurity intervention, such as inspection, is necessary. Other information associated with consignments includes details such as the country of origin, supplier, and importer, for example. The choice of which consignments to inspect has typically been informed by historical records of intercepted material. Fortunately for regulators, interception is a rare event; however, this sparsity undermines the utility of historical records for deciding which containers to inspect. In this article, we report on an analysis that uses more detailed information to inform inspection. Using quarantine biosecurity as a case study, we create statistical profiles using generalized linear mixed models and compare different model specifications with historical information alone, demonstrating the utility of a statistical modeling approach. We also demonstrate some graphical model summaries that provide managers with insight into pathway governance.

KEY WORDS: Biosecurity; border inspection; prediction

1. INTRODUCTION

Efficient and effective border biosecurity strategies are essential for protecting ecosystems and economies from invasive pests. The annual cost of invasive species generally is estimated to be over USD$200bn[(1)] in the United States, and at least USD$4bn in Australia.[(2)] In Australia, the Department of Agriculture and Water Resources (the department) is both the regulatory authority and the inspectorate for biosecurity protection, carrying out both preborder and border intervention on a range of imported goods, based on the risk profile of the goods and international agreements. The objective of these interventions is to minimize the risk of biosecurity risk material (BRM) entering the country.

Here, we focus on border inspection and the challenge of deciding which of the tens of millions of consignments that arrive should be inspected. Before arrival, consignments of goods are classified according to their economic tariff number,[(3)] and this classification is used, with other information, as a screening step to determine whether further biosecurity intervention, such as inspection, is necessary. Other information associated with consignments includes details such as the country of origin, supplier, and importer, for example.

Border inspection for quarantine biosecurity is carried out for a number of reasons, namely, (i) to

[1]Centre of Excellence for Biosecurity Risk Analysis, University of Melbourne, Parkville, Victoria, Australia.

[2]Department of Agriculture and Water Resources, Canberra, Australian Capital Territory Australia.

[*]Address correspondence to Stephen E. Lane, Centre of Excellence for Biosecurity Risk Analysis, University of Melbourne, Parkville, Victoria 3010, Australia; lane.s@unimelb.edu.au.

0272-4332/18/0100-0653$22.00/1

verify the effectiveness of mandated prearrival treatments; (ii) to detect and intercept BRM; (iii) to provide information about the intrinsic contamination rate of the activity; and (iv) to deter potential malefactors. As noted above, tens of millions of consignments arrive every year, so the challenge is to determine which should be inspected.

We define a *pathway* as a collection of activities that culminates in the arrival to Australia of a set of alike consignments. The pathways are hierarchical, so we may consider a pathway of all consignments of a commodity, or all consignments of that commodity for a specific country, or even for a specific supplier. For example, the plant product pathway, which is the focus of this article, includes goods such as kiwi fruit and cashew nuts, which can themselves be considered pathways. Statistically, pathways can be thought of as processes.

Pathways can be classified as either high risk or a low risk, based on the probability that a consignment contains BRM, called the *approach rate*. For example, in Australia, kiwi fruit is a high-risk plant product pathway, with an approach rate of 55.8%, whereas cashew nuts is a low-risk plant product pathway, with an approach rate of 1.3%. Importantly, the degree of severity of the detected BRM in cashews has been identified as very low. The risk severity classification is important because the department may apply different interventions to low-risk than to high-risk pathways, as discussed. The identification of pathways as high or low risk is called *profiling*, and is an essential step in the efficient management of biosecurity intervention.

Traditionally, profiling has been applied by using records of interception of regulated pests on the pathway. This application is based on the assumption that future biosecurity compliance can be predicted by past biosecurity compliance, at least for some periods in the past and the future. However, interception of regulated pests is a rare event, which is good news from the point of view of biosecurity protection, but makes profiling more difficult, especially in sparse pathways, because reliable estimates of pathway risk are hard to obtain. This observation motivates the following question: whether future biosecurity compliance can be predicted by other characteristics as well as by past biosecurity compliance.

Historically, all consignments of imported plant product pathways were subjected to mandatory inspection. As part of a comprehensive review of Australia's biosecurity system,[4] the authors recommended establishing a science-based system for managing biosecurity issues, noting that zero risk is both unattainable and undesirable. With the full inspection strategy, pathways that have lower approach rate cost considerably more inspection effort to intercept BRM. For example, 4,623 consignments were inspected in the cashew pathway over four years, of which BRM was detected in 59, so the average number of inspections per detection (IPD) was about 78, compared to about 2 for the kiwi fruit pathway.

We now introduce the case study that motivates the research. The inspection work flow of imported plant product pathways comprises three components, namely, suppliers that export plant products, importers that import the products from suppliers, and border inspections that attempt to detect as much as possible BRM. Inspections at the border can be stratified by supplier or importer, that is, unique inspection regimes may be applied to individual importers or suppliers.

The department currently uses the continuous sampling plan (CSP) algorithm, specifically, CSP-3, to manage the biosecurity risk of low-risk pathways.[5,6] The CSP family of algorithms allocates intervention effort within pathways according to recent inspection history. The department has implemented CSP-3 for the inspection of a range of low-risk pathways, including dried apricots, green coffee beans, raisins, cashews, and some nuts. This particular approach to profiling has been shown to result in reductions of both leakage (how much BRM is missed in the inspection process) and IPD relative to random sampling plans.[7–9]

A wrinkle in the application of the CSP algorithm is that although it is implicit that the analysis of inspection history would take account of only the kinds of contamination that are of specific regulatory interest, in fact, any aspect of the inspection history can be used as an indicator of future risk. That is, although the department may be specifically concerned about intercepting regulated pests, the inspection history provides a much richer view of the pathway because it includes information about other incidents, such as the interception of nonregulated pests, failures of documentation, and so on, which may arguably and testably be related to the chances of failure types that are of regulatory concern. The question that motivated this study is: What data provide the most useful information about the pathway — the relatively sparse history of interception of regulated pests, or the more complete picture of the relative performance on the pathway, or some combination? Furthermore, can insight into the future performance in a given pathway be provided by information about historical performance in other, possibly related pathways?

This article reports an analysis of the use of auxiliary information to try to improve profiling. The objective is to distinguish high-risk and low-risk pathways, where risk refers to the interception rate of regulated pests, based on a range of characteristics of the pathway, including the interception rate of regulated pests, nonregulated pests, administrative failures, and supplier and tariff information. We aim to form a picture of the governance of the pathway and use that picture as a basis for assessing the relative biosecurity risk. The balance of the article is organized as follows. In the next section we introduce the data set and the models used to test our conjectures. We then present the results and a discussion and conclusion.

2. MATERIALS AND METHODS

We used a number of analytical approaches to assess the conjecture. First, we tested the association between nonregulated pest and administrative (more frequent, low-severity) failures and regulated pest inspection (less frequent, high-severity) failures. If such an association was found, then we reasoned that historical governance-related variables may be used to predict future high-severity biosecurity failures. Second, we used historical failure rates to create profiles, and investigated performance using receiver operating characteristic (ROC) curves. Finally, we constructed statistical models that would predict future regulated pest interception probabilities as a function of previous regulated pest interception probabilities and other, governance-related predictor variables.

All data preparation and modeling were performed using R Version 3.3.0[10] with the generalized additive mixed models of Section 2.4 using R package `rstanarm`.[11]

2.1. Data

The data for the analysis comprise the inspection history for all consignments classified as fruit Chapter 8,[3] that arrived between January 2007 and December 2011, a period of five years. The pathway is a complex one, comprising 80 different tariff codes, 3,150 unique importers, and 3,655 unique suppliers from 127 countries. For the purposes of this study we will assume that all significant biosecurity contamination has been captured by the regulatory border inspection. There were approximately 48,300 inspections of more than 75,000 goods. Approximately 5,300 inspections resulted in interception of a regulated pest, 8,500 inspections resulted in interception of a nonregulated pest, and 5,900 inspections recorded some administrative failure.

For modeling (see Section 2.4), we aggregated the data by year, tariff, and supplier. This aggregation was done for two reasons: first, it allowed us to create models that account for both supplier and tariff effects, and second, aggregating by year limits the effects of any seasonality. We use *interceptions* to refer to both interceptions of pests and administrative failures throughout the study. An appropriately formatted data set for modeling was constructed as follows.

For each year y within 2008 to 2011:

- Compute interception/fail rates for year $y-1$ by tariff, supplier, and year for:
  - (i) administrative interceptions,
  - (ii) nonregulated pest interceptions,
  - (iii) regulated pest interceptions.

We denote by X_{sty} the number of interceptions out of n_{sty} inspections from tariff t performed in year y from supplier s. Correspondingly, each inspection has a probability p_{sty} of being intercepted in one of the ways listed above. Then X_{sty} was modeled as:

$$X_{sty} \stackrel{d}{=} \text{Binomial}(p_{sty}, n_{sty}).$$

Computing interception rates by tariff, supplier, and year sometimes resulted in very small binomial denominators, due to the sparse history of inspection and interceptions produced. For this reason, rather than *raw* interception rates, we calculated *smoothed* interception rates using parametric empirical Bayes.[12] In particular, we used the Beta-binomial model to smooth interception rates for suppliers within tariffs and years; we provide the full details in Appendix A.

2.2. Association between Low-Severity and High-Severity Interceptions

To investigate the association between low-severity and high-severity interceptions, we calculated odds ratios and 95% confidence intervals (using a normal approximation for the log-odds) for the odds of a regulated pest interception for consignments with or without nonregulated pest or administrative interceptions.

For each inspected consignment, suppose Y denotes the outcome of inspection, so that $Y = 1$ indicates a regulated pest was intercepted. Further, let X denote whether the consignment contained a nonregulated pest (or had an administrative failure), so that $X = 1$ indicates the consignment contains a nonregulated pest (or had an administrative failure). The odds ratio is:

$$\mathrm{OR} = \left[\frac{\Pr(Y=1|X=1)/\Pr(Y=0|X=1)}{\Pr(Y=1|X=0)/\Pr(Y=0|X=0)}\right].$$

2.3. Profiling Using Annual Inspection Data

We created profiles using annual inspection data, and compared performance using ROC curves. For each year y in 2007 to 2010 and each kind of interception rate (regulated pest, nonregulated pest, and administrative) we compute ROC curves against year $y+1$ biosecurity inspection outcomes for regulated pests and, further, calculate the area under the curve (AUC).

We also computed ROC curves within tariff, due to the suspicion that the tariff-to-tariff variation would dominate the ROC signal, due to the differences of interception rates between the tariffs, rather than for the importers within the tariffs. That is, if we only ran the profiles across the tariffs then a naive assessment of the performance would look very good because we would expect the differences between the risks of the tariffs to be reasonably stable from year to year. Hence, assessing the model within tariffs provides a more reasonable assessment.

2.4. Profiling Using Statistical Modeling

We chose to construct models using generalized additive mixed model formulations with the linear predictors for the logit probability, $\log(\frac{p_{sty}}{1-p_{sty}})$ specified as:

$$\begin{aligned}
\text{Base}: &\ \beta_0 + \gamma_s + \tau_t\\
\text{M1}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y\\
\text{M2}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y + \varphi_{st}\\
\text{M3}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y + \varphi_{st} + \kappa_{sy}\\
\text{M4}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y + \varphi_{st} + \kappa_{sy} + b_3(p_{R,st(y-1)})\\
\text{M5}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y + \varphi_{st} + \kappa_{sy} + b_3(p_{N,st(y-1)})\\
\text{M6}: &\ \beta_0 + \gamma_s + \tau_t + \alpha_y + \varphi_{st} + \kappa_{sy} + b_3(p_{A,st(y-1)}),
\end{aligned}$$

where β_0 is a fixed process constant to be estimated; γ_s is a supplier-level effect; τ_t is a tariff-level effect; α_y is a effect for year of interception; φ_{st} is an effect for the supplier-tariff cross-classification; κ_{sy} is an effect for the supplier-year cross-classification; $b_3(\cdot)$ represent cubic regression splines for the previous year's regulated pest interception rate $p_{R,st(y-1)}$, nonregulated pest interception rate $p_{N,st(y-1)}$, and administrative interception rate $p_{A,st(y-1)}$. Bayesian logistic regression models were fit using `rstanarm`.[11] We used student-t priors for all coefficients, setting the scale for the intercept prior at 10, and for all other coefficients at 2.5.

To be more descriptive, M4 tests whether historical regulated pest interception rates can be used to predict future regulated pest interception probability, while M5 and M6 test effect of historical nonregulated pest and administrative interception rates on probability of future regulated pest interception.

Comparison of the statistical profiling results was made via a combination of: LOOIC comparisons,[13] and predictive log-likelihood via repeated five-fold cross-validation. LOOIC is similar to AIC in that it estimates out-of-sample prediction accuracy; however, LOOIC integrates over uncertainty in the parameters, and does not assume multivariate normality as the AIC does. We used 20 repeats, resulting in 100 training/testing data sets for comparison. To ensure balance across the data sets, sampling was performed within years. All models from Section 2.4 were fit to each training data set, and predictions made on the testing data sets.

3. RESULTS

We present the results in three sections: the association between low-severity and high-severity interceptions; the operational AUC tests; and the statistical modeling results. We finish this section with an in-depth look at the information gained from the modeling procedures.

3.1. Association between Low-Severity and High-Severity Interceptions

Fig. 1 shows the odds ratios, along with 95% confidence intervals, for the association between regulated pest (high-risk) interceptions and nonregulated pest and administrative (low-risk) interceptions both overall and by year. All estimates and lower bounds

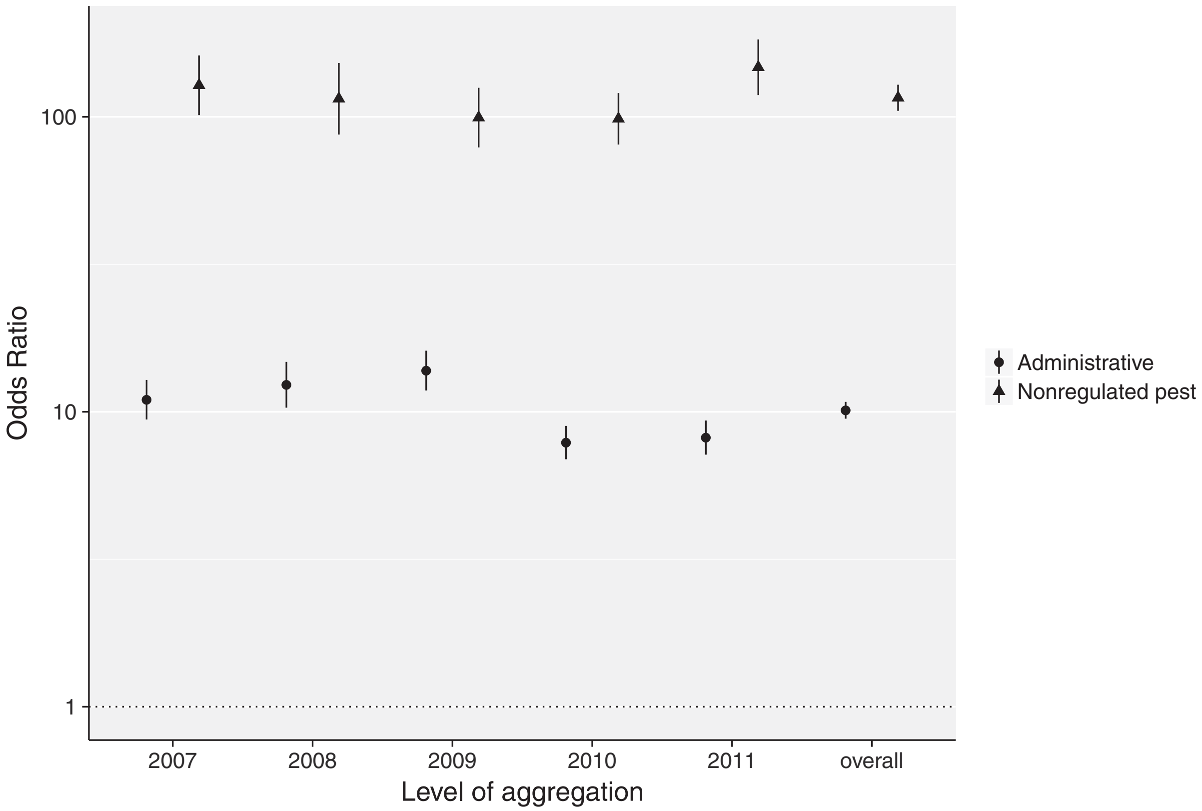

Fig. 1. Estimates (95% confidence intervals) of the odds ratios between low- and high-risk interceptions overall and by year (2007–2011). The odds ratios are calculated between regulated pest (high-risk) interceptions, and nonregulated pest and administrative (low-risk) interceptions.

of the confidence intervals are well above 1, showing there is a large association between low- and high-risk interceptions.

3.2. Comparison of Profiles Using Annual Inspection Data

3.2.1. Across Tariffs

Fig. 2 presents ROC curves that compare how well the different profiles perform. As per Section 2.3, the profiles are generated from the previous year's interception rates. All profiling approaches are substantially better than random, and the administrative profile consistently led to the weakest performance across each year. We have also shown the performance from a *combined* profile in Fig. 2; this is simply the profile using interception rates calculated from a variable indicating if *any* of the interception types occur. Clearly, the combined interception profile offers little performance over the regulated pest profile.

The profiles derived from nonregulated pest and administrative interception rates were consistently slightly worse than those based on regulated pest interception rates. Table I presents the AUC values for each of the curves presented in Fig. 2. The values are consistently close to 1, which suggests that the relative interception rates are very stable from year to year, and that the interception rates for each year y are a very good indicator for year $y + 1$. We also derived profiles using data without empirical Bayes smoothing (see Section 2.1); however, these profiles underperformed compared to the profiles using empirical Bayes smoothing. The results without empirical Bayes smoothing are shown in Appendix B, Table B1.

3.2.2. Within Tariffs

As noted in Section 2.3, our suspicion was that tariff-to-tariff variation would dominate the ROC signal, evidence for which was supported via

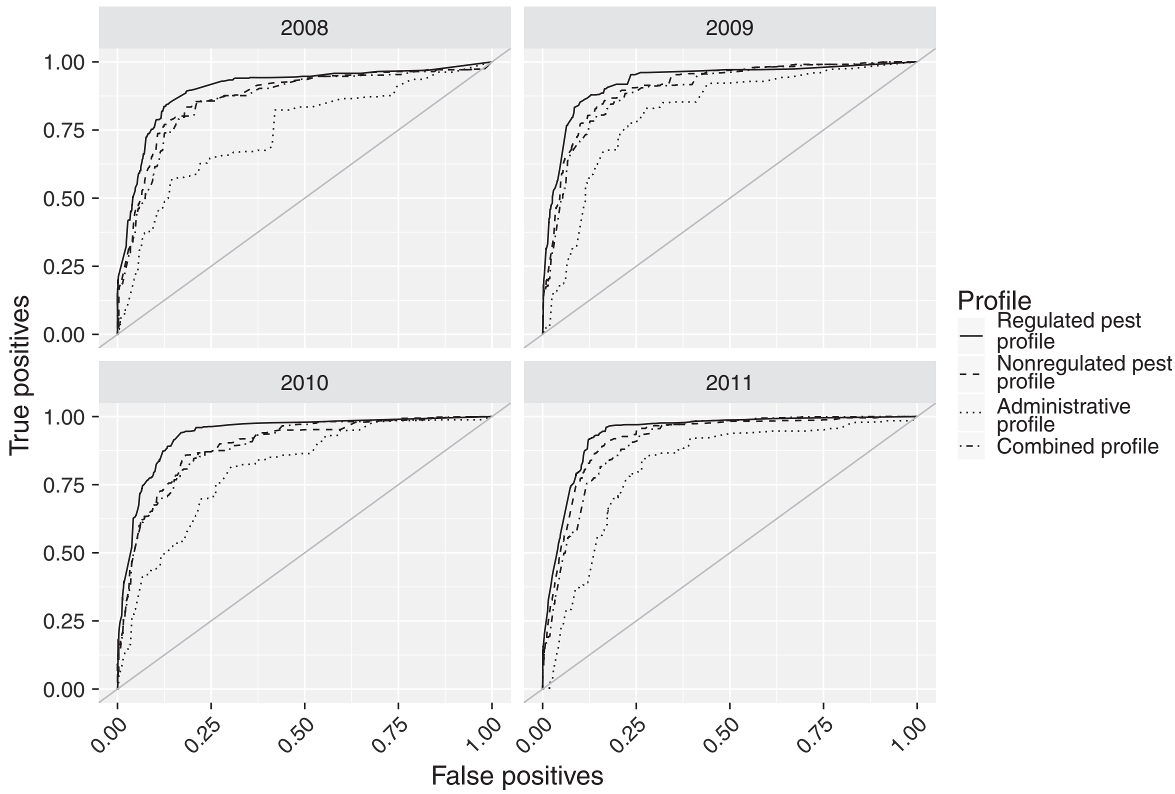

Fig. 2. ROC curves showing the performance of four profiling strategies for four years of data (2007–2011). The profiles are constructed by tariff and importer using the previous year's inspection data. A line is added at $x = y$ to facilitate comparison.

Table I. Summary of AUC Values for Profiling Strategies, by Year

Profile	2008	2009	2010	2011
Regulated pest	0.902	0.929	0.935	0.939
Nonregulated pest	0.870	0.909	0.892	0.920
Administrative	0.743	0.815	0.803	0.813
Combined	0.859	0.896	0.890	0.905

Note: The profiles are as follows: *Regulated pest* refers to using the previous year's regulated pest interception rate; *Nonregulated pest* refers to using the previous year's nonregulated pest interception rate; *Administrative* refers to using the previous year's administrative interception rate; and *Combined* refers to using the previous year's combined interception rate. Each AUC is computed using the data from the *following* year's inspections.

modeling (Table II). Fig. 3 plots the AUCs arising from the regulated pest profile versus the AUCs arising from the nonregulated pest, administrative, and combined profiles, respectively, within tariff. Each point represents an ROC curve applied to a single tariff, where the entities within the tariff that are being profiled are the suppliers. The size of each point indicates the number of regulated pest interceptions in the tariff, providing a sense of importance of that tariff.

A relationship between the number of regulated pest interceptions and AUC is not apparent in Fig. 3; we would expect larger points in the top-right corner if this were the case. However, within-tariff variation is considerable, providing a measure of conservatism against the strong performance of the across-tariff comparisons shown in Table I. This suggests that any profiling undertaken would need to take account of the tariff being profiled.

Correlation between the regulated pest profile and the other profiling strategies appears strong, especially between the nonregulated pest and combined profiles. The administrative profile results, however, show that many of the regulated pest AUCs lie above the $y = x$ line. This suggests that the administrative profiles are likely to perform worse than the nonregulated pest profiles. This observation is supported by the findings of the statistical modeling (Table II), where the model that included nonregulated pest interception rates performed better

Table II. LOOIC-Based Comparison of Statistical Profiling Models

Model	LOOIC	se(LOOIC)	Eff. P	se(Eff. P)	ΔLOOIC	se(ΔLOOIC)
M3	2, 788	111	458	23.3		
M4	2, 886	117	470	25.0	98.3	17.0
M2	2, 980	127	419	23.6	192.2	29.4
M1	3, 093	133	377	22.5	304.9	36.6
M5	3, 141	146	449	26.6	353.7	49.6
M6	3, 154	145	446	26.9	366.2	49.6
Base	3, 250	148	392	23.7	462.7	52.3

Note: The model with the smallest LOOIC (M3) is shown first, with subsequent rows ordered by increasing LOOIC. ΔLOOIC shows the difference in LOOIC between all models and Model M3; se(ΔLOOIC) shows the estimated standard error of the difference. Eff. P gives the estimated effective number of parameters; se(Eff. P) shows its standard error.

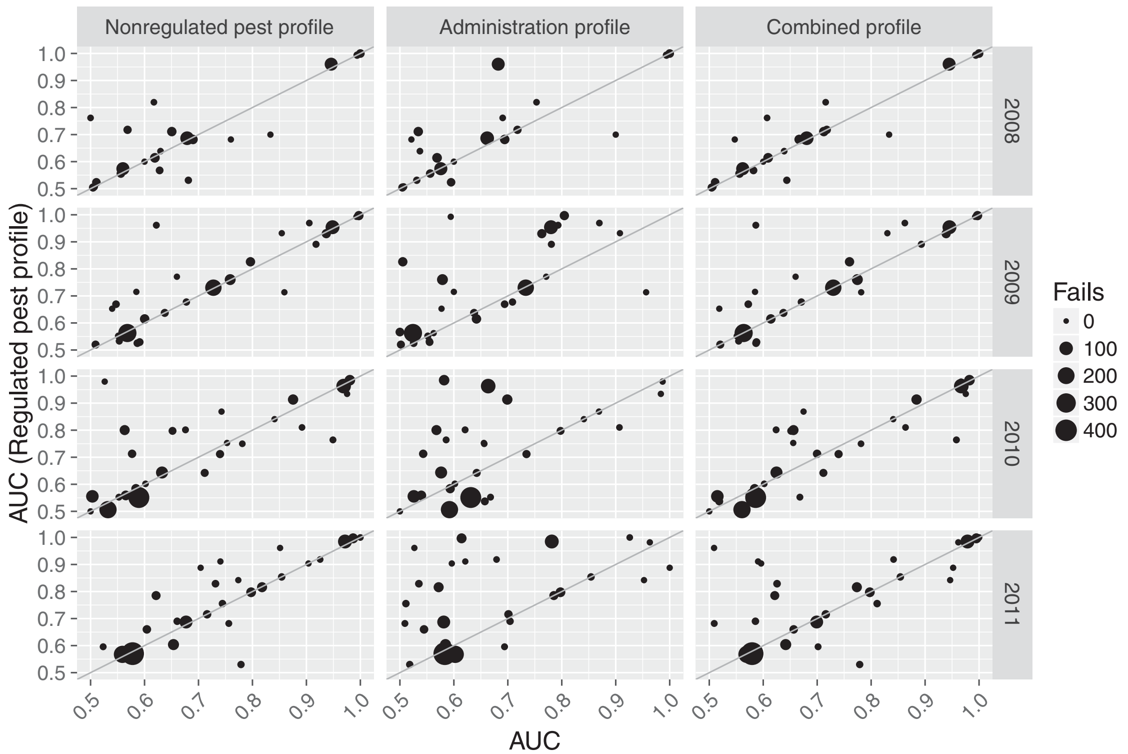

Fig. 3. AUCs computed for each tariff code in the data to assess within-tariff profiling operationally, by year. The y-axis is the AUC using the previous year's regulated pest interception rate to set the profile. The x-axis in each panel is the AUC using the previous year's nonregulated pest, administration, and combined interception rates to set the profile, assessed on the same inspection data that are used for the y-axis. The size of the point is related to the number of fails within the profile, and a line has been added at $x = y$ to facilitate comparison.

than the model including administrative interception rates.

3.3. Comparison of Profiles Using Statistical Modeling

Comparison of the models from Section 2.4 is reported in Table II. Model M3 has the lowest LOOIC, and the difference in LOOIC between M3 and M4 (98.3) is much larger than the standard error of its difference (17). These results show that supplier and tariff information are important for predicting regulated pest interception probability. The models are greatly improved with the addition of interaction terms between suppliers and tariffs, and suppliers and years. After allowing for the effects of suppliers,

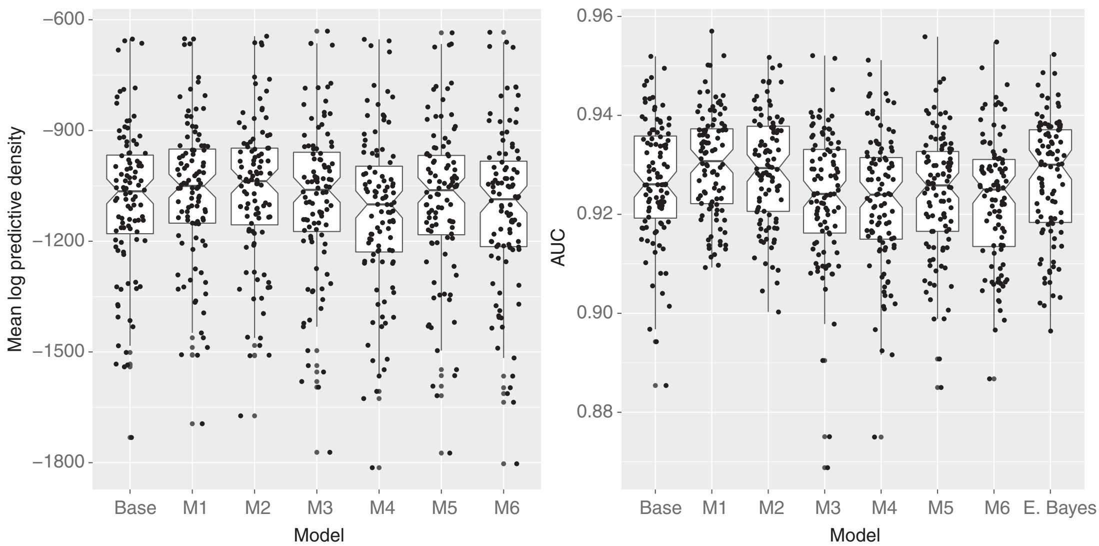

(a) Mean log predictive density (b) AUCs

Fig. 4. Mean log predictive density and AUCs for statistical profiling. Also shown are the AUCs from the regulated pest profile (E. Bayes).

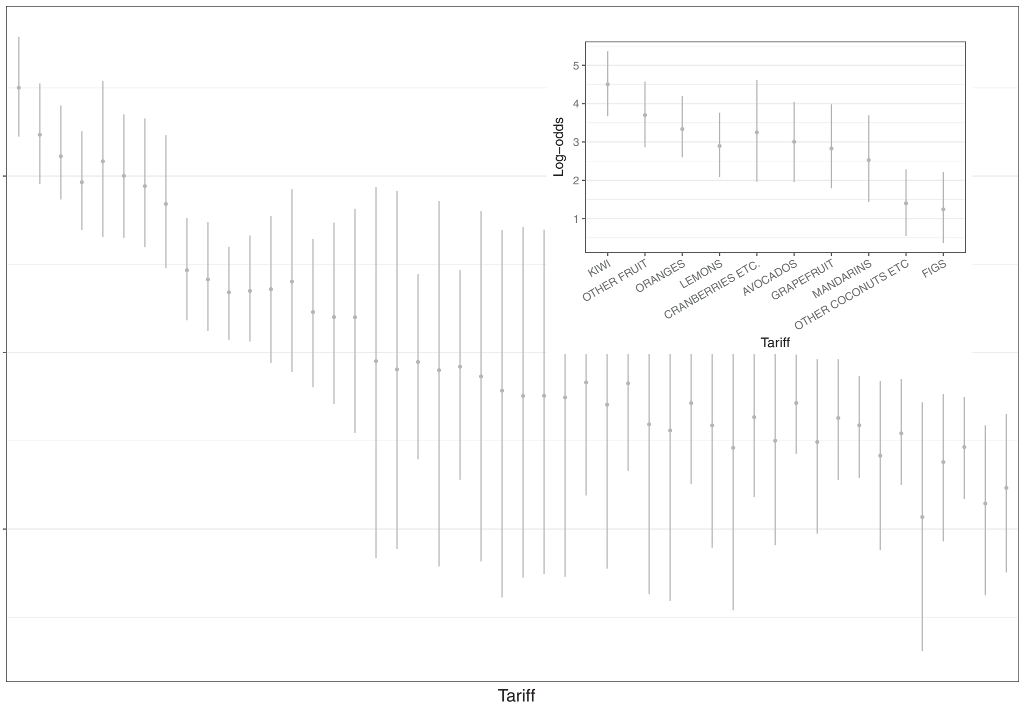

Fig. 5. Marginal odds of the tariff effect in Model 3. Tariffs are ordered left-to-right by decreasing probability of their marginal odds ratio being greater than 1, with bars showing 90% posterior credible intervals. The inset shows the top 10 tariffs.

tariffs, and years, the addition of the previous year's regulated pest interception rate (M4 vs. M3), the previous year's nonregulated pest interception rate (M5 vs. M3), and the previous year's administrative interception rates (M6 vs. M3) do not improve the model.

Fig. 4 shows the out-of-sample mean log predictive density and AUCs (Section 2.3) for all statistical profiling methods. Also shown in Fig. 4(b) are the AUCs from the regulated pest profile. Models Base–M3 perform the best in terms of predictive

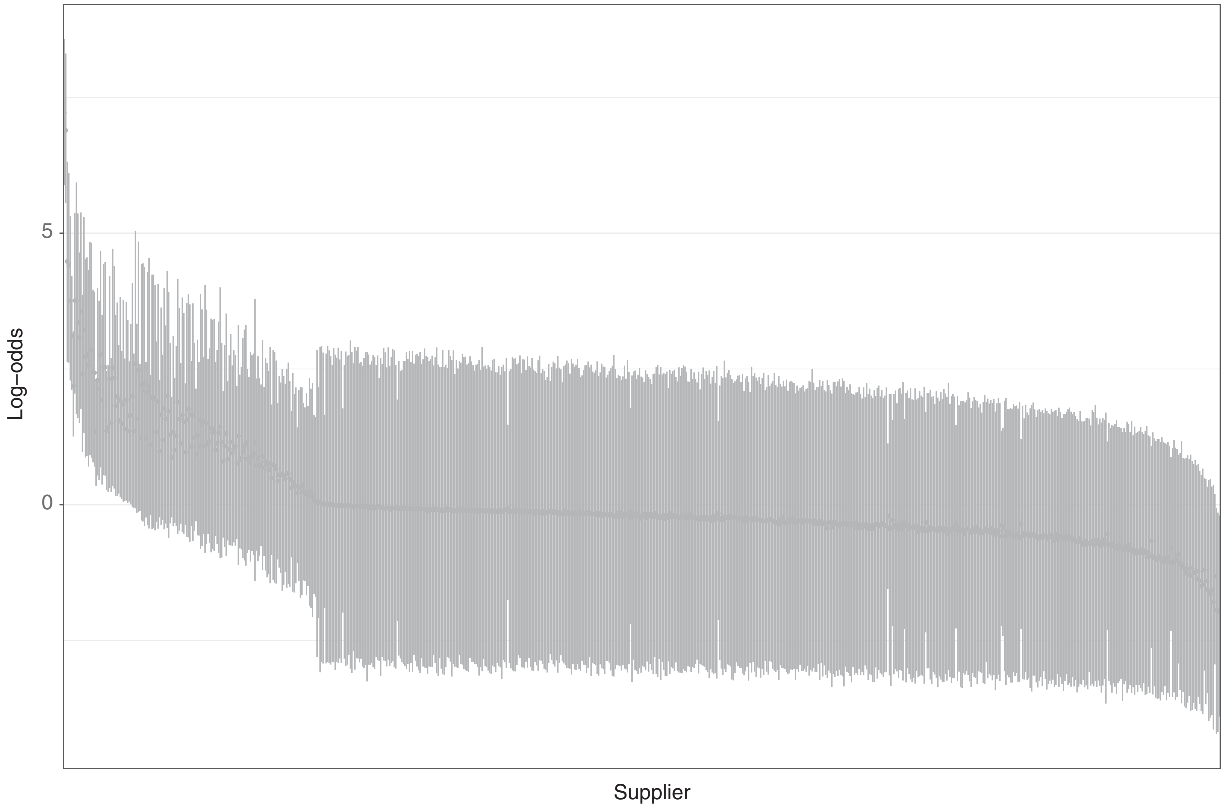

Fig. 6. Marginal odds of the supplier effect in Model M3. Suppliers are ordered left-to-right by decreasing posterior probability of their marginal odds ratio being greater than 1, with bars showing 90% posterior credible intervals.

log-likelihood (larger values are better), with no clear demarcation between them. In comparison, Models M1 and M2, as well as the Base model and the empirical Bayes profile, perform best on AUC.

3.4. Model Examination

In this section, we present an investigation of the effects from Model M3. We decided to investigate Model 3 further due to its superior performance in LOOIC (Table II), as well as the within-supplier examinations that would be available due to the interaction term. Fig. 5 shows the marginal odds ratio for tariffs from Model M3, ordered left-to-right by decreasing probability of their marginal odds ratio being greater than 1; bars in the figure show 90% posterior credible intervals. The inset shows the top 10 tariffs, and, as to be expected, Kiwi fruit is the tariff that contributes the highest risk.

Fig. 6 shows the marginal odds ratio for suppliers from the model, ordered left-to-right by decreasing posterior probability of their marginal odds ratio being greater than 1; bars in the figure show 90% posterior credible intervals. Supplier labels have been masked for privacy reasons. The suppliers to the left of the figure are those predicted to have a large increase in probability of regulated pest interception, all else being equal. It is these suppliers that would naturally be the first targets in an operational capacity.

Fig. 7 provides a closer examination of the risky suppliers. We have selected the top 25 suppliers (by the probability of their marginal odds ratio being greater than 1) and calculated their posterior probability of a regulated pest being present in a consignment, averaged over all years from Model M3. The panel on the left shows their posterior probability for each tariff that the supplier imports, whilst the panel on the right shows the observed proportion (averaged over years) of regulated pest interceptions by tariff. This figure shows that these highest risk suppliers import a range of tariffs—i.e., their poor performance is not necessarily due to importing one or two of the highest risk tariffs (as shown in Fig. 5).

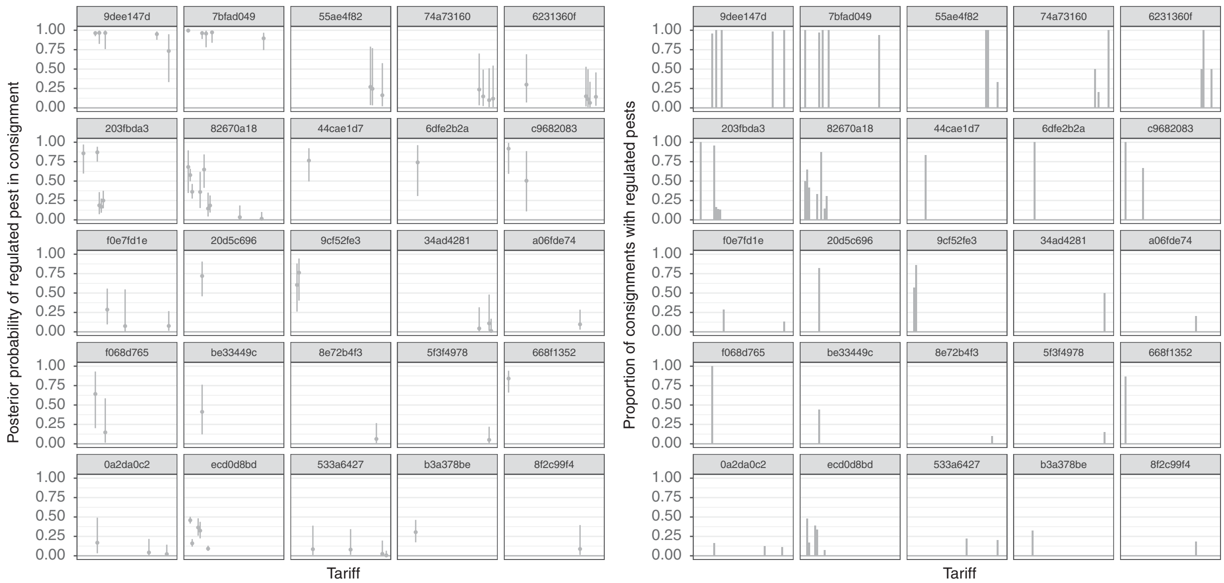

(a) Posterior probability of a regulated pest interception.

(b) Observed proportion of consignments with regulated pests.

Fig. 7. Posterior probability of a regulated pest interception in the top 25 suppliers, along with the observed proportion of consignments with regulated pests. Panels are ordered left-to-right, top-to-bottom by the probability of their marginal odds ratio being greater than 1; bars in the left panel show 90% posterior credible intervals.

Further, in the tariffs they do import, they have consistently high levels of consignments with regulated pest contamination (right panel, Fig. 7).

Fig. 8 provides a closer examination of suppliers who pose minimal risk. We have selected the bottom 25 suppliers (by the probability of their marginal odds ratio being greater than 1) and calculated their posterior probability of a regulated pest being present in a consignment, averaged over all years from Model M3. The panel on the left shows their posterior probability for each tariff that the supplier imports, while the panel on the right shows the observed proportion (averaged over years) of consignments that *did not* contain regulated pests by tariff. Similar to Fig. 7, these suppliers import a range of tariffs—i.e., their good performance is not necessarily due to importing lower risk tariffs. However, in comparison to the risky suppliers, in the tariffs they do import, they have consistently high levels of consignments without regulated pest contamination (right panel, Fig. 8).

4. DISCUSSION AND CONCLUSION

There was a strong association between regulated pest interceptions and the lower-risk administrative and nonregulated pest interceptions (Section 3.1). This association was also observed when using administrative interceptions as a predictor for operational profiling (Fig. 2), demonstrating the utility of the operational profiling approaches. However, we note that this does not carry over into the statistical models (Section 3.3), for which including historical rates as predictor variables did not improve model fits.

The statistical profiles still performed well using the cross-validated AUCs (Fig. 4) as well as the predictive log-likelihood. Thus, in answer to our motivating question of which data provide the most useful information about the pathway, we would conclude that it is the knowledge of particular suppliers, tariffs, and their combination that is most informative. The previous year's regulated pest profile performed well based on AUC; however, adding this to the statistical profiles gave no benefit (Table II). Furthermore, there is limited scope for investigating why a particular supplier may be problematic. Statistical profiling, in comparison, allows decisions to be based on posterior probabilities. For example, we could calculate a supplier's (marginal) probability of having a regulated pest interception. Intervention could then be planned on either the top ranked suppliers (if funds are limited), or all suppliers that meet a threshold.

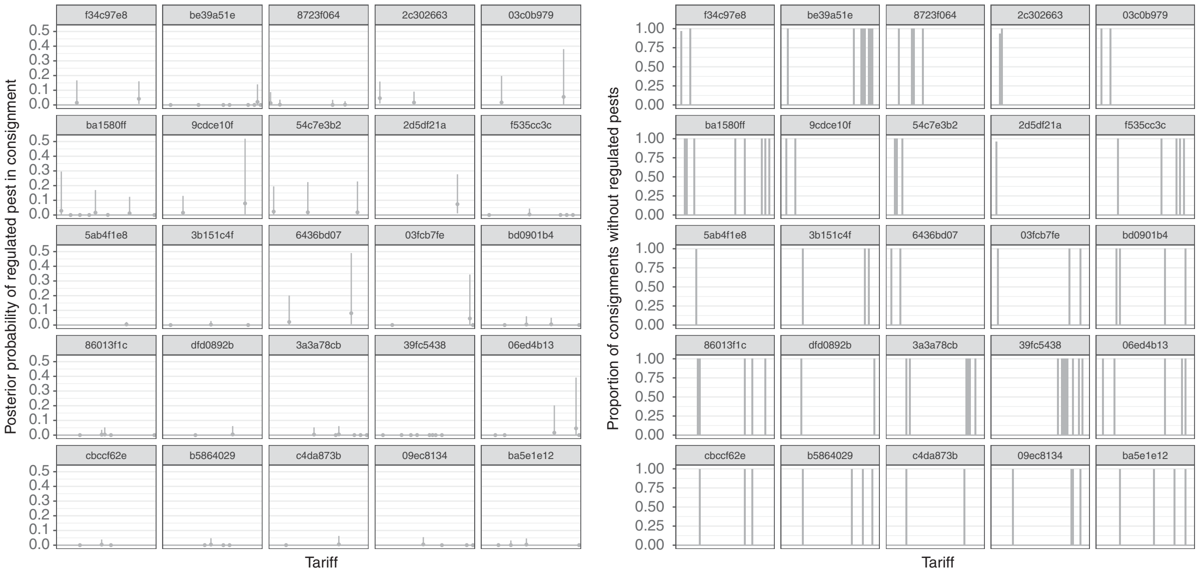

(a) Posterior probability of a regulated pest interception.

(b) Observed proportion of consignments without regulated pests.

Fig. 8. Posterior probability of a regulated pest interception in the bottom 25 suppliers, along with the observed proportion of consignments without regulated pests. Panels are ordered left-to-right, top-to-bottom by the probability of their marginal odds ratio being less than 1; bars in the left panel show 90% posterior credible intervals.

A benefit of using a statistical modeling approach to investigate profiling is the added level of interrogation possible from the fitted model. In Section 3.4, we demonstrated how we can gain a clearer insight into the governance of this pathway. First, by studying the marginal effects of suppliers and tariffs, we can build up a picture of risk without relying on observed rates, which are noisy due to sampling and process error. We can pinpoint which tariffs and suppliers contribute to excessive risk, essentially by an ordering of the marginal odds ratio, and then choose to investigate those that have a posterior probability higher than a predefined cutoff set by management.

With a list of potentially risky suppliers, we further demonstrated how a manager could gain information into the governance of those suppliers by investigating the posterior predicted probabilities of regulated pest interceptions (Fig. 7). This information could be used to initially examine why a particular supplier may be having trouble with contaminated consignments, and be used to help improve its processes. Similarly, looking at the less risky importers (Fig. 8) may provide information on good process that can be shared with the riskier importers.

Our use of interception records as a proxy for biosecurity risk increases the probability of false negatives, which are undetected risks. This effect may vary systematically across unmeasured variables. For example, the estimated interception rate of a particular supplier may be exaggerated if it supplies predominantly to a port that employs particularly effective inspectors. The approach suggested in this article could be expanded to include port effects in the statistical model; if there are particularly effective inspectors at a port, then this will be reflected in the model's port-level estimates. The outcomes of suppliers that provide predominantly to that port would then be ameliorated by the port effect. The data set we have used does not allow this comparison.

The effect of systematic inaccuracies in recording pest interceptions can be estimated. The positive predictive value (PPV) is defined as the proportion of recorded pest interceptions that contain pests; perfect inspectors would have a PPV equal to 1. PPV depends upon the true contamination rate and the effectiveness of the inspectors. Effectiveness is measured by the inspectors' *sensitivity* in detecting pests: the proportion of contaminated consignments that are correctly identified (the true positive rate), and *specificity*: the proportion of consignments not contaminated that are correctly identified (the true negative rate). Particularly effective inspectors will have high sensitivity, and high specificity. Mathematically,

$$\text{PPV} = \frac{\text{sensitivity} \times \text{contamination rate}}{\text{sensitivity} \times \text{contamination rate} + (1 - \text{specificity}) \times (1 - \text{contamination rate})}.$$

An example follows. We assume that the true contamination rate is 0.1, and the rate at which inspectors make false positives (i.e., the rate at which they record a pest when it is not there) is 0.05. We can compare the PPV for two different ports. Assume that the sensitivity of inspectors at Port A is 0.9, i.e., they are particularly effective, and the sensitivity of inspectors at Port B is 0.7, i.e., they are less effective; all else being equal, the PPV of the inspectors at Port A is 0.67, and at Port B is 0.61. In other words, if inspectors at Port B are 22% less effective than those at Port A, then the PPV is only affected by 8.7%. Consequently, large changes in effectiveness result in smaller changes in outcomes for suppliers.

To summarize, statistical models constructed using inspection data provide sufficient information for profiling purposes, and can be used to further interrogate the governance of multiple pathways and to help identify the processes underlying poor performance on these pathways.

ACKNOWLEDGMENTS

The authors are grateful to Lindsay Penrose (ABARES, Australian federal Department of Agriculture and Water Resources) for statistical analysis on an early draft not reported here and Brendan Woolcott (Plant Division, Australian federal Department of Agriculture and Water Resources) for providing the data for the analysis. The authors would also like to thank the anonymous reviewers for comments that improved the article. AR and SL's contributions were funded in part by the Centre of Excellence for Biosecurity Risk Analysis.

APPENDIX A: CALCULATION OF RATES USING EMPIRICAL BAYES

In this appendix, we detail the procedure used for calculating the smoothed rates in Section 2.1 via empirical Bayes. Recall that we have X_{sty} the number of failures out of n_{sty} inspections from tariff t performed in year y from supplier s; we assume that $X_{sty} \stackrel{d}{=} \text{Binomial}(p_{sty}, n_{sty})$.

To find the empirical Bayes estimate of p_{sty} for supplier s, in tariff t and year y, assume that the binomial proportions p_{sty} have a prior Beta distribution: $p_{sty} \sim \text{Beta}(\alpha_{ty}, \beta_{ty})$. Then X_{sty} has a Beta-binomial distribution, with probability mass function

$$\Pr\ (X_{sty} = k) = \binom{n_{sty}}{k} \frac{\Gamma\,(n_{sty}+1)}{\Gamma\,(k+1)\,\Gamma\,(n_{sty}-k+1)} \times \frac{\Gamma\,(k+\alpha_{ty})\,\Gamma\,(n_{sty}-k+\beta_{ty})}{\Gamma\,(n_{sty}+\alpha_{ty}+\beta_{ty})} \frac{\Gamma\,(\alpha_{ty}+\beta_{ty})}{\Gamma\,(\alpha_{ty})\,\Gamma\,(\beta_{ty})}.$$

The parameters α_{ty} and β_{ty} are found using maximum likelihood:

$$(\hat{\alpha}_{ty}, \hat{\beta}_{ty}) = \arg\max_{\alpha_{ty},\beta_{ty}} \left\{ -\sum_{s=1}^{S} \log \Pr\,(X_{sty} = x_{sty}) \right\},$$

where x_{sty} is the observed value of X_{sty}, and S is the number of suppliers. To complete the calculation, the rates for supplier s in tariff t and year y are updated using the following formula:

$$\tilde{p}_{sty} = \frac{x_{sty} + \hat{\alpha}_{ty}}{n_{sty} + \hat{\alpha}_{ty} + \hat{\beta}_{ty}}.$$

APPENDIX B: FURTHER RESULTS

Table B1. Summary of AUC Values for Profiling Strategies, by Year

Profile	2008	2009	2010	2011
Tariff and Supplier				
Regulated pest	0.881	0.899	0.902	0.917
Nonregulated pest	0.849	0.878	0.859	0.890
Administrative	0.728	0.767	0.759	0.776
Combined	0.833	0.861	0.854	0.867
Supplier within Tariff				
Regulated pest	0.861	0.897	0.903	0.906
Nonregulated pest	0.839	0.880	0.864	0.907
Administrative	0.795	0.841	0.824	0.842
Combined	0.836	0.881	0.871	0.901
Supplier within Tariff, Smoothed				
Regulated pest	0.902	0.929	0.935	0.939
Nonregulated pest	0.870	0.909	0.892	0.920
Administrative	0.743	0.815	0.803	0.813
Combined	0.859	0.896	0.890	0.905

Note: The profiles are as follows: Tariff and Supplier refers to profiles constructed for the interaction of tariff and supplier, Supplier within Tariff refers to averaging the supplier interception rates within tariffs, and Supplier within Tariff, Smoothed refers to using the empirical Bayes estimate of the suppliers within tariffs and years. *Regulated pest* refers to using the previous year's regulated pest interception rate; *Nonregulated pest* refers to using the previous year's nonregulated pest interception rate; *Administrative* refers to using the previous year's administrative interception rate; and *Combined* refers to using the previous year's combined interception rate. Each AUC is computed using the data from the *following* year's inspections.

REFERENCES

1. Pimentel D. Biological Invasions: Economic and Environmental Costs of Alien Plant, Animal, and Microbe Species, 2nd ed. Hoboken, NJ: CRC Press, 2011.
2. Sinden JA, Jones R, Hester S, Odom D, Kalisch C, James R, et al. The economic impact of weeds in Australia. Plant Protection Quarterly, 2005; 20(1):25–32. Available at: https://doi.org/e-publications.une.edu.au/1959.11/331
3. Department of Immigration and Border Protection. Customs Tariff Act 1995, 2016. Available at: https://doi.org/www.border.gov.au/Busi/Tari/Curr
4. Beale R, Fairbrother J, Inglis A, Trebeck D. One biosecurity: A working partnership. Commonwealth of Australia, 2008.
5. Dodge H. A sampling inspection plan for continuous production. Annals of Mathematical Statistics, 1943; 14(3):264–279.
6. Dodge HF, Torrey MN. Additional continuous sampling plans. Industrial Quality Control, 1951; 7(5):7–12.
7. Robinson A, Bell J, Woolcott B, Perotti E. DAFF Biosecurity: Plant-Product Pathways. Australian Centre of Excellence for Risk Analysis, 2012. 1001J 1.
8. Robinson A, Woolcott B, Holmes P, Dawes A, Sibley J, Porter L, et al. Plant quarantine inspection and auditing across the biosecurity continuum. Australian Centre of Excellence for Risk Analysis, 2013. 1101C 1.
9. Arthur T, Zhao S, Robinson A, Woolcott B, Perotti E, Aston C. Statistical modelling and risk return improvements for the plant quarantine pathway. Australian Centre of Excellence for Risk Analysis, 2013. 1206F 1.
10. R Core Team. R: A Language and Environment for Statistical Computing. Vienna, Austria, 2016. Available at: https://doi.org/www.r-project.org/
11. Gabry J, Goodrich B. rstanarm: Bayesian Applied Regression Modeling via Stan, 2016. R package version 2.12.1. Available at: https://doi.org/CRAN.R-project.org/package=rstanarm
12. Robinson AP, Chisholm M, Mudford R, Maillardet R. Ad hoc solutions to estimating pathway non-compliance rates using imperfect and incomplete information. Pp. 167–180 in Jarrad F, Low-Choy S, Mengersen K (eds). Biosecurity Surveillance: Quantitative Approaches. CABI Invasive Series. Boston: CABI, 2015.
13. Vehtari A, Gelman A, Gabry J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Statistics and Computing, 2016; 30:1–20. Available at: https://doi.org/link.springer.com/article/10.1007/s11222-016-9696-4