*Ecological Applications*, 21(4), 2011, pp. 1410–1417


# Allocating surveillance resources to reduce ecological invasions: maximizing detections and information about the threat

Andrew Robinson,[1,4] Mark A. Burgman,[2] and Rob Cannon[3]

[1]*Australian Centre of Excellence for Risk Analysis, Department of Mathematics and Statistics, University of Melbourne, Parkville, Victoria 3010 Australia*
[2]*Australian Centre of Excellence for Risk Analysis, School of Botany, University of Melbourne, Parkville, Victoria 3010 Australia*
[3]*Biosecurity Services Group, Department of Agriculture, Fisheries, and Forestry, Canberra, Australian Capital Territory 2601 Australia*

*Abstract*. Allocating resources to detect invasive pests, diseases, and pathogens on exposure pathways requires a trade-off between the need to detect as many contaminated items as possible and the need to acquire knowledge about contamination rates. We develop a model and an algorithm that provide guidance for the allocation of inspection resources across multiple dynamic pathways in cases where not every item can be inspected. The model uses a null hypothesis that the contamination rate of a pathway is above a specified level: a risk cutoff. Pathways with a risk above the cutoff are fully inspected, and those with a risk below the cutoff level are monitored at a rate that would detect a change of the risk to being above the cutoff level with high probability. We base our decision on the 95% upper confidence limit for the contamination rate. We demonstrate via simulations and a data set that focusing inspection resources on specific pathways can result in substantially more effective intervention, and that the reduction in overall effectiveness of monitoring low-risk pathways need not be substantial. Use of the model demands the selection of the risk cutoff, and this limit can be set according to projected consequences.

*Key words: contamination rates; detection resources; diseases; estimating rates; full inspection vs. sampling; pathogens; pests; quarantine; risk cutoff; species introductions; surveillance effectiveness.*

## Introduction

Nonindigenous pests, weeds, diseases, and pathogens invade marine and terrestrial environments, with significant ecological, economic, and human health consequences (Savidge 1987, OTA 1993, Pimentel et al. 2000, 2005, Hulme 2009). Invasive species have been shown to adversely affect native species, ecosystems, and ecosystem services (Xie et al. 2001, Gurevitch and Padilla 2004, Simberloff 2006, Langor and Sweeney 2009). In economic terms, Colautti et al. (2006) estimated a range for costs of between Canadian $13.3 and $34.5 billion annually for 16 invasive species in Canada. The impact of invasive species on the environment and agricultural productivity in the United States is estimated at US$138 billion annually (Pimentel et al. 2005, Pimentel 2007).

Many pests and diseases disperse into environments associated with commodities that are part of burgeoning international trade (Costello et al. 2007). Border-based interceptions play a critical role in the prevention of these invasions. For example, Haack (2001, 2006) reported 8341 beetle interceptions for the United States in a 15-year period (cf. Work et al. 2005, Brockerhoff et al. 2006).

Manuscript received 28 January 2010; revised 26 July 2010; accepted 23 August 2010. Corresponding Editor: M. P. Ayres.
[4] E-mail: A.Robinson@ms.unimelb.edu.au

Effective surveillance systems detect and eliminate as many threats as possible for a given budget. To develop effective strategies that avoid or mitigate the consequences of invasive species associated with trade, it is also necessary to have an understanding of the rate at which various pests contaminate the traded commodities (Kenis et al. 2007). Thus inspections at the border play two roles: they detect and eliminate potential environmental pests and they accumulate information about the rate of contamination of the units inspected. This paper focuses on the allocation of inspection resources to implement policy for acceptable levels of protection against invasive species and to learn about contamination rates.

Some previous analytical work has addressed related topics. Press (2009) proposed a method for allocating inspection resources for fixed, known prior probabilities of malfeasance. Cox (2009) recommended a portfolio approach to deal with dependencies between pathways. Ramirez-Marquez (2008) focused on preventing attacks through container cargo using a decision tree based on multiple sensors, for situations in which historical data were unavailable or unreliable. Surkov et al. (2008) integrated the costs of invasion and estimated the marginal benefits of border inspection. Sikder et al. (2006) and Moffitt et al. (2008) used rough sets and information gap, respectively, to offset the estimated costs of the effects of invasion, about which little is known, against the cost of inspections. Finally, Govindaraju et al. (2010) reviewed the partial inspection and skipping strategy that has been adopted by the New Zealand Food Safety Authority.

Here we address the simple and general problem of how best to allocate surveillance resources among pathways when historical information is available, when contamination rates may vary in time, and when we wish to satisfy the joint objectives of being as effective as possible and of learning about uncertain and changing numbers of invasive species in the different pathways. Although the techniques that we propose are illustrated in a quarantine setting, they are relevant for a wide array of problems, including customs and immigration, disease surveillance in human, animal, and plant populations, pest and disease eradication systems, and other contexts that rely on the deployment of limited surveillance resources among a set of statistical populations.

## The Model

Briefly, the surveillance problem that we address is as follows. An inspectorate identifies a suite of independent pathways, each comprising sampling items (e.g., containers, packets of seed, vehicles, individual animals). Pathways may refer to importation of different products via different transport routes, or from different trading partners, or combinations thereof (e.g., Costello et al. 2007). In general, a pathway is an aggregation or stream of like items that can be interpreted as a population for statistical purposes.

Some of the items in the pathway are expected to contain a potentially damaging, invasive species. Such items are referred to as "contaminated" items. The inspectorate must identify as many of these contaminated items as possible across all of the pathways. Identification of contaminated items is carried out by inspecting the items on the pathway. Inspection resources are finite. We define the "approach rate" as the number or rate of contaminated items arriving in a pathway, and "leakage" as the number or rate of contaminated items remaining in a pathway after inspection and intervention.

In general, when only a proportion of incoming goods can be inspected, efficiency (the number of detections per unit effort) will be enhanced by reducing inspection rates in areas where the risks are smallest and increasing inspection rates where the risks are highest. However, this approach is counterbalanced by the need to monitor the pathways for any change in the risk posed. The challenge is to determine what proportion of a pathway should be inspected. Should we inspect everything in a pathway? If we only sample a portion of the pathway, then what portion should be sampled? Or should the pathway be ignored?

We define "managing the risk" of the pathways, in a quarantine context, as knowing that the contamination rate is below a determined threshold, either by knowing that the pathway is clean (by monitoring) or by applying an intervention with statistically predictable effects. The threshold below which the contamination rate must be kept is determined with reference to the consequences of undetected contamination.

We now describe essential characteristics for the scenarios that are in the scope of this study. First, inspections are the only means by which the degree of contamination of a pathway can be estimated. Second, inspection can only be "done" or "not done." There can be occasions when partial inspection is an option; see the *Discussion* for further comments on such strategies. Third, the usefulness of information that has been collected by inspecting the pathways decreases over time, because the risk associated with pathways may change. Finally, the inspectorate operates with a flexible budget, so that the firm resource limitations associated with scenarios such as those reported in Cannon (2009) do not apply.

Given a fixed amount of inspection resources and exact information on the contamination rate for each pathway, it is known that the highest number of contaminated items will be intercepted by inspecting the pathways that have the highest contamination rate (see, e.g., Cannon 2009), using the following algorithm. Resources (for the year) would be allocated to inspect all of the items that will arrive on the pathway that has the highest contamination rate. It is assumed that we have a reasonable estimate of the number of items arriving along each pathway and that the system is sufficiently flexible to cope with any change. If sufficient unused resources remain, the pathway with the second highest contamination rate would be completely inspected. The procedure is repeated until resources are depleted. The lowest-rate pathways may not be inspected at all. This procedure gives rise to three classes of pathways: one for which all items are inspected, one for which no items are inspected, and another for which some items are inspected. In applications, the limitation on resources is usually flexible enough to ensure that the one pathway for which available resources run out is fully inspected.

Note that we can only use historical data to classify the pathways, but we perform the classification intending to maximize the future number of intercepted contaminated items. The information that is used for the classification has unknown value, and this value will decrease with the age of the information. This characteristic leads to the possibility that pathways with high contamination rates might be ignored by inspectors because either the historical contamination rate was low by chance, or the contamination rate of the pathway was previously low but has changed in time.

We now develop a statistical model for the inspection process. A number $N$ of items is expected on a pathway over some fixed time period. The unknown probability that an individual item will contain at least one contaminant is denoted $p$. We have historical data comprising $n$ inspections, from which $x$ contaminated items are observed. Although we may know exactly how many contaminated items did actually arrive during the year, we shall assume that the consignments were a random sample of size $n$ from the vast number of items exported from the source country. Hence, we shall assume that

$$x \sim \mathrm{Binomial}(n, p). \tag{1}$$

Subscripts will be used to denote different times of measurement as necessary.

We now introduce the idea of "predicted risk" for a pathway. This risk will be defined as the estimated probability of contamination of an item from the pathway, inflated to reflect uncertainty about the accuracy of the estimation. Our decisions on what sample sizes we shall use will be based on the predicted risk, denoted by $f$, rather than the estimate of $p$.

By using the concept of predicted risk, we are arguing for the need to consider both the estimate of the contamination probability $p$ and the accuracy of the estimate when allocating resources. Given two commodities with the same contamination rate but different inspection counts, we would be less certain about the estimate of the contamination rate for the commodity that has the lower inspection count; the standard error of the estimate would be higher. A risk-sensitive solution requires that we take account of the uncertainty of the estimate, effectively forcing us to pay for the higher uncertainty in the same way that we pay for the higher contamination rate, by allocating additional inspection resources.

One way to achieve a risk-sensitive solution is to use a quantile of a one-tailed prediction interval of $p$ to represent the predicted risk. That is, instead of choosing the value that is best supported by the data, we choose a value that represents the upper limit of the interval with a specified probability of including the rate. For example, we might choose a predicted risk that corresponds to the endpoint of the upper 95% prediction interval of the true contamination rate. Informally, we could describe this by saying that we are 95% confident that the current contamination rate is less than $f$ (Fig. 1).

The behavior of different confidence intervals for the rate parameter of the binomial distribution has been the subject of some research (see, e.g., Madden et al. 1996, Brown et al. 2001, Cai 2005). There is considerable literature to suggest that traditional Wald-style confidence intervals are flawed (see, e.g., Brown et al. 2001, Cai 2005). For one-tailed confidence intervals with robust statistical properties, Cai (2005) recommends Jeffrey intervals, which are Bayesian intervals with a prior distribution on the binomial parameter equal to Beta(0.5, 0.5). Jeffrey intervals do not absolutely guarantee $(1 - \alpha)\%$ coverage for all possible combinations of samples and $\alpha$, but they represent a good compromise between coverage and parsimony.

Jeffrey intervals are easily computable in standard spreadsheets. The upper one-tailed confidence interval is computed from the inverse of the cumulative density function of the Beta distribution. The Beta distribution requires two shape parameters $a$ and $b$; for our model (Eq. 1), they are computed as $a = x + 0.5$ and $b = n - x + 0.5$. Thus if we wished to use an upper 95% confidence interval as our definition for the predicted risk, then the function call in popular spreadsheet programs for $f(n, x, 0.95)$, the estimated 95% upper confidence interval, would be

$$f(n, x, 0.95) = \mathrm{BetaInv}(0.95, x + 0.5, n - x + 0.5) \tag{2}$$

where $x$ is the observed number of contaminated items, $n$ is the number of items inspected, and BetaInv is the inverse of the cumulative distribution function for a Beta distribution.

Our resource allocation strategy relies on classification of the pathways into two classes: "fully inspected" and "sampled." The classification approach that we recommend follows hypothesis-testing to guide decision-making, as originally advocated by Neyman and Pearson (e.g., Neyman and Pearson 1933). The classification proceeds as follows. We classify pathways as either being of such a high risk that all items should be inspected or being of lesser risk so that only a proportion of items are inspected in order to confirm that this low-risk status remains valid. If there are no restrictions on resources, then the inspectorate nominates a risk cutoff, $r$, e.g., 1%. The value of $r$ may be determined with input from other stakeholders, and may vary across pathways as a function of expected consequence. If there are restrictions on resources, then $r$ would be determined iteratively. In either case, the value of the predicted risk ($f$) would be determined for each pathway and then (1) every pathway with $f \geq r$ would be fully inspected, and (2) every pathway with $f < r$ would be sampled at a rate to be determined, which we will examine further.

For pathways for which $f$ is close to $r$ or for which only a few items arrive, the sampling rate must be 100% in order to gain adequate information about the contamination rate. We shall still consider such pathways as being sampled rather than fully inspected.

We now examine the use of information measured on the pathway over a period of time to classify the pathway for the purposes of future inspection. Our goal is to determine a sampling rate that will estimate and constrain the probability that pathways are misclassified. We will use time steps of one year, but only for the sake of demonstration.

Because all of the items in the higher-risk pathways are inspected, we shall have an estimate of the underlying proportion of contaminated items in that pathway that is as accurate as possible. Even so, occasionally we will conclude that the predicted risk is lower than $r$ and erroneously classify the pathway as lower risk.

In assigning the sampling rates for the lower-risk pathways, we try to ensure that the pathway will not be

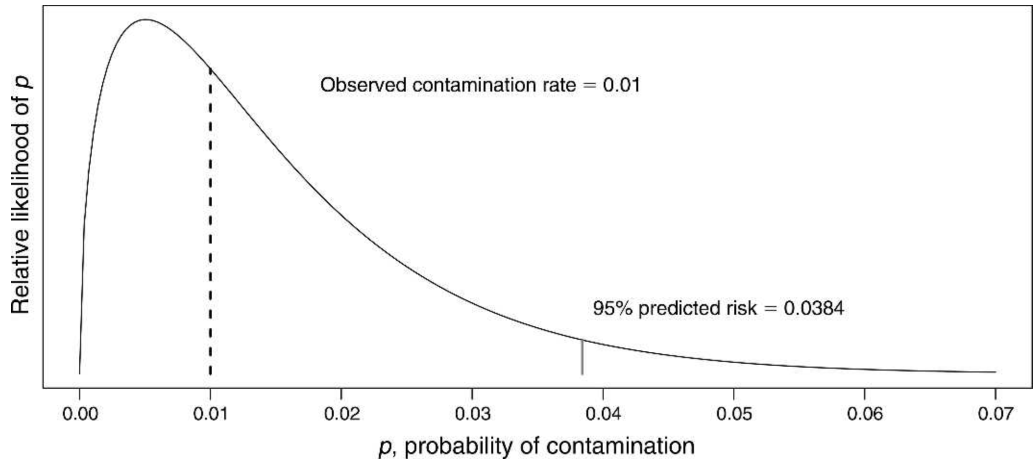


FIG. 1. Demonstration of the inflation of risk to account for uncertainty. The data are simulated inspections of 100 items, of which one was contaminated (i.e., contamination rate of 0.01 indicates failure to detect one contaminated item out of 100). The curved line shows the likelihood function of the estimated rate, using Jeffrey's approach (Bayesian intervals; see Cai 2005). The dashed vertical line shows the estimated contamination rate. The short, solid vertical line shows the one-sided, upper 95% confidence limit on the rate, which we define to be the "predicted risk."

misclassified as higher risk if the underlying contamination rate remains the same as represented in the historical data. In other words, we aim to inspect sufficient items next year so that there is only a small probability that the predicted risk observed in that year will exceed $r$ for those pathways that have lower risk. In Neyman-Pearson terms, we would like the classification to have high power, that is, we would like the test to have a high probability of correctly rejecting the null hypothesis that the pathway contamination is higher than $r$.

The model development follows. We have observed $x_1$ contaminated items from $n_1$ inspected items in the current year. We use Eq. 2 to assess whether or not the pathway is high risk. Higher-risk pathways are completely sampled in the next period, that is, $n_2 = N_2$, where $n_2$ is the number of items to be inspected in the next year.

For the lower-risk pathways we wish to determine $n_2$ so that there is high probability that the predicted risk that will be calculated from observing $x_2$ contaminated items from the $n_2$ inspected items will be below the acceptable risk cutoff $r$. We use the information provided by $x_1$ and $n_1$ from this year's data to estimate how likely different values of $x_2$ would be for a given value of $n_2$. This is done with the following ad hoc algorithm, which might be described as a "double-percentile" approach, to determine a suitable value of $n_2$.

For any given value of $n_2$, there is a corresponding (non-integer) value of $x_2$ for which the predicted risk would be exactly $r$. Eq. 2 provides one relationship between $x_2$ and $n_2$, although for speed of calculation it is better expressed in the form

$$\text{Betadist}(r, x_2 + ½, n_2 - x_2 + ½) = 1 - \alpha_2 \qquad (3)$$

where Betadist is the cumulative density function for the Beta distribution and $\alpha_2$ (for example, 0.05) is used to quantify our definition of predicted risk. Based on past experience, we want to choose $n_2$ so that the likelihood that we would observe more than this value of $x_2$ contaminated items is small, say $\alpha_2$. In order to resolve this equation, we need another way to relate $x_2$ and $n_2$.

We use the observed proportion of contaminated items ($x_1/n_1$) to calculate a value $y_2$ that is an upper percentile ($\alpha_1$) for the number of contaminated items that might be observed with the inspection of $n_2$ items:

$$y_2 = n_2 \times \text{BetaInv}(1 - \alpha_1, n_2 \times x_1/n_1 + 0.5, \\ n_2 - n_2 \times x_1/n_1 + 0.5). \qquad (4)$$

If we sample $n_2$ items, there would be only a small chance that we would observe more than $y_2$ contaminated items.

By equating $x_2$ and $y_2$, the value of $n_2$ that is unlikely to result in the pathway being classified as "higher risk" can be determined. This can be done by solving Eqs. 3 and 4, using a goal-seek tool. The values of $\alpha_1$ and $\alpha_2$ will most likely be identical, and choices for this value will depend on how concerned we are about wrongly classifying a lower-risk pathway as being of higher risk. This point will be discussed later.

Software to perform these operations in the open-source statistical environment R (R Development Core Team 2010) is available in the Supplement.

In the case of allocation under exact resource constraints, we fix either the total number of inspections that can be done ($\Sigma\ m$), or the nominated sampling rate ($\Sigma\ m/\Sigma\ n$), and use goal-seek software to find the value for $r$ that is satisfied by the available resources.

We note that the method suggested for choosing the sample size emphasizes the uncertainty that arises from small sample sizes. There are alternative ways of calculating a sample size that may cause more frequent

TABLE 1. Example of risk-sensitive allocation of detection resources.

| Commodity | $x$ | $n$ | $f$ (%) | TFSR(%) |
|---|---|---|---|---|
| A | 2 | 100 | 5.425 | 100 |
| B | 20 | 10 000 | 0.285 | 10.9 |

*Note:* From the observed number, $x$, of contaminated items from $n$ inspections, $f$ is the predicted risk, and TFSR is the tentative future sampling rate (expressed as a percentage) based on a 1% risk cutoff.

switching between whether the group should be sampled at 100% or at a lesser rate, but will also be less conservative in terms of penalizing the uncertainty.

## EXAMPLES

Assume that the historical data are as represented in Table 1. Using the approach we have just outlined, the solution to maximizing effectiveness in the short term is to allocate inspection efforts to commodity A because it has the higher contamination rate.

If we decide on a cutoff point of 1%, then the method suggests that we would inspect every instance of commodity A because the predicted risk for commodity A is higher than 1%. However, the predicted risk for commodity B is less than 1%, and so we would only sample commodity B, taking 1088 items in this case (a sampling rate of 10.9%). This is the sampling rate that will result in a predicted risk less than 1% with at least 95% probability, as long as the contamination rate does not increase.

### *ULD (unit loading devices) data*

To illustrate the proposed inspection strategy in a more realistic setting, we show a case study drawn from quarantine inspection in Australia (see Robinson et al. 2008). Under nationally mandated Increased Quarantine Intervention (IQI), implemented by the Australian Government in 2001 in response to the outbreak of foot-and-mouth disease in the United Kingdom, 100% of unit loading devices (ULDs; containers used for air transportation) were externally inspected.

Using historical inspection data provided by the Australian Quarantine Inspection Service, the national predicted risk for the external inspection for 2007 was 0.0918%. The regional and national predicted risks are presented in Table 2. The highest regional contamination rate for external inspection for the year of 2007 was 1.5%, in the Far North region. To simplify the analysis, we assumed that all items on each pathway were inspected, and that the inspections were 100% effective, that is, that no contaminated items escaped detection. AQIS leakage survey records suggest that the rate at which inspections miss contamination is sufficiently low that this assumption is reasonable for our purposes. For example, DAFF (2009:161, Table 16) notes that effectiveness for ULD inspections was consistently >90%.

The results of this analysis have two potential applications. (1) The predicted risk can be used to support an assessment of the utility of 100% inspections. Our results suggest that, because the approach rate is low, a risk-sensitive approach may recommend sampling these items at less than 100% and focusing inspection resources in other pathways if they are identified as being higher risk. (2) The results can be used to identify those regions where arriving items have higher contamination rates. Table 2 identifies Far North as having higher contamination approach rates than the other regions. For example, if we were to take 1% as a suitable risk cutoff, then on the strength of these results, it may be reasonable to maintain the 100% inspection rates in the Far North region, but to consider reducing the inspection rates for the other regions.

### *Simulations*

Finally, we use regional historical quarterly ULD inspection data to compare the effectiveness of the naive strategy of allocating all inspection resources into the pathways that have the highest predicted contamination rates at a given point in time (strategy 1) with our proposed allocation strategy, which monitors all pathways (strategy 2). Our goal is to use this real data set to provide an example of how much effectiveness is lost by demanding up-to-date information about the contami-

TABLE 2. Suggested inspection rates of ULDs (unit loading devices) in Australia.

| Region | Contaminated | Inspected | $p$ (%) | $f$ (%) | TFSR (%) |
|---|---|---|---|---|---|
| Southeast Queensland | 58 | 37 743 | 0.154 | 0.190 | 2.46 |
| Far North Queensland | 33 | 2 957 | 1.116 | 1.470 | 100.00 |
| New South Wales | 137 | 207 764 | 0.066 | 0.076 | 0.33 |
| South Australia | 59 | 17 510 | 0.337 | 0.415 | 10.10 |
| Victoria | 24 | 91 491 | 0.026 | 0.036 | 0.65 |
| Western Australia | 0 | 14 067 | 0.000 | 0.014 | 3.83 |
| National | 311 | 371 532 | 0.084 | 0.092 | 0.20 |

*Notes:* For each region, the number of contaminated and inspected ULDs in 2007 was used to determine the average contamination rate, $p$, and the predicted risk, $f$, both expressed as percentages. Eq. 4 was used to calculate the tentative future sampling rate, TFSR, based on a risk cutoff of 1%. The table also gives the sampling rate based on the national figures taken as a whole. It should be noted that this national sampling rate is not the same as that obtained by aggregating the regional sampling rates.

Table 3. Summary of simulated comparison of inspection allocation strategies.

| Strategy | Inspection rate (%) | Effectiveness (%) | Leakage |
|---|---|---|---|
| Inspect random items | 6.4 | 6.4 | 0.095% |
| 1) Only inspect the most risky pathways | 6.4 | 24.7 | 0.076% |
| 2) Reserve some of the inspection capability to monitor lower-risk pathways | 6.4 | 21.4 | 0.079% |

*Note:* The inspection rate is the percentage of items that are inspected and was constant for each strategy; effectiveness is the percentage of contaminated items that are intercepted; and leakage is the average rate at which items are still contaminated after intervention has been performed.

nation rate for all of the pathways, and how much effectiveness is gained by detecting and responding to changes in approach rate. We view strategy 1 as a strawman argument; it is impossible to allocate inspection resources reliably without at least some information on the pathways.

The inspection data cover seven regions, which we treat as pathways, from Quarter 3 of 2003 to Quarter 1 of 2008, inclusive (19 quarters). The data comprise external inspections of 1 786 858 ULDs, of which 1805 were contaminated, representing a 0.10% approach rate for contamination.

Strategy 1 provides an algorithm for allocating resources to different pathways, but no way to set the total available resources, other than by budgetary constraints. In order to provide a balanced comparison, we ran strategy 2 first, using 1% as the risk cutoff, and recorded the mean total number of inspections across all of the simulations. We then used the same average number of inspections for strategy 1.

The following algorithms were used to compute the efficiency of the two strategies. For strategy 1, we wish to allocate the inspection resources to the riskiest pathways. In order to deploy this strategy it is essential that we have an estimate of the risk of each pathway to begin with. To simulate this process we used the inspection results for each of the 19 quarters as the estimate of risk to be used for allocating inspection resources for all of the other quarters. The allocated sampling rates were then held constant throughout the simulation, as prescribed by the strategy. This simulation approach resulted in 19 distinct inspection time lines, the results of which we summarized by their mean. The small variation in effectiveness between simulation runs arose from finding a slightly different number of contaminated items in the one region that was only partially inspected (because resources ran out).

The effectiveness of strategy 2 was estimated by simulating the following allocation algorithm 500 times. Again, we assumed that all items were inspected in the first quarter and determined the amount of sampling for the second quarter. We then simulated the number of contaminated items detected using random draws from the hyper-geometric distribution. Thereafter, the simulated results of inspection for one quarter were used to determine the sampling rate for the next quarter.

Results of the comparison are reported in Table 3, they do not include the initial quarter. In terms of categorization, the contamination rate for South Australia was close to the risk cutoff, and the state switched from the 100% inspection category to the sampled category an average of 5.9 times over the 17 possible quarters (switching is impossible in the first quarter, and meaningless in the second). The Far North switched on average 1.1 times, and switching rates were negligible for the other regions.

## Discussion

Table 3 shows that the positive consequences for effectiveness and leakage of adopting either strategy 1 or 2, as opposed to inspecting a similar number of items that are chosen randomly, are substantial. It is clear that there are systematic differences between the inherent contamination rates of the pathways for this example, and that these differences persist sufficiently to justify taking account of them in allocating inspection resources. Furthermore, for this example, the loss of effectiveness that results from monitoring the less-risky pathways compared to concentrating solely on the highest-risk pathways is negligible. Recall that strategy 1 allocates inspection resources optimally based on complete knowledge at a specific point in time. The differences between the effectiveness and leakage of strategy 1 and strategy 2 are modest.

Generally speaking, the algorithm that we have presented will allocate more resources to the monitoring of pathways that are known to have high contamination rates and to those pathways with a poorly known contamination rate. Only pathways that are reasonably confidently known to be contaminated at low rates will be allocated comparatively fewer resources. However, it is likely that there will be pathways that have risk close to the cutoff. Such pathways will flip between the high and low classes. In such cases, the proportions of items that are contaminated in the pathways will be similar, and the consequences of the misclassification will be small in terms of the number of contaminated items that are not detected.

It is inevitable that estimates and arbitrary trade-offs will be made in balancing risk management objectives. Our approach to allocating inspection resources differs from solutions advanced by other authors, for example Surkov et al. (2008), in that it provides a simpler and more straightforward framework for handling questions of risk. The difference is in the level of abstraction that is invoked. The cutoff can be set relative to economic information about consequences of incursion if it is available. Organizations that lack the means or the data to deploy a complete economic approach will find our simpler strategy more feasible.

The model we have outlined does not deal explicitly with uncertainty in the number of items arriving in each pathway. We have assumed that, in the first instance, the current volume (or a percentage change) is a reasonable estimate of the future volume. If the cutoff value $r$ is determined without consideration of the number of items that will be inspected, the number of items in a pathway only enters into the calculations in converting the number of samples required into a sampling rate. If, during the year, it is thought that the number of items arriving has materially changed, the sampling rate can be recalculated. In contrast, if the value of $r$ is determined by the amount of inspections that can be done, the consequences of uncertainty about the number of items arriving is much greater. In such a case, a more frequent analysis of the data to recalculate sampling rates is probably the simplest approach.

We have considered a situation in which it is not possible to "partly" inspect: an item is either inspected or it is not. There are cases in which one can set the level of intervention for each pathway, for example by inspecting a different number of fruits in a carton and/or inspecting a different number of cartons (see, e.g., Cannon 2009). Although setting the level of intervention is an attractive approach to allocating resources if it is possible, there is still the same problem about needing to monitor adequately for changes in the contamination rate. The optimal solution may not provide sufficient information and a compromise solution must be accepted.

In our examples, we nominated the 0.95 quantile in defining predicted risk. However, it is not clear that this amount represents an appropriate penalty for uncertainty, and the most appropriate value could be more or less. Strictly speaking, this consideration is a function of the consequences of the type 1 and type 2 errors, and estimates of the marginal benefits of inspection, such as those published by Surkov et al. (2008), may be useful in determining the best trade-off. Otherwise, we have assumed that the inspectorate will account for the consequences of invasion by setting an appropriate risk level.

If we have a fully automatic system for determining the rate at which each pathway is sampled, misclassification of a lower-risk pathway as a higher-risk "100% inspected" pathway is not a substantial problem. We would have the ability to recalculate the sampling rates, possibly not instantaneously, but certainly quite frequently. Hence any pathway falsely classified as higher risk would soon revert to a lower-risk category. Such frequent changes of sampling rates may be difficult to manage in a manual system where inspectors need to decide which items to inspect.

We emphasize that the prescriptions from analyses like these should be treated as guidelines, rather than hard-and-fast rules. The practicalities of the physical constraints of inspection will inevitably outweigh some elements of the risk analysis. Even if the 100% inspection class remains relatively constant, the sample rate for the non-100%-sampled categories will change each year. This regular change may lead to complications that outweigh the benefits. Under these circumstances, using a constant rate (say 5% or 10%) would be more practical, perhaps with the proviso that if the number of items approaching is low, the sampling rate should be increased. Furthermore, although it is not considered in this analysis, a system must always be flexible enough to cope with intelligence from external sources that the risk has changed, for example because of an outbreak of disease in the source country.

We have used annual and quarterly data as a basis for inspection and allocation, and considered each pathway distinctly. Future work will determine whether some temporal smoothing approach will allow for improved use of resources. We will also examine the question of whether some aggregation across pathways, or alternatively smoothing across estimates of pathway contamination rates, provides greater stability without sacrificing too much specificity.

To sum up, in order to provide security against mistakenly ignoring contaminated pathways, we believe that no pathway should go entirely uninspected. This position differs from advice by, for example, Horan et al. (2002), who advocate allocation of few or no resources to confronting events that are less possible, regardless of the expected damages of those events. We agree with Cannon (2009), who describes leaving some pathways uninspected as unsatisfactory. Our prescriptions differ from those of Hua et al. (2006) and Surkov et al. (2008) in this way; no pathway is left uninspected. Naive risk analysis focuses on maximizing detection rates in the present. Strategic risk analysis demands both detection and estimation, recognizing that future resource allocation will require up-to-date information about the relative risks of pathways. We provide an algorithm that allows for this strategic approach.

## Acknowledgments

The authors thank the following people who have provided invaluable advice and input: Wayne Atkinson, Ferne Clark, Greg Hood, Robert Langlands, and Chris Miller. This research was supported by the Australian Centre of Excellence for Risk Analysis and the Biosecurity Services Group. The authors are grateful to two anonymous reviewers for detailed and thoughtful comments. The opinions remain solely those of the authors.

## Literature Cited

Brockerhoff, E., J. Bain, M. Kimberley, and M. Knizek. 2006. Interception frequency of exotic bark and ambrosia beetles (Coleoptera: Scolytinae) and relationship with establishment in New Zealand and worldwide. Canadian Journal of Forest Research 36:289–298.

Brown, L. D., T. T. Cai, and A. DasGupta. 2001. Interval estimation for a binomial proportion (with discussion). Statistical Science 16:101–133.

Cai, T. T. 2005. One-sided confidence intervals in discrete distributions. Journal of Statistical Planning and Inference 131:63–88.

Cannon, R. M. 2009. Inspecting and monitoring on a restricted budget: where best to look? Preventive Veterinary Medicine 92:163–174.

Colautti, R. I., S. A. Bailey, C. D. A. van Overdijk, K. Admunsen, and H. J. MacIsaac. 2006. Characterised and projected costs of nonindigenous species in Canada. Biological Invasions 8:45–59.

Costello, C., M. Springborn, C. McAusland, and A. Solow. 2007. Unintended biological invasions: Does risk vary by trading partner? Journal of Environmental Economics and Management 54:262–276.

Cox, L. A., Jr. 2009. What's wrong with hazard-ranking systems? An expository note. Risk Analysis 29:940–948.

DAFF (Department of Agriculture, Fisheries and Forestry). 2009. Annual report 2007–08. Technical report. Australian Government, Department of Agriculture, Fisheries and Forestry, Canberra, Australian Capital Territory, Australia.

Govindaraju, K., M. Bebbington, and T. Wrathall. 2010. Statistical evaluation of the New Zealand Food Safety Authority sampling protocol for imported food. Risk Analysis 30: 817–826.

Gurevitch, J., and D. K. Padilla. 2004. Are invasive species a major cause of extinctions? Trends in Ecology and Evolution 19:470–474.

Haack, R. 2001. Intercepted Scolytidae (Coleoptera) at US ports of entry: 1985–2000. Integrated Pest Management Reviews 6:253–282.

Haack, R. 2006. Exotic bark- and wood-boring Coleoptera in the United States: recent establishments and interceptions. Canadian Journal of Forest Research 36:269–288.

Horan, R. D., C. Perrings, F. Lupi, and E. H. Bulte. 2002. Biological pollution prevention strategies under ignorance: the case of invasive species. American Journal of Agricultural Economics 84:1303–1310.

Hua, Z., S. Li, and Z. Tao. 2006. A rule-based risk decision-making approach and its application in China's customs inspection decision. Journal of the Operational Research Society 57:1313–1322.

Hulme, P. E. 2009. Trade, transport and trouble: managing invasive species pathways in an era of globalization. Journal of Applied Ecology 46:10–18.

Kenis, W., W. Rabitsch, M. A. Auger-Rozenberg, and A. Roques. 2007. How can alien species inventories and interception data help us prevent insect invasions? Bulletin of Entomological Research 97:489–502.

Langor, D. W., and J. Sweeney. 2009. Ecological impacts of non-native invertebrates and fungi on terrestrial ecosystems. Biological Invasions 11:1–3.

Madden, L. V., G. Hughes, and G. P. Munkvold. 1996. Plant disease incidence: inverse sampling, sequential sampling, and confidence intervals when the observed mean incidence is zero. Crop Protection 15:621–632.

Moffitt, L. J., J. Stranlund, and C. Osteen. 2008. Robust detection protocols for uncertain introductions of invasive species. Journal of Environmental Management 89:293–299.

Neyman, J., and E. S. Pearson. 1933. On the problem of the most efficient tests of statistical hypotheses. Philosophical Transactions of the Royal Society of London A, Containing Papers of a Mathematical or Physical Character 231:289–337.

Office of Technology Assessment (OTA). 1993. Harmful non-indigenous species in the United States. Technical Report OTA-F-565. OTA, Washington, D.C., USA.

Pimentel, D. 2007. Area-wide pest management: environmental, economic, and food issues. Pages 35–47 in M. J. B. Vreysen, A. S. Robinson, and J. Hendrichs, editors. Area-wide control of insect pests. Springer Verlag, Dordrecht, The Netherlands.

Pimentel, D., L. Lach, R. Zuniga, and D. Morrison. 2000. Environmental and economic costs of nonindigenous species in the United States. BioScience 50:53–65.

Pimentel, D., R. Zuniga, and D. Morrison. 2005. Update on the environmental and economic costs associated with alien-invasive species in the United States. Ecological Economics 52:273–288.

Press, W. H. 2009. Strong profiling is not mathematically optimal for discovering rare malfeasors. Proceedings of the National Academy of Sciences USA 106:1716–1719.

R Development Core Team. 2010. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ⟨http://www.R-project.org⟩

Ramirez-Marquez, J. 2008. Port-of-entry safety via the reliability optimization of container inspection strategy through an evolutionary approach. Reliability Engineering and System Safety 93:1698–1709.

Robinson, A., M. Burgman, W. Atkinson, R. Cannon, C. Miller, and H. Immonen. 2008. AQIS import clearance data framework. Technical Report 0804. Australian Centre of Excellence for Risk Analysis, University of Melbourne, Parkville, Victoria, Australia.

Savidge, J. A. 1987. Extinction of an island forest avifauna by an introduced snake. Ecology 68:660–668.

Sikder, I. U., S. Mal-Sarkar, and T. K. Mal. 2006. Knowledge-based risk assessment under uncertainty for species invasion. Risk Analysis 26:239–252.

Simberloff, D. 2006. Risk assessments, blacklists, and white lists for introduced species: Are predictions good enough to be useful? Agricultural and Resource Economic Review 35:1–10.

Surkov, I. V., A. G. J. M. Oude Lansink, O. van Kooten, and W. van der Werf. 2008. A model of optimal import phytosanitary inspection under capacity constraint. Agricultural Economics 38:363–373.

Work, T. T., D. G. McCullough, J. F. Cavey, and R. Komsa. 2005. Arrival rate of nonindigenous insect species into the United States through foreign trade. Biological Invasions 7:323–332.

Xie, Y., Z. Y. Li, W. P. Gregg, and L. Dianmo. 2001. Invasive species in China: an overview. Biodiversity and Conservation 10:1317–1341.

## SUPPLEMENT

R code used to create the results that are presented in Table 2 (*Ecological Archives* A021-066-S1).