Optimal Sampling Strategy for Probability Estimation: An Application to the Agricultural Quarantine Inspection Monitoring Program

Huidi Ma[1], Benjamin D. Leibowicz[1*], and John J. Hasenbein[1]

[1]Operations Research and Industrial Engineering, The University of Texas at Austin, ETC 5.128D, 204 E. Dean Keeton St., Stop C2200, Austin, TX 78712, United States.

[*]Corresponding author. Email: bleibowicz@utexas.edu.

Abstract

Imported agricultural pests can cause substantial damage to agriculture, food security, and ecosystems. In the United States, the Agricultural Quarantine Inspection Monitoring (AQIM) program conducts random sampling to estimate the probabilities that cargo and passengers arriving at ports of entry carry pests. Assessing these risks accurately is critical to enable effective policies and operational procedures. In this paper, we formulate an optimization model that minimizes the mean squared error of the probability estimates that AQIM obtains. The central decision-making tradeoff that the model explores is whether it is preferable to sample more arriving containers (and fewer boxes per container) or more boxes per container (and fewer containers), given limited resources. We first derive an analytical solution for the optimal sampling strategy by leveraging several approximations. Then, we apply our model to a numerical case study of maritime cargo sampling at the Port of Long Beach. We find that, across a wide range of parameter settings, the optimal strategy samples more containers (but fewer boxes per container) than the current AQIM protocol. The difference between the two strategies and the accuracy improvement with the optimal approach are larger if the pest statuses of boxes in the same container are more strongly correlated.

Keywords: sampling; probability estimation; inspections; agricultural pests; risk assessment

1 Introduction

The high-volume international flow of goods and people that characterize the modern global economy heightens the risks of importing agricultural pests and diseases across national borders. Once established within a country, these pests can cause significant damage to agriculture, food security, and ecosystems. The Food and Agriculture Organization (2021) recently indicated that invasive pests and diseases account for approximately 40% of annual global crop losses, inflicting a direct economic impact exceeding $220 billion, underscoring the urgent need for effective cross-border pest control measures to safeguard international trade, biodiversity, and public health. The occurrence of African Swine Fever (ASF) outbreaks in Asia, as documented by Gale, Kee, and Huang (2023), has led to the loss of millions of pigs and substantial disruptions in the pork industry. Illustratively, the swift expansion of the Fall Armyworm in Africa and Asia constitutes a significant menace to cereal crops and the livelihoods of numerous smallholder farmers, as highlighted by Makgoba, Tshikhudo, Nnzeru, and Makhado (2021). Within the United States (U.S.), apprehensions about honey production and pollination services have surged due to a 40% reduction in bee populations attributed to Colony Collapse Disorder and Varroa mite infestations, as indicated in a report from the National Honey Bee Health Stakeholder Conference Steering Committee (2012).

In the U.S., the Plant Protection and Quarantine (PPQ) division of the U.S. Department of Agriculture (USDA) is tasked with safeguarding the nation's agriculture and natural resources. PPQ partners with Customs and Border Protection (CBP) to administer and carry out the Agricultural Quarantine Inspection (AQI) program. AQI performs targeted inspections of cargo and people entering the U.S. at ports of entry (POEs) in order to detect agricultural pests and diseases and prevent them from being imported. While AQI's goal is to detect pests, it must pursue this mission with limited inspection resources and while facilitating international trade. This is a careful balancing act that makes it important for AQI to have reliable information about the rates of pest arrivals via various pathways for deciding where and what to focus its inspections on.

The Agricultural Quarantine Inspection Monitoring (AQIM) program is responsible for providing this information to AQI, as outlined in the Agricultural Quarantine Inspection Monitoring (AQIM) Handbook (U.S. Department of Agriculture, 2021). While AQI's mission is to *mitigate risk* through inspections targeted as intelligently as possible, AQIM's mission is to *assess risk* by conducting random sampling of cargo and people arriving at U.S. POEs. Random sampling allows AQIM to estimate the probability that a unit of a particular commodity arriving at a given POE contains an actionable pest. These probability estimates inform how AQI deploys its inspection resources for risk mitigation and allow the U.S. government to evaluate the performance of AQI by examining the disparity between estimated and intercepted pest arrivals.

Two illustrative examples demonstrate the considerable value of AQIM data for enhancing policies and operations designed to counter the agricultural pest threat. The first example concerns the recent African Swine Fever outbreak in the Caribbean (U.S. Department of Agriculture, 2019). During this outbreak, data from AQIM regarding passenger traffic passing through Puerto Rico en route to the U.S. mainland facilitated a comprehensive analysis. This analysis focused on the encountered products, the prevalence of incoming passengers carrying items of concern, and their countries of origin. This information obtained by AQIM resulted in PPQ modifying passenger inspection rates in its predeparture inspection program. Second, in evaluating the feasibility of implementing a passenger preclearance initiative, particularly in terms of agricultural pest concerns, in a Southeast Asian country, PPQ relied on AQIM data (Animal and Plant Health Inspection Service, 2017). Examination of AQIM data revealed a significant risk level, culminating in the decision that the proposed preclearance program was not viable.

The random sampling protocol currently employed by AQIM, and defined in its handbook (U.S. Department of Agriculture, 2021), specifies a certain number of *containers* to inspect each week for each commodity class at each POE. Ultimately, AQIM reports the fraction of these sampled containers in which it finds an actionable pest as its probability estimate for that commodity class and POE. It is important to note that AQIM inspectors do not inspect all *boxes* that are inside a sampled container. Instead, AQIM follows a "hypergeometric" table that indicates how many boxes to inspect based on the total number of boxes in the container. We will describe this current AQIM protocol in detail in Section 3. Importantly, we note that the current division of inspection resources between containers and boxes is not tied to any mathematical notion of optimal resource allocation. Perhaps more accurate probability estimates could be obtained by sampling more containers but fewer boxes per container, or vice versa. This is the central strategy decision that we focus on in this paper.

To investigate this strategy decision and optimize the AQIM random sampling approach, we formulate an optimal resource allocation model that minimizes the mean squared error (MSE) of the pest probability estimate. The model decides how many containers to sample and how many boxes to inspect per sampled container, subject to a budget constraint. The budget constraint can be specified as a maximum number of boxes that can be inspected or expanded to incorporate a fixed cost for sampling a container. By introducing several approximations, we are able to derive an analytical solution for the optimal sampling strategy. Following this theoretical analysis, we parameterize a numerical case study based on data from the Port of Long Beach maritime POE. For this case study, we present optimal solutions, conduct sensitivity analysis on input parameters that are difficult to assign values to using empirical data, and compare the performance of our model's optimal sampling strategy to that of the current AQIM protocol.

Our main findings can be succinctly summarized as follows:

1. The optimal sampling strategies derived from our model tend to sample more containers, but inspect fewer boxes per sampled container, than the current AQIM protocol. This difference is quite robust to variations in parameter settings.

2. The optimal sampling strategies typically produce probability estimates with substantially lower MSEs than the current AQIM protocol.

3. The differences between the optimal sampling strategy and the current AQIM protocol, in terms of strategy and performance, are greater when the correlation between the pest statuses of boxes in the same container is higher.

While our analysis was inspired by the AQIM application and we describe the model using language tailored to its context, we emphasize that our model and analysis are actually quite general. They could certainly apply to other applications of two-level random sampling for probability estimation, with two levels that are analogous to the containers and boxes in the AQIM case. For example, in quality control, a manufacturer may want to estimate the probability that a shipment contains a defective unit, in which case it could randomly sample a certain number of shipments and inspect some number of units per shipment. The findings and insights that we uncover in this paper would be relevant to this other application, with shipments analogous to containers and units analogous to boxes.

The remainder of this paper is structured as follows. We review the most relevant literature in Section 2, then provide technical background on the current AQIM protocol in Section 3. Section 4 presents our optimization model, which we analyze theoretically in Section 5. We outline our Port of Long Beach case study in Section 6 and then present and discuss results for it in Section 7. Section 8 concludes with a summary of our most important findings and future research directions.

2 Literature Review

In this section, we review three branches of literature that are relevant to our work along different dimensions: (1) optimal inspection strategies for national security, (2) operations research for invasive species management, and (3) sampling.

2.1 Optimal Inspection Strategies for National Security

Numerous studies have analyzed optimal inspection strategies for various national security applications. In the area of agricultural pest inspections, Chen, Epanchin-Niell, and Haight (2018) formulate a model that minimizes the expected cost of pest importations by choosing how many units to sample from each lot. They apply this model to a case study based on plant shipments imported from Costa Rica to Miami. Trouvé and Robinson (2024) similarly investigate the optimal allocation of sampling effort, but their model incorporates the overdispersion phenomenon whereby different consignments within the same pathway exhibit varying pest infestation rates. They demonstrate how considering overdispersion as well as other features of the pest arrival process when determining the inspection strategy can substantially reduce propagule pressure. It is important to note that both Chen et al. (2018) and Trouvé and Robinson (2024) optimize sampling strategies in order to minimize pest importations and their impacts (consistent with the goal of AQI inspections), with infestation rates included in their models as input parameters. By contrast, in this paper, we optimize decisions about how to randomly sample cargo in order to estimate infestation rates as accurately as possible (consistent with the goal of AQIM inspections).

Batabyal and Beladi (2006) investigate optimal resource allocation decisions for preventing biological invasions that occur via international trade by conducting a queueing-theoretic analysis. Employing queueing theory, the authors analyze the prevention problem from a long-term perspective, characterizing regulatory regimes as distinct queues. They formulate a publicly owned port manager's decision problem as an optimization challenge using queuing-theoretic techniques. The study compares and contrasts optimality conditions of different inspection regimes. Building on the foundational work by Batabyal and Beladi (2006), Yamamura, Katsumata, Yoshioka, Yuda, and Kasugai (2016) extend this line of inquiry by elucidating various statistical theories incorporated into Japan's import plant quarantine systems. These encompass import inspections, early detection procedures, and emergency control. The authors illustrate these theories through real instances of quarantine measures adopted against the infiltration of plum pox virus disease and citrus huanglongbing. A key finding underscored by the authors is the critical significance of implementing suitable import quarantine systems to avert the inadvertent introduction of invasive alien pests.

Beyond cross-border pest control, optimal inspection strategies in other national security contexts have garnered attention from researchers. Bagchi and Paul (2014) investigate how changes in intelligence gathering impact potential terrorists' behavior and airport security. It suggests that increasing intelligence spending may raise the likelihood of attacks by highly motivated terrorists but decrease the chances of attacks by less motivated ones. Additionally, higher intelligence spending leads to shorter airport security lines, reducing congestion issues. The paper also highlights that lowering the cost of screening unequivocally improves overall social welfare. Techniques that make screening more cost-effective can enhance social welfare without necessarily cutting intelligence spending. Moreover, Clauset and Woodard (2013) present a statistical algorithm for estimating the probability of large terrorist events, like the 9/11 attacks, using semi-parametric models and a nonparametric bootstrap. The algorithm is also employed for a data-driven statistical forecast of a similar event in the next decade. Accurate probabilities for significant terrorist events contribute valuable insights for global risk assessment, guiding long-term planning and response efforts. Jacobson, Kobza, and Nakayama (2000) introduce a non-intrusive sampling method, utilizing the observed numbers of alarms and clears in security-system operations, to estimate the probabilities of threat, false alarm, and false clear in access-control security systems. The derived estimators offer a means to assess the performance of such systems. Additionally, the paper delves into the sampling procedure employed for estimating threat and false alarm probabilities, along with the probability of a false clear. It establishes convergence properties and confidence intervals for these estimators, and an illustrative example is provided to showcase their utility.

Container inspections for addressing other threats at ports also play a pivotal role in optimal inspection strategies. Bakshi, Flynn, and Gans (2011) examine the operational impact of container inspections at international ports, and their paper specifically focuses on the tradeoff between the security derived from inspections and the resulting congestion in the system. The study recognizes the vulnerability of the U.S. to maritime terrorism, highlighting the potential for terrorists to conceal nuclear devices within shipping containers, thereby causing widespread disruption to global supply chains.

2.2 Invasive Species Management and Operations Research

The field of invasive species management has seen significant contributions from operations research methodologies. Büyüktahtakın and Haight (2018) conduct a comprehensive review of operations research models in invasive species management, emphasizing the need for efficient decision tools to prioritize actions and minimize adverse impacts. They illustrate various biological and economic aspects of invasive species management by introducing a spatio-temporal optimization model. The paper categorizes relevant literature based on modeling methods, critically examines existing research, identifies limitations, and proposes directions for further exploration in optimizing invasive species management planning.

One crucial aspect of optimization models in invasive species management is the formulation of control models. The optimization problem within these models revolves around the efficient allocation of resources among various control activities, with the overarching goal of minimizing invasion damage over time, adhering to a designated control budget, and considering the biological dynamics of the invader. To tackle this dynamic optimization problem, researchers employ various approaches, including dynamic programming (DP), mathematical programming, and optimal control. For instance, Billionnet (2013) describes the application of mathematical programming to support decision makers in safeguarding biodiversity. The study presents compelling examples of optimization problems linked to conservation planning, such as the careful selection of nature reserves, effective control of landscape fragmentation, ecologically sustainable forest utilization, combating invasive species, and preserving genetic diversity. To address the complexity of controlling invasive species, Kıbış and Büyüktahtakın (2017) introduce a mixed-integer programming (MIP) model designed for controlling sericea infestation, demonstrating its applicability through a case study. Solving the problem as a full dynamic optimization model, the MIP model outperforms its mixed-integer nonlinear programming (MINLP) equivalent and nonlinear programming (NLP) relaxation in solution quality. A comparison of five linearization methods reveals that the proposed approach consistently produces higher-quality solutions. Incorporating binary treatment decisions, dispersal factors, and probabilities, the model linearizes nonlinear aspects, providing computational solvability for practical-sized invasive species management problems. Examining the national-scale economic impact of terrestrial invasive species, Olson (2006) provides a comprehensive review of recent studies on the economics of invasive species management. It examines the economic literature concerning the control and prevention of biological invasions, as well as literature on international trade and trade policy regarding invasive species. Additionally, the paper provides an overview of selected studies focusing on terrestrial invasive plants, animals, and microbes.

In the realm of invasive species management, an array of prevention strategies, beyond border inspections, has garnered attention. These comprehensive measures include quarantines and formidable barriers strategically deployed to prevent the emigration of invasive species from their established populations. Optimization studies explore the establishment of barrier zones, containment, quarantine of existing populations, and preemptive removal of hosts to minimize damages and control costs. Epanchin-Niell and Wilen (2012) introduce a bioeconomic model for bioinvasions, considering spatial-dynamic processes and explicitly addressing spatial aspects. It explores optimal control policies based on bioeconomic parameters, spatial configurations, and initial invasion types. The geometry of the initial invasion and landscape influences optimal policies, which may involve leveraging landscape features or altering the invasion shape to minimize costs. Notably, these policies exhibit forward-looking behavior, aiming to slow and redirect the invasion away from areas with high potential damages or toward areas with low control costs. Sharov and Liebhold

(1998) demonstrate the effective use of barrier zones to slow the spread of the gypsy moth population, resulting in a significant reduction in its advancement rate. Furthermore, Moore et al. (2010) provide a comprehensive review of classical biological control programs, emphasizing the need for more projects to combat the increasing number of invasive species worldwide. Their categorization of projects into complete control, partial control, and ongoing initiatives highlights the significance of preserving biodiversity, products, and ecosystem services. Lastly, Kovacs, Haight, Mercader, and McCullough (2014) contribute a multidimensional perspective by exploring bioeconomic analyses in the context of an emerald ash borer invasion in an urban forest with multiple jurisdictions. They emphasize the importance of efficient resource allocation strategies to limit the impact of invasive species, particularly in scenarios with imperfect information about local payoffs and the intricate relationship between native species and timber harvesting. The primary focus is on maximizing the net benefits of ash trees while considering the impact of emerald ash borer infestation. The paper addresses the optimality of removing and replacing ash trees, suggesting that all affected trees should be removed if they are at risk of being killed by the emerald ash borer.

2.3 Sampling

The AQIM sampling protocol, which we will describe in detail in Section 3, has elements in common with several classical sampling methods including cluster sampling, stratified sampling, and two-stage sampling. In this subsection, we briefly summarize each of these sampling designs and clarify how the AQIM procedure is similar to, and different from, each of them. Then, we highlight several papers that specifically pertain to sampling for agricultural pests.

As discussed by Thompson (2012), cluster sampling involves dividing a population into clusters and randomly selecting some of these clusters to include all members within them in the sample. An extension of cluster sampling is two-stage cluster sampling, where a sample of secondary units is selected from each of the previously selected primary units. The AQIM sampling protocol is similar to two-stage cluster sampling in that containers are randomly sampled in the first stage, and then boxes from each selected container are randomly sampled in the second stage. Another sampling method that is relevant to AQIM is stratified sampling, where the population is divided into distinct non-overlapping groups called strata, and a sample is chosen from each stratum using a specific design. In stratified random sampling, the entire sample frame is partitioned into separate subgroups, and a simple random sample is independently taken within each subgroup. In the AQIM application, the pathways defined by different POE-commodity type combinations can be viewed as strata created by dividing the overall universe of arriving cargo and passengers along these lines.

A major difference between the AQIM sampling protocol and two-stage cluster sampling is that, in the AQIM application, the ultimate goal is to estimate the pest probability at the primary unit (container) level instead of at the secondary unit (box) level. By contrast, in two-stage cluster sampling, the definition and selection of clusters is generally a means to an end, a way to conveniently sample the whole population of secondary units in order to estimate probabilities at the secondary unit level. Similarly, AQIM procedures differ from classical stratified sampling in that, once the AQIM strata are defined, sampling and probability estimation take place independently in each stratum to derive stratum-level results for each pathway (POE-commodity type combination). In classical stratified sampling, the goal is usually to combine stratum-level sampling results to estimate a single probability for the entire population.

In terms of sampling for agricultural pests, several papers in the literature explore how to estimate the underlying parameters of pest arrival processes based on observed inspection results. Trouvé and Robinson (2021) provide evidence showing that overdispersion is a common feature of pest infestation rates across consignments from the same pathway. As a result, estimating the infestation rate using binary data on inspection outcomes (without information on how many units failed an inspection) without accounting for overdispersion will tend to underestimate pest threats, because the binary inspection data are right-censored. Montgomery, Petras, Takeuchi, and Katsar (2023) introduce Pest or Pathogen Spread (PoPS), an open-source simulation platform for pest arrivals and inspections. As they demonstrate, this simulator has a number of potential uses. One use is to estimate the true pest infestation rate and other underlying parameters of pest arrival processes, as the simulation parameters can be adjusted until the simulated inspection results roughly match the empirical inspection data collected in the real world. PoPS can also be used to simulate the effects of different inspection strategies and quantify tradeoffs between inspection effort and effectiveness. Kim, Hong, Egger, Katsar, and Griffin (2019) propose and test methods for assigning categorical risk scores to commodity-country combinations based on pest inspection results. They account for both the estimated pest interception rates and their confidence intervals, and find that a generalized linear mixed effects model performs well at prediction. In contrast to our work, none of these papers explicitly optimizes the sampling strategy with respect to an objective, and none of them considers the budget allocation tradeoff between sampling more containers versus inspecting more boxes per container.

3 Technical Background: Current AQIM Sampling Protocol

In this section, we describe the current sampling protocol employed by AQIM and outlined in its handbook (U.S. Department of Agriculture, 2021). It is important to understand how AQIM sampling currently takes place in order to establish the notions of a container, a box, and a sampling strategy that are embedded in our model in Section 4. Furthermore, we will eventually compare our optimal sampling strategies and their performance to those of the current AQIM protocol.

For a detailed explanation of the current AQIM sampling protocol, consider maritime cargo as an example pathway. Figure 1 provides a visual overview of how AQIM random sampling currently functions at a maritime cargo POE. AQIM first categorizes the entire universe of cargo arriving at this POE into three subgroups: Commercial Perishable Agricultural Cargo (e.g., fresh fruit, vegetables, cut flowers), Wood Packaging Material, and Italian Tile. For each subgroup, the AQIM Handbook dictates that AQIM randomly inspect two containers per week (eight per four-week period or "month") at each POE. Ultimately, the container-level probability estimates that AQIM reports will be specific to the POE and subgroup (e.g., Commercial Perishable Agricultural Cargo arriving at the Port of Seattle). The AQIM manager on site has some discretion for deciding how to randomly select containers to sample, but a common method is to generate random inspection times and

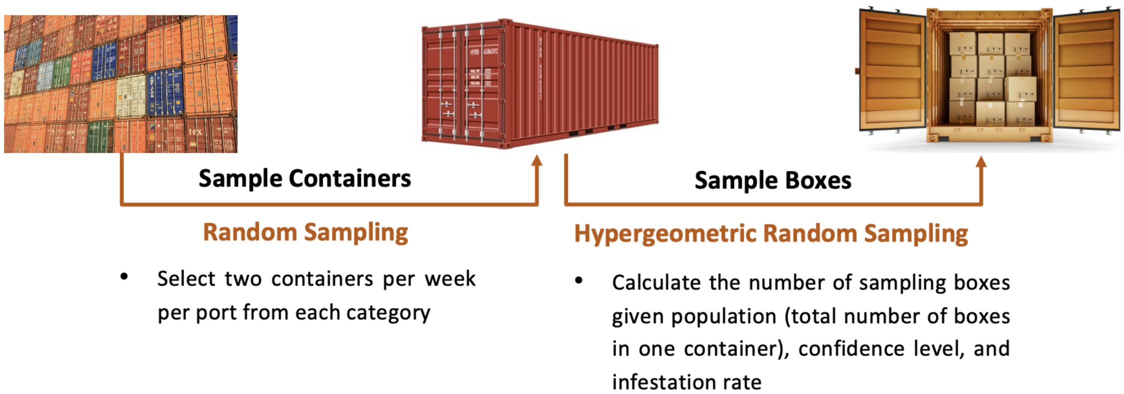

Figure 1: Illustration of the current AQIM sampling protocol.

sample the next container that arrives after each time.

Once AQIM randomly chooses a container to sample, the next step is to randomly select some boxes from the container to inspect for pests. The underlying logic for not inspecting every box in a container is that containers often carry many boxes, and inspecting all of them would be very resource-intensive while not necessarily being that much more effective at detecting pests than inspecting a subset of them. The precise number of boxes that AQIM inspects is based on the total number of boxes in the container, as prescribed in Table 1 from the AQIM Handbook (U.S. Department of Agriculture, 2021). The numbers prescribed in the table were calculated to ensure that, if 10% of boxes in a container carry a pest, then inspecting the given number of boxes will result in at least a 95% probability of detecting a pest in at least one inspected box. Since the number of inspected boxes that contain a pest is a random variable that follows a hypergeometric distribution (which describes sampling without replacement), AQIM refers to Table 1 as the *hypergeometric table*, and the process of randomly selecting boxes to inspect as hypergeometric sampling. During our site visits to the Port of Long Beach maritime cargo POE and Los Angeles International Airport air cargo POE, our conversations with AQIM managers and observations of inspections confirmed that AQIM personnel adhere to the AQIM Handbook's guidelines very closely.

Table 1: The hypergeometric table from the AQIM Handbook that specifies how many boxes from a container to inspect.

Total number of boxes	Number of boxes to sample
1-10	All boxes in the shipment
11-12	11
13	12
14-15	13
16-17	14
18-19	15
20-22	16
23-25	17
26-28	18
29-32	19
33-38	20
39-44	21
45-53	22
54-65	23
66-82	24
83-108	25
109-157	26
158-271	27
272-885	28
886 and up	29

Before proceeding, we emphasize that the current AQIM protocol is not tied to any notion of optimal resource allocation that would estimate pest probabilities as accurately as possible given the available inspection resources. The determination of the number of containers that AQIM samples lacks mathematical underpinnings. Rather, it is based on an intuitive assessment of the available inspection resources and the desire to avoid disrupting trade to an unwarranted degree. While the hypergeometric sampling of boxes is grounded in probability concepts, the 10% box infestation assumption and targeted 95% confidence level for detection are essentially arbitrary. Without an in-depth analysis, it is unclear whether 95% is the “right” confidence level to aim for in order to obtain the most accurate container-level pest probability estimates. Perhaps it would be better to reduce this confidence level in order to spread box inspections out over a larger sample of containers, or to increase it by inspecting more boxes per container even if it means sampling fewer containers. This is the central strategy decision that our model presented in the next section is formulated to analyze.

4 Model Description and Formulations

In this section, we present our optimization model and describe its basic structure and assumptions. Then, in the next section, we analyze it in order to derive an analytical solution for the optimal AQIM sampling strategy (with the help of some approximations that make the problem analytically tractable). Table 3 summarizes the nomenclature that we introduce in this section and use in our model formulations and subsequent analysis.

From a decision-making standpoint, we assume that AQIM's objective is to obtain the most accurate estimate of the probability that a container in a given subgroup (e.g., commercial perishable agricultural cargo) arriving at a given POE (e.g., Port of Long Beach) carries an actionable pest. First, let p represent the true probability that such a container carries a pest. Note that p is a parameter with a fixed value, but one that is unknown from AQIM's perspective; after all, p is what AQIM is trying to estimate. Now, let $\hat{p}$ denote the (container-level) estimate of the pest probability that AQIM obtains via its random sampling. Specifically, AQIM will report as its estimate $\hat{p}$ the number of sampled containers in which it finds a pest divided by the total number of containers that it samples. In other words, $\hat{p}$ is the empirical fraction of sampled containers in which AQIM detects a pest. Note that $\hat{p}$ is a random variable because it depends on random elements, as we will describe below. Due to random chance, AQIM may calculate a high $\hat{p}$ in one month, a lower $\hat{p}$ in the next month, and so on due to the random pest statuses of the containers that it samples and the boxes that it inspects from each container.

We assume that AQIM measures the accuracy of its pest probability estimate $\hat{p}$ in terms of its mean squared error (MSE) relative to the true pest probability p. The MSE that AQIM seeks to minimize in its optimization problem can be expressed as $E\left[(\hat{p}-p)^2\right]$. This is the objective function to be minimized in formulations (1) and (2). In the next section, we will derive an expression for the MSE as a function of the problem's variables and parameters.

Table 3: Nomenclature table

Parameters	Description
N	Total number of containers arriving per period
b	Number of boxes in each container
p_0	Probability that a container comes from a pest source
r	Probability that a box in a container from a pest source carries a pest
ρ	Correlation between the pest status random variables (0-1) of any two boxes from the same container (derived parameter)
B	Total budget
$cost_c$	Fixed cost of sampling a container
$cost_b$	Cost of inspecting a box
Decision Variables	**Description**
n	Number of containers to sample per period
i	Number of boxes to inspect from each sampled container

The strategy that AQIM chooses is defined by two decision variables that represent the number of containers that it randomly samples in each period (denoted by n) and the number of boxes that it randomly inspects from each sampled container (denoted by i). We assume that these two decision variables are bounded from above by the total number of containers in the subgroup that arrive at the POE per period (N) and the total number of boxes in each container (b), respectively. In practice, AQIM has limited resources that constrain how intensively it can randomly sample arriving cargo and passengers. Therefore, we incorporate a budget constraint with a total budget of B. The "costs" of sampling a container and inspecting a box are respectively represented as the $cost_c$ and $cost_b$ parameters. As we will demonstrate in our numerical case study, we design the budget constraint to be quite flexible. For instance, a situation with no fixed cost for sampling a container can be represented by setting $cost_c = 0$ and $cost_b = 1$, and measuring the budget B in units of box inspection capacity. On the other hand, if there is a fixed cost per container and various monetary cost components, $cost_c$, $cost_b$, and B can all be represented in dollars.

The first formulation that we use to represent AQIM's optimization problem is the following, which we refer to as the Integer Model:

Model 1: Integer Model

$$\text{Objective:} \quad \text{Minimize } E\left[(\hat{p}-p)^2\right] \tag{1a}$$

$$\text{SUBJECT TO:} \quad n \cdot \text{cost}_c + n \cdot i \cdot \text{cost}_b = B \tag{1b}$$

$$0 < n \leq N \quad \text{(Container constraint for } n\text{)} \tag{1c}$$

$$1 \leq i \leq b \quad \text{(Box constraint for } i\text{)} \tag{1d}$$

$$i \in \mathbb{Z}^+, n \in \mathbb{R}^+ \quad \text{(Feasible region for } i \text{ and } n\text{).} \tag{1e}$$

Constraint (1e) requires that the number of boxes inspected from each sampled container be an integer. This is realistic in that it would be hard to interpret a strategy that involves inspecting a fraction of a box. On the other hand, we allow n to take on a continuous value so that AQIM can exhaust all of its budget (i.e., have the budget constraint bind) in order to minimize the MSE. In practice, AQIM could exhaust whatever budget remains after sampling $\lfloor n \rfloor$ containers and i boxes per container by inspecting an "extra" box from a few containers, for example, but we do not represent this intricacy in our formulations in order to make the problem more tractable.

To enable the formal analysis that we carry out in the next section, we also consider the Continuous Model (2), which is defined as follows:

Model 2: Continuous Model

$$\text{Objective:} \quad \text{Minimize } E\left[(\hat{p}-p)^2\right] \tag{2a}$$

$$\text{SUBJECT TO:} \quad n \cdot \text{cost}_c + n \cdot i \cdot \text{cost}_b = B \tag{2b}$$

$$0 < n \leq N \quad \text{(Container constraint for } n\text{)} \tag{2c}$$

$$1 \leq i \leq b \quad \text{(Box constraint for } i\text{)} \tag{2d}$$

$$i, n \in \mathbb{R}^+ \quad \text{(Feasible region for } i \text{ and } n\text{).} \tag{2e}$$

This Continuous Model is exactly the same as the Integer Model (1), except that it changes the constraint (1e) to (2e). While (2) is in some sense less "realistic" than (1), it has the advantage of permitting us to obtain an analytical solution for the optimal AQIM sampling strategy, which we will derive in Section 5.

Now, we will describe the process by which pests are randomly present in containers and boxes. The randomness inherent in the pest statuses of the containers and boxes that AQIM inspects is what makes $\hat{p}$ a random variable, and based on the probabilistic process of pest arrivals, we will derive an equivalent expression for the MSE objective function, $E\left[(\hat{p}-p)^2\right]$, in the next section. Our representation of random pest arrivals allows us to vary the correlation between the pest statuses of boxes in a container, a parameter that AQIM currently has little information about but that could exert a strong influence on the optimal sampling strategy.

We assume that each container either comes from a *pest source*, with probability p_0, or does not come from a pest source, with probability $1 - p_0$. For example, a pest source could be a farm or warehouse where the shipment originated and that is infested with a pest. We assume that whether one container comes from a pest source is independent of whether any other container comes from a pest source. Conditional on a container coming from a pest source, each box in the container carries a pest with probability r. For containers that do not come from a pest source, the probability that any of its boxes features a pest is zero. We assume that the pest statuses of boxes in the same container are conditionally independent. To be more specific, given that a container is from a pest source, the events representing the presence (or absence) of pests in each box are all mutually independent. Of course, if a container does not come from a pest source, then no box contains a pest (hence, in this case, the pest events are technically all conditionally independent). However, in general, box pest statuses can exhibit some (unconditional) dependence. This is because knowing the pest statuses of some boxes in a container provides useful information about the likelihood that the container came from a pest source (which affects the probability of any unopened box carrying a pest). This dependence can be measured by the correlation coefficient (denoted ρ) between the two random variables that represent the binary pest statuses of any two boxes from the same container. In the next section, we will derive an expression for ρ in terms of p_0 and r.

5 Analysis

5.1 Preliminary Derivations

We begin our analysis by deriving the true probability that a container carries a pest (p), the distribution of the number of sampled containers in which AQIM detects a pest (a random variable we denote as R), and the correlation coefficient between the pest statuses of two boxes in the same container (ρ). These preliminary derivations allow us to write down an expression for the MSE objective function and interpret parameter settings by referring to ρ, which is somewhat more intuitive and directly measurable (from AQIM's point of view) than p_0 and r.

The true probability that a container features a pest is

$$p = p_0 \left(1 - (1 - r)^b\right). \tag{3}$$

Essentially, (3) says that in order to carry a pest, a container must come from a pest source and at least one of the b boxes in the container must feature a pest.

AQIM calculates its estimate of the container pest probability by dividing the number of sampled containers in which it detects a pest (R) by the number of containers that it samples (n). Therefore, according to its definition,

$$\hat{p} = \frac{R}{n}. \tag{4}$$

While n is a decision variable that AQIM chooses, R is a random variable. In order to derive the distribution of R, we first determine an expression for the probability that AQIM detects a pest in a container that it samples, which we represent as p_d. This detection probability can be written as

$$p_d = p_0\left(1-(1-r)^i\right). \tag{5}$$

The logic underlying (5) is that in order for AQIM to detect a pest in a container, the container must come from a pest source and at least one of the i boxes that AQIM inspects must feature a pest. Note that (5) is similar in form to (3), except the former reflects the fact that AQIM chooses how many boxes to inspect per container (i) and typically this will be less than b.

Given our expression for p_d in (5), whether AQIM detects a pest in each one of the n containers that it samples is essentially an independent coin flip, with probability p_d of yielding a pest detection. Therefore, R follows a binomial distribution, specifically $R \sim \text{Binomial}\left(n, p_0\left(1-(1-r)^i\right)\right)$. The probability mass function of R is thus

$$P(R=k) = \binom{n}{k} p_d^k \left(1-p_d\right)^{n-k}. \tag{6}$$

We now turn our attention to ρ, which denotes the correlation coefficient between the two random variables (call them X and Y) that indicate whether two boxes in the same container feature a pest (value 1 for pest, 0 otherwise). We derive an expression for ρ as a function of p_0 and r as follows:

$$\begin{aligned}
\rho &= \frac{Cov(X,Y)}{\sqrt{Var(X)Var(Y)}} \\
&= \frac{E(XY)-E(X)E(Y)}{p_0 r(1-p_0 r)} \\
&= \frac{P(X=1,Y=1)-p_0^2r^2}{p_0 r(1-p_0 r)} \\
&= \frac{P(Y=1|X=1)P(X=1)-p_0^2r^2}{p_0 r(1-p_0 r)} \\
&= \frac{p_0 r^2 - p_0^2 r^2}{p_0 r(1-p_0 r)} \\
&= \frac{r(1-p_0)}{1-p_0 r}.
\end{aligned} \tag{7}$$

Examining (7), we see that there is a one-to-one correspondence between the values of ρ and r, given some fixed value of p_0. Therefore, setting a value for r in our model is essentially equivalent to setting the correlation coefficient ρ. In the other direction – which we follow to parameterize our case study in Section 6 – if AQIM has beliefs about the values of p_0 and ρ, then it can use the relationship in (7) to establish the value of r.

5.2 Derivation of the Mean Squared Error

For the remainder of this section, we analyze the Continuous Model (2) with the assumption that there is no fixed cost to sample a container ($cost_c = 0$) and the budget constraint is specified in terms of the maximum total number of boxes that can be inspected. In other words, (2b) reduces to $n \cdot i = B$. For now this simplifies our analysis, and as we will describe in Section 6, the $cost_c = 0$ assumption more or less holds for our Port of Long Beach case study.

The MSE of the estimated container pest probability $\hat{p}$ is the objective function to be minimized in our model. Here, we start with the definition of the MSE and then derive an expression for it in terms of the decision variables and input parameters, as follows:

$$\begin{aligned} E\left[(\hat{p}-p)^2\right] &= E\left[\hat{p}^2 - 2\hat{p}p + p^2\right] = \frac{1}{n^2}E\left(R^2\right) - 2p\cdot\frac{1}{n}E(R) + p^2 \\ &= \frac{1}{n}\left(p_0 - p_0\left(1-r\right)^i\right)\left(1 - p_0 + p_0\left(1-r\right)^i\right) + \left(p_0 - p_0\left(1-r\right)^i\right)^2 \\ &\quad - 2\left(p_0 - p_0\left(1-r\right)^b\right)\left(p_0 - p_0\left(1-r\right)^i\right) + \left(p_0 - p_0\left(1-r\right)^b\right)^2. \end{aligned} \tag{8}$$

We can now write (8) as a function of a single variable by using the budget constraint (which holds with equality) to substitute $n = \frac{B}{i}$. Denoting the resulting expression for the MSE as $f(i)$ and letting $q = 1 - r$ for shorthand, the MSE is

$$f(i) = \frac{i}{B}p_0\left(1-q^i\right)\left(1 - p_0\left(1-q^i\right)\right) + \left(p_0\left(1-q^i\right)\right)^2 - 2p_0\left(1-q^b\right)p_0\left(1-q^i\right) + \left(p_0\left(1-q^b\right)\right)^2. \tag{9}$$

5.3 Approximate Analytical Solution

When we take the derivative of (9) and set it equal to zero to try to find the optimal sampling strategy (i^*), we obtain

$$f'(i) = \frac{ip_0q^i\ln(q)}{B}\left[2p_0(1-q^i) - 1\right] + \frac{p_0(1-q^i)(1-p_0(1-q^i))}{B} + 2p_0^2q^i\ln(q)\left(q^i - q^b\right) = 0. \tag{10}$$

This equation does not have an analytical solution. In order to make it analytically tractable, we introduce the following approximation for the q^i terms in (9) based on the first two terms of the Taylor expansion:

$$q^i \approx \frac{1}{1 + i\cdot\ln\frac{1}{q}}. \tag{11}$$

Figure 2 illustrates how the approximate expression in (11) compares to the exact value of q^i for different combinations of q and i values. In general, the plots show that the approximation is fairly accurate overall, especially when q is low and i is high. In addition, we set $q^b = 0$ in (9). Since b is often a large value (in the hundreds, or even thousands), q^b is essentially 0 for practical computations.

After introducing these substitions, we obtain the following as an approximation to f:

$$\begin{aligned} \hat{f}(i) = \frac{i}{B}p_0\left(1 - \frac{1}{1+i\cdot\ln\frac{1}{q}}\right)\left(1 - p_0\left(1 - \frac{1}{1+i\cdot\ln\frac{1}{q}}\right)\right) + \left(p_0\left(1 - \frac{1}{1+i\cdot\ln\frac{1}{q}}\right)\right)^2 \\ - 2p_0\left(1-q^b\right)p_0\left(1 - \frac{1}{1+i\cdot\ln\frac{1}{q}}\right) + \left(p_0\left(1-q^b\right)\right)^2. \end{aligned} \tag{12}$$

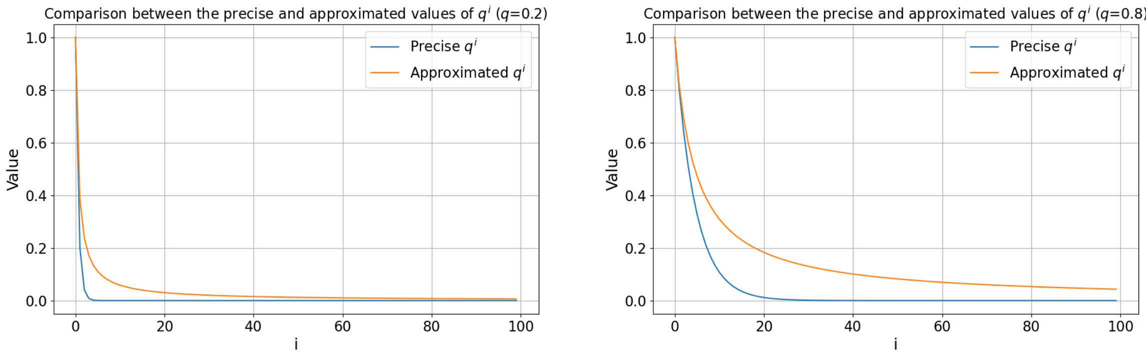

(a) $q = 0.2$

(b) $q = 0.8$

Figure 2: Comparison between the exact value of q^i and the approximate expression in (11).

We then take the derivative of the function in (12) and set it to zero:

$$\hat{f}'(i) = \frac{p_0 \ln q \left(-2i + 2Bp_0 + i^2(-1+p_0)\ln q(i \ln q - 3)\right)}{B(1 - i\ln q)^3} = 0. \tag{13}$$

The solution to (13), $\hat{i}$, is a critical point of $\hat{f}$:

$$\hat{i} = \frac{1}{\ln q} + \frac{1-3p_0}{3^{\frac{1}{3}}A} - \frac{A}{3^{\frac{2}{3}}(p_0-1)(\ln q)^2}$$

where $A = \ln q \cdot [9(1-p_0)^2 p_0(B\ln q - 1) - \sqrt{3(p_0-1)^3(1-9p_0-27B(p_0-1)p_0^2(2-B\ln q)\ln q)}]^{\frac{1}{3}}$. (14)

It turns out that $\hat{i}$ is the only real solution to (13) and hence the only critical point of $\hat{f}$. Therefore, the minimizer of $\hat{f}$, in the feasible region, is either 1 (lower bound on i), $\hat{i}$, or b (upper bound on i). In practice, it is of course easy to check which of these values yields the minimum MSE. Then, this minimum provides an approximately optimal solution to the original continuous optimization problem in (2). In Section 7.3 we will provide a computational analysis of the accuracy of the Continuous Model (with and without the approximations introduced in this section) relative to the Integer Model.

In situations where AQIM believes that the approximations introduced in this section are valid, (14) provides a direct mapping from the values of the input parameters B, p_0, and r (captured in $q = 1 - r$) to the optimal number of boxes to inspect from each sampled container. Then, the optimal number of containers to sample is determined via the budget constraint as $n^* = \frac{B}{\hat{i}}$.

6 Case Study: Port of Long Beach

To obtain additional insights from our model and to demonstrate its applicability to AQIM's real-world operations, we parameterize a numerical case study of maritime cargo sampled at the Port of Long Beach. The Port of Long Beach, located in the Los Angeles metropolitan area in California, is one of the 25 busiest container cargo ports in the world (Port of Long Beach, 2023). Together with the neighboring Port of Los Angeles, the combined port complex ranks ninth in the world. The Port of Long Beach alone handles one out of every five containers that move through U.S. ports. Containers arriving at the Port of Long Beach represent cargo valued at $200 billion annually. Roughly 90% of the trade passing through the Port of Long Beach is between the U.S. and East Asia, with China, Vietnam, Thailand, South Korea, and Taiwan the top five trading partners (Port of Long Beach, 2023). Given its prominence in U.S. maritime cargo imports and exports, the Port of Long Beach is an ideal case study to focus on with our model.

To help us assign values to our model's input parameters, AQIM management provided us with all AQIM inspection data collected during the calendar year 2022 at the Port of Long Beach. Each data entry describes one maritime cargo container that AQIM randomly sampled at the port. For each sampled container, the data include information such as the vessel name, country of origin, commodity, number of boxes/units in the container, number of boxes/units inspected (based on the hypergeometric table), and whether an actionable pest was detected.

To further improve our understanding of AQIM operations at the Port of Long Beach, we conducted a site visit to the port (along with air cargo at Los Angeles International Airport) in August 2023. During this site visit, we had informative discussions with the local AQIM Coordinator, CBP officers who supervise agriculture activities, and agriculture specialists who inspect cargo for AQIM and enter the inspection data. We also observed several real and mock AQIM inspections that included the selection of boxes to inspect, the inspection of their contents, and data entry.

Using the data provided by AQIM and drawing on the knowledge gained during our site visit, we established the Port of Long Beach case study parameterization summarized in Table 5. In the paragraphs below, we elaborate on how we chose each of the parameter values. It is important to emphasize that our model entails some necessary simplifications and abstractions relative to real-world AQIM operations at the Port of Long Beach, and that the values of some parameters are difficult to pinpoint. Therefore, we will present results for sensitivity analyses on several parameters in Section 7.

Table 5: Parameter settings for the Port of Long Beach case study.

Parameter	**Description**	**Value**
$cost_c$	Fixed cost of sampling a container	0
$cost_b$	Cost of inspecting a box	1
B	Budget	232
b	Number of boxes in each container	1500
p_0	Probability that a container comes from a pest source	0.025
r	Probability that any box in a container from a "pest source" contains a pest	0.306
ρ	Correlation coefficient between any two boxes in one infected container	0.3

We set the $cost_c$ parameter to zero because we learned during our site visit that AQIM does not directly incur the cost of transporting a sampled container to the inspection warehouse, unloading all of its contents, and staging the boxes for inspection. In practice, a contracted logistics company performs these tasks and charges fees for its services to the consignee receiving the container shipment. In other words, the fixed cost of sampling a container is passed on to the importer as part of the cost of doing business and bringing their goods through U.S. customs. While sampling more containers has a downside in terms of AQIM's qualitative imperative of facilitating trade, it does not require AQIM to directly commit any resources. Even though we let $cost_c = 0$ in our reference case study parameterization, we conduct a sensitivity analysis to explore how incorporating a fixed cost per container changes the optimal sampling strategy.

Since $cost_c = 0$, we can simplify the budget constraint by setting $cost_b = 1$ and quantifying the budget B in units of boxes per month. Currently, the AQIM sampling protocol calls for eight containers to be sampled per month at the Port of Long Beach. Our site visit revealed that AQIM personnel inspect 29 boxes from the vast majority of containers that they sample, since they are large enough to correspond to the last row of the hypergeometric table (see Table 1). Therefore, we set $B = 8 \cdot 29 = 232$ to reflect the current availability of resources for AQIM sampling. By calibrating our budget constraint to match the current commitment of AQIM resources each month, our case study allows us to investigate whether AQIM could achieve better probability estimates without increasing its capacity to inspect boxes.

In reality, the number of boxes varies from one container to another, but one simplification that our model makes is to assume that all containers carry b boxes. Our empirical AQIM data from the Port of Long Beach includes the number of boxes in each sampled container. The mean value is roughly 1500 boxes per container, so we adopt this value for b. Note that, when calculating this mean, we only considered data entries where the unit used to measure the container's contents was "box/cartons" and not some other unit, such as "kilograms" for loose cargo.

In the empirical AQIM data from the Port of Long Beach, the fraction of all containers sampled in 2022 in which an actionable pest was detected is approximately 0.02. The way that we model random pest arrivals, in order to carry a pest, a container must come from a pest source. The reverse, however, is not necessarily true. A container from a pest source may not actually feature a pest, which happens if the "coin flips" that determine the pest statuses of individual boxes all yield no pest. Therefore, we set the probability that a container comes from a pest source at $p_0 = 0.025$ to be similar to, but slightly higher than, the empirical fraction of sampled containers in which AQIM detected a pest.

Recall from (7) that, given a fixed p_0, there is a one-to-one correspondence between values of r and ρ. So, we need only to assign a value to one of these parameters, and then the value of the other is fully determined by (7). We choose to first establish ρ, which is the correlation coefficient between the two random variables indicating whether two boxes in the same container feature pests. AQIM does not calculate ρ or collect the data necessary to estimate it empirically because it only records pest inspection outcomes at the container level, not at the box level. However, during our site visit, AQIM personnel suggested that there often appears to be some positive correlation, since with certain commodities, inspecting only a few boxes is sufficient to assess whether a pest is pervasive throughout the container or not. Based on their descriptions, we set $\rho = 0.3$ in our reference case study parameterization, but we vary its value in a sensitivity analysis to examine how the correlation affects the optimal sampling strategy. With $p_0 = 0.025$ and $\rho = 0.3$, (7) requires that we set $r = 0.306$.

7 Numerical Results

In this section, we mainly present and discuss numerical results obtained for our Port of Long Beach case study. Beyond merely reporting the optimal sampling strategy for our reference case study parameterization described in Section 6, we focus on scenario and sensitivity analyses that allow us to explore how the correlation coefficient (ρ) and the fixed cost per container ($cost_c$) affect the optimal strategy and its performance. Unless otherwise noted, the results that we report are determined using the Integer Model (1) in which i is restricted to be an integer. Toward the end of this section, we briefly compare the numerical solutions of the Integer Model to those of the Continuous Model with and without the approximations introduced for the analysis in Section 5.

7.1 Case Study Results: Scenario Analysis

Table 7 displays results obtained for five parameter settings that are all based on the reference parameterization from Table 5, but with different values of the correlation coefficient ρ (and, by extension, r). For all of the cases shown in Table 7, it is assumed that there is no fixed cost for sampling a container ($cost_c = 0$). Note that Case 3 is our reference parameterization with its assumption that $\rho = 0.3$. Case 1 represents the special case in which the box pest events are independent. In particular, for $\rho = 0$, we assume that the presence of a pest in a specific box is independent of the pest statuses of other boxes. In this case, equation (7) does not apply. Rather, we model this special "independent" case within the structure of our model by setting $p_0 = 1$ and interpreting r as the probability that any box features a pest. Then, the true probability that a container features a pest is given by $p = 1-(1-r)^b$. For Case 1 we thus calculate r by assuming that $p = 0.02$ and $b = 1500$. Case 2 assumes a weak (but non-zero) correlation. Cases 4 and 5 assume stronger correlations. For each case, Table 7 reports the current AQIM protocol (essentially just the strategy defined in the AQIM Handbook) and its associated MSE, our model's optimal sampling protocol and the MSE that it achieves, and the percentage reduction in MSE achieved by following the optimal strategy instead of the current AQIM protocol.

Table 7: Case study results for the reference parameterization and different values of ρ.

Case#	Parameters		AQIM Protocol			Optimal Protocol			Improve%
	ρ	r	n	i	MSE	n^*	i^*	MSE	
1	$\rho = 0$	$r = 1.34 \times 10^{-5}$	$n = 8$	$i = 29$	0.00044	$n^* = 46.4$	$i^* = 5$	0.00039	9.96%
2	$\rho = 0.02$	$r = 0.0205$	$n = 8$	$i = 29$	0.0015	$n^* = 46.4$	$i^* = 5$	0.0005	64.08%
3	$\rho = 0.3$	$r = 0.305$	$n = 8$	$i = 29$	0.033	$n^* = 116$	$i^* = 2$	0.0002	99.28%
4	$\rho = 0.5$	$r = 0.506$	$n = 8$	$i = 29$	0.083	$n^* = 116$	$i^* = 2$	0.0002	99.76%
5	$\rho = 0.8$	$r = 0.804$	$n = 8$	$i = 29$	0.189	$n^* = 232$	$i^* = 1$	0.0001	99.94%

We make several important observations based on the results in Table 7. In all five cases, the optimal sampling strategy samples many more containers – but inspects fewer boxes per container – than the current AQIM sampling protocol. The ability for AQIM to estimate container pest probabilities more accurately by moving its sampling strategy in this direction from current practices thus appears to be robust to the particular value of ρ. The difference between the current protocol and the optimal strategy grows with ρ. This behavior is logical when considering the extreme case of $\rho = 1$, where the optimal strategy would intuitively be to set $i = 1$ (because observing whether one box has a pest reveals with certainty whether all boxes in the container do or do not have a pest) and sample as many containers as the budget allows. The percentage improvement in the MSE of the container pest probability estimate is also larger when ρ is higher. In the independent case with $\rho = 0$, the optimal protocol is only 6.11% more accurate than the current one. However, even a weak correlation of $\rho = 0.02$ is sufficient to make the optimal protocol achieve a 64.08% lower MSE than the current strategy. At the reference ρ value of 0.3 and above, our model's optimal strategy reduces the MSE by more than 99%.

Clearly, it would be helpful for AQIM to measure ρ by collecting box-level (as opposed to just container-level) inspection data in order to determine the best sampling strategy. But simply knowing that $\rho > 0$, which our intuition and our conversations with AQIM personnel certainly suggest, provides a strong justification for sampling more containers but fewer boxes per container in order to significantly improve probability estimates. In essence, knowing that the pest statuses of boxes in the same container are positively correlated means that an inspector does not need to inspect very many pest-free boxes in order to be fairly confident that the container does not feature a pest. Additional inspection time and effort can then be allocated to sampling additional containers.

Table 8: Case study results for the reference parameterization, but with $cost_c = 3$, and different values of ρ.

Case#	Parameters		AQIM Protocol			Optimal Protocol			Improve%
	ρ	r	n	i	MSE	n^*	i^*	MSE	
1	$\rho = 0$	$r = 1.34 \times 10^{-5}$	$n = 8$	$i = 29$	0.00044	$n^* = 36.57$	$i^* = 4$	0.00040	9.91%
2	$\rho = 0.02$	$r = 0.0205$	$n = 8$	$i = 29$	0.0016	$n^* = 36.57$	$i^* = 4$	0.0006	62.68%
3	$\rho = 0.3$	$r = 0.305$	$n = 8$	$i = 29$	0.035	$n^* = 51.2$	$i^* = 2$	0.0004	98.85%
4	$\rho = 0.5$	$r = 0.506$	$n = 8$	$i = 29$	0.083	$n^* = 64$	$i^* = 1$	0.0003	99.60%
5	$\rho = 0.8$	$r = 0.804$	$n = 8$	$i = 29$	0.175	$n^* = 64$	$i^* = 1$	0.0003	99.80%

Table 8 is analogous to Table 7, but for cases with a fixed cost per container of $cost_c = 3$ instead of zero as in the reference parameterization. While we explained our justification for setting $cost_c = 0$ for the reference parameter setting in Section 6, it is still worthwhile to investigate the impact of incorporating a positive fixed cost for sampling a container. There may be a small fixed cost at present for the time that the inspector requires to enter the inspection results, done once for each container. Furthermore, given AQIM's desire to minimize disrupting trade, it could choose to internalize the logistics costs of transporting, unloading, and staging each container for inspection, even though it does not incur these costs directly. Lastly, our model is general enough to apply to other sampling processes beyond our AQIM application, and other applications could entail fixed costs.

The patterns in Table 8 are more or less the same as those seen in Table 7 as far as the effects of the correlation coefficient ρ on the optimal sampling protocol and its improvement in the MSE. Once again, we find that it is optimal to sample more containers but inspect fewer boxes per container than in the current protocol, and that the difference between the current and optimal sampling strategies, and the difference in their MSEs, are larger when ρ is higher. Unsurprisingly, incorporating a fixed cost for sampling a container leads to optimal solutions with lower n^* and higher i^* values than in the corresponding cases without a fixed cost. However, even with the fixed cost included, our model indicates that the optimal sampling strategy involves sampling considerably more containers than the current AQIM approach. Therefore, the evidence for concluding that better probability estimates could be achieved by moving the sampling approach in this direction is quite strong.

7.2 Case Study Results: Sensitivity Analysis

In this subsection we continue to explore how the correlation coefficient (ρ) and fixed cost per container ($cost_c$) influence the optimal sampling strategy and its MSE. We conduct sensitivity analyses that consider a broader range of parameter values than those included in the cases in the previous subsection, and allow us to visualize the impacts of these parameters. Unless otherwise noted, all parameters are assumed to take on their reference values from Table 5.

Figure 3 visually depicts how the optimal n^* (blue bars, left y-axis) and i^* (red bars, right y-axis) vary with ρ when all other parameters are held at their reference values. Consistent with what we observed for the scenarios in the previous subsection, as ρ increases, n^* increases and i^* decreases. At correlation coefficients of $\rho = 0.6$ and above, it is optimal to inspect only one box per container. In other words, when box pest statuses are correlated this strongly, it is best to make the container sample size as large as possible and determine the pest status of each container based entirely on inspecting one box from the container.

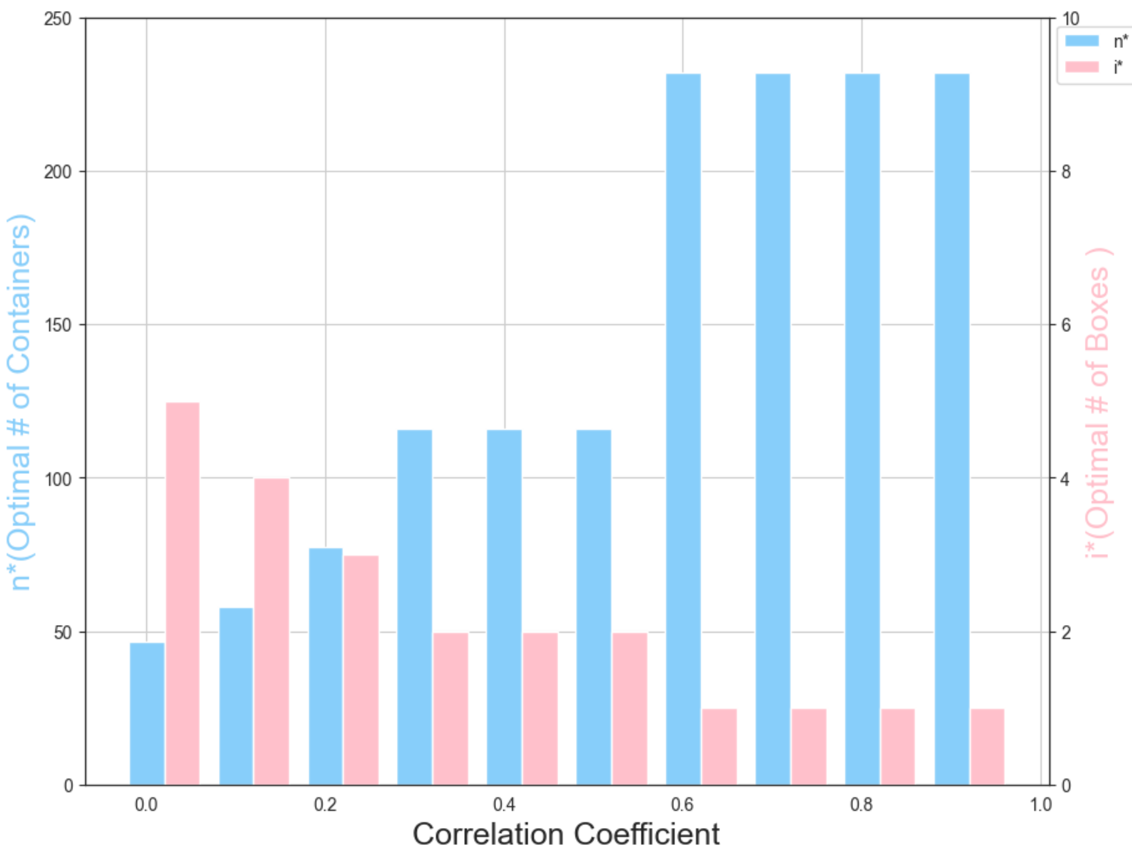

Figure 3: Optimal sampling strategy for the reference parameterization but with varying values of ρ.

Table 9 reports the MSEs of the container pest probability estimates obtained via the current AQIM protocol and our model's optimal sampling protocol. It also reports the ratio of the former to the latter. Once again, we see that the accuracy improvement that can be gained by switching from the current to the optimal sampling strategy is small when $\rho = 0$ but quickly increases with ρ. Even with a weak correlation of $\rho = 0.1$ the optimal strategy reduces the MSE by 95% compared to current practices. When $\rho = 0.3$ and greater, the improvement is above 99%. The stronger the correlation between box pest statuses, the more this information can be leveraged to inspect fewer boxes per container without sacrificing much container-level confidence, which means that more containers can be sampled.

Table 9: Comparing the MSEs achieved by the current AQIM protocol and the optimal protocol, and their ratio (AQIM MSE/Optimal MSE), for the reference parameterization with varying values of ρ.

ρ	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
AQIM MSE	0.00044	0.0078	0.0178	0.033	0.0572	0.083	0.112	0.145	0.179	0.219
Optimal MSE	0.00039	0.0004	0.0003	0.00026	0.0002	0.00019	0.00016	0.00013	0.0001	0.000097
Ratio	0.88	0.051	0.017	0.0079	0.0035	0.0023	0.0014	0.0009	0.00056	0.0004

Next, we investigate the sensitivity of our results to variation in the fixed cost of sampling a container ($cost_c$). Holding all other parameter values fixed at their reference settings (including $\rho = 0.3$), Figure 4 illustrates how n^* and i^* change as $cost_c$ increases from 0 to 9.

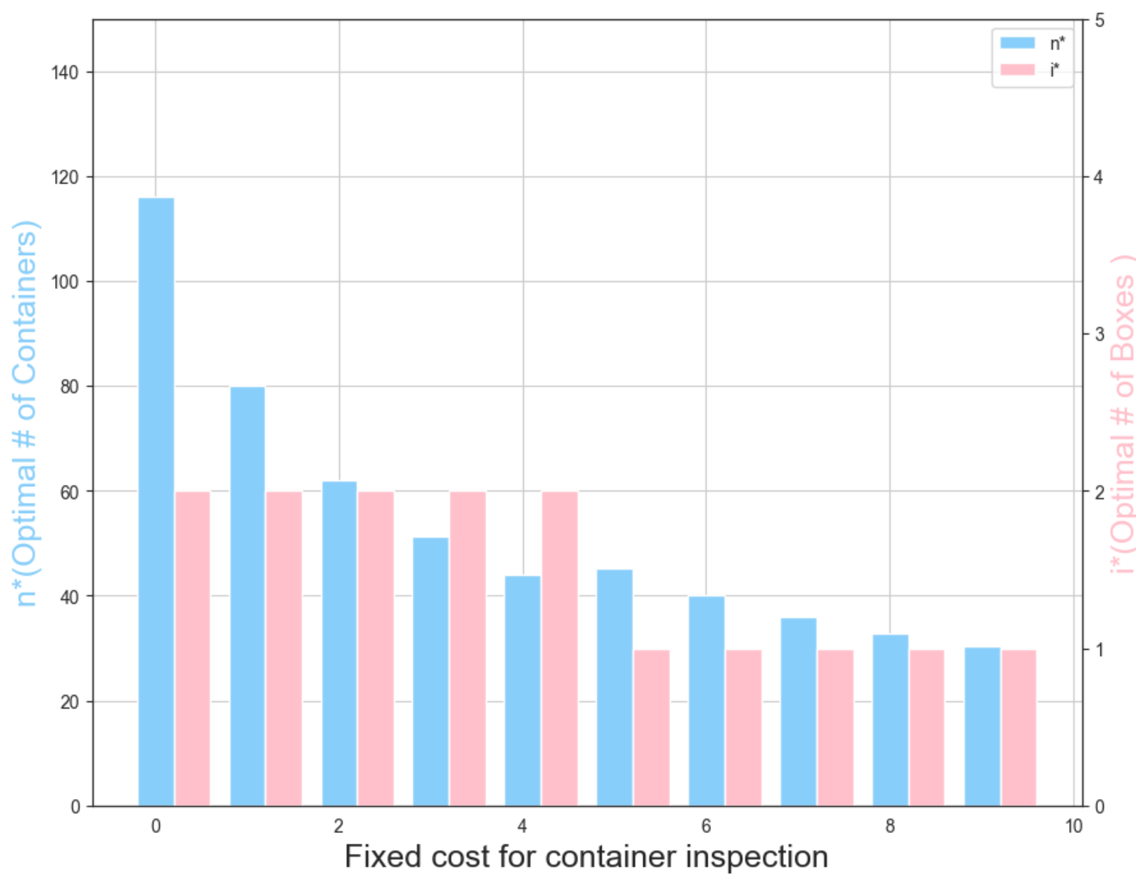

Figure 4: Optimal sampling strategy for the reference parameterization but with varying values of $cost_c$.

Intuitively, we see in Figure 4 that increasing the value of $cost_c$ makes the optimal solution sample fewer containers. A higher value of $cost_c$ makes it relatively more costly to sample a container than to inspect a box, which causes the optimal sampling strategy to shift to sampling fewer containers. However, it is crucial to point out that even with $cost_c = 9$ (i.e., sampling a container requires nine times the resources of inspecting a box), the optimal protocol still samples more containers than the current AQIM approach. Therefore, our finding that AQIM should sample more containers than it currently does in order to improve its probability estimates appears robust to incorporating a fairly large fixed cost per container.

What is less intuitive about the results in Figure 4 is that, when $cost_c$ increases, i^* either stays the same or decreases. Instead of moving in the opposite direction of n^*, which decreases with $cost_c$, to reflect the tradeoff between the two strategy levers, i^* moves in the same direction as n^*, if at all. This seemingly counterintuitive behavior can be understood by interpreting it in terms of the *substitution effect* and *income effect* from economics. The substitution effect based on the relative costs of sampling a container versus inspecting a box tends to decrease n^* and increase i^* as $cost_c$ increases. The income effect based on the real purchasing power of the budget B tends to decrease both n^* and i^* as $cost_c$ increases, since the latter effectively makes a fixed budget worth less in real terms. From Figure 4 we can see that the income effect evidently outweighs the substitution effect, resulting in i^* either remaining constant or declining with $cost_c$.

Table 10: Comparing the MSEs achieved by the current AQIM protocol and the optimal protocol, and their ratio (AQIM MSE/Optimal MSE), for the reference parameterization with varying values of $cost_c$.

$cost_c$	0	1	2	3	4	5	6	7	8	9
AQIM MSE	0.033	0.034	0.035	0.035	0.034	0.036	0.034	0.039	0.035	0.037
Optimal MSE	0.0002	0.0003	0.0003	0.0004	0.00041	0.00047	0.00050	0.00050	0.00054	0.00055
Ratio	0.006	0.009	0.008	0.011	0.012	0.013	0.014	0.012	0.015	0.013

Table 10 reports the MSEs associated with the current AQIM protocol and our model's optimal sampling strategy for different values of $cost_c$, with all other parameter settings held fixed at their reference values. Once again, we find that switching from the current approach to the optimal sampling protocol would lead to a significant reduction in the MSE in all cases. However, in contrast to what we observed for the correlation coefficient, the value of the fixed cost per container does not exert a qualitatively strong influence on the optimal strategy's MSE or on the percentage improvement between the current AQIM and optimal MSEs. Therefore, the ability of AQIM to improve its container pest probability estimates by implementing the optimal strategy, or at least moving in the direction of sampling more containers, seems quite robust to incorporating a container fixed cost.

7.3 Solutions Comparison

All of the Port of Long Beach case study results that we have presented up to this point were obtained using our Integer Model (1) of the problem. As we discussed in Section 4, this is in some sense the most realistic formulation because it restricts i to be an integer. For our case study, obtaining optimal solutions to the Integer Model was not computationally difficult, so its solutions are the ones that we focused on. However, to enable us to derive analytical insights in Section 5, we also considered the Continuous Model formulation (2), as well as a further simplified version of the Continuous Model with the approximations introduced in Section 5. Here, we briefly conclude this section on numerical results by comparing the optimal solutions obtained from the Integer Model, the Continuous Model, and the Continuous Model with approximations.

Table 11 reports the optimal i^* value for these three different solutions for four of the cases with different ρ values that were earlier included in Table 7. Before even comparing the particular values of i^* determined for each version of the formulation, it is important to note that all three solutions yield i^* values that are considerably lower than the current AQIM protocol's i value of 29, for all four cases. So, our finding that the optimal strategy samples more containers but inspects fewer boxes per container than the current AQIM approach holds regardless of which formulation is adopted. Comparing the solutions, the Continuous Model and Continuous Model with approximations i^* values are always within 0.5 units of the Integer Model i^* value, so that rounding their solutions would yield matches with the Integer Model solutions. The Continuous Model with approximations always produces an i^* value that is higher than that of the Integer Model; there is no such obvious trend for the Continuous Model i^* values. All in all, the optimal solutions obtained using the three subtly different formulations are close enough that the choice of formulation would be unlikely to affect any of the qualitatively important findings (both theoretical and numerical) that we have uncovered.

Table 11: Comparison of three different solutions for the reference parameterization and varying values of ρ.

Case #	Parameters		Integer Model i^*	Continuous Model i^*	Continuous Model with approximations i^*
	ρ	r			
1	0.02	0.0205	5	5.1	5.01
2	0.3	0.305	2	2.15	2.27
3	0.5	0.506	2	1.50	1.69
4	0.8	0.804	1	1	1.11

8 Conclusions

In this study, we formulated an optimization model to minimize the MSE of the container-level pest probability estimates calculated by the Agricultural Quarantine Inspection Monitoring (AQIM) program in the U.S. With limited resources, AQIM faces a tradeoff between sampling more containers but inspecting fewer boxes per container, or sampling fewer containers but inspecting more boxes per container. This is the central strategy decision that our model captures. We first conducted a theoretical analysis of the problem by introducing several approximations to aid tractability and deriving an analytical expression for the optimal sampling strategy. Then, we parameterized a numerical case study of sampling maritime cargo at the Port of Long Beach, with parameter values based on empirical AQIM data and a site visit. We presented and discussed case study results to investigate how critical parameters affect the optimal sampling strategy, how the optimal approach differs from the current AQIM protocol, and how much accuracy can be gained by shifting from current practices to the optimal approach. While our model is motivated by the AQIM application and we base our numerical case study on it, we emphasize that our model and insights could apply more generally to other probability estimation applications with two levels analogous to our containers and boxes.

Our main findings can be succinctly summarized as follows:

1. The optimal sampling strategies derived from our model tend to sample more containers, but inspect fewer boxes per sampled container, than the current AQIM protocol. This difference is quite robust to variations in parameter settings.

2. The optimal sampling strategies typically produce probability estimates with substantially lower MSEs than the current AQIM protocol.

3. The differences between the optimal sampling strategy and the current AQIM protocol, in terms of strategy and performance, are greater when the correlation between the pest statuses of boxes in the same container is higher.

Practically, we understand that it may not be possible for AQIM to implement the exact optimal sampling strategies recommended by our model. The values of the input parameters are uncertain, changing over time, and different from one POE and pathway to another. Some of our simplifying assumptions are departures from reality, such as the homogeneity of containers in terms of their number of boxes. Furthermore, AQIM might determine that increasing the number of containers that it samples to the full extent suggested by our results would conflict too much with its imperative to facilitate trade. However, despite these caveats, we strongly recommend that AQIM move in the direction of sampling more containers in order to achieve more accurate pest probability estimates, even if it comes at the expense of inspecting fewer boxes per container due to resource limitations. We certainly feel confident that the sampling strategy ought to shift in this direction.

Another practical recommendation that we offer AQIM based on our findings is that the program could gain a lot of valuable information, for relatively little effort, by recording pest inspection results at the box level instead of only doing so at the container level. This would merely require inspectors to enter whether each box has an actionable pest, yes or no. Using this data, AQIM could calculate the correlation coefficient ρ. As our analysis has shown, knowing ρ would go a long way toward helping AQIM optimize its sampling procedures and advocate for their implementation, given that stronger correlations make the probability estimation accuracy improvement larger.

Looking ahead, we see two future research directions as particularly fruitful. First, a natural extension of this study would be to optimize the allocation of inspection resources among multiple pest importation pathways instead of only one. From one common budget, the decision-maker must sample containers and inspect boxes from multiple pathways with different parameter values (e.g., probabilities that a container comes from a pest source, correlations between box pest statuses, etc.) with the goal of minimizing an aggregate error metric. This extension could help AQIM concentrate resources more or less heavily depending on the pathway, and update resource allocations as new data are collected. Second, we recognize that it is not necessarily sensible from a risk mitigation standpoint to estimate all probabilities equally accurately. For instance, to aid downstream AQI decision-making, it could be more helpful to estimate the container pest probability for pathways that have had the greatest pest risk in the past, that serve as conduits for particularly problematic pests, and so on. Therefore, it would be valuable to link a model of probability estimation to the downstream model for risk mitigation in order to gain insights into which probabilities are most important to assess accurately (and thus devote ample resources to estimating).

9 Acknowledgments

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 18STCBT00001 through the Cross-Border Threat Screening and Supply Chain Defense Center of Excellence (CBTS). The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

We thank Gregory Pompelli and Chris Scarmardo of CBTS for coordinating and managing this project. We also thank Andrea Sinclair of CBP and Edna Cintron-Velazquez of USDA APHIS PPQ for sharing their knowledge of AQIM, providing data, organizing our site visit to the Port of Long Beach, and commenting on our work while it was in progress.

References

Animal and Plant Health Inspection Service. (2017). *Report on Plant Protection and Quarantine (PPQ)* (Tech. Rep.). U.S. Department of Agriculture.

Bagchi, A., & Paul, J. A. (2014). Optimal allocation of resources in airport security: Profiling vs. screening. *Operations Research*, *62*(2), 219–233.

Bakshi, N., Flynn, S. E., & Gans, N. (2011). Estimating the operational impact of container inspections at international ports. *Management Science*, *57*(1), 1–20.

Batabyal, A. A., & Beladi, H. (2006). International trade and biological invasions: A queuing theoretic analysis of the prevention problem. *European Journal of Operational Research*, *170*(3), 758–770.

Billionnet, A. (2013). Mathematical optimization ideas for biodiversity conservation. *European Journal of Operational Research*, *231*(3), 514–534.

Büyüktahtakın, I. E., & Haight, R. G. (2018). A review of operations research models in invasive species management: State of the art, challenges, and future directions. *Annals of Operations Research*, *271*, 357–403.

Chen, C., Epanchin-Niell, R. S., & Haight, R. G. (2018). Optimal inspection of imports to prevent invasive pest introduction. *Risk Analysis*, *38*(3), 603–619.

Clauset, A., & Woodard, R. (2013). Estimating the historical and future probabilities of large terrorist events. *The Annals of Applied Statistics*, *7*, 1838–1865.

Epanchin-Niell, R. S., & Wilen, J. E. (2012). Optimal spatial control of biological invasions. *Journal of Environmental Economics and Management*, *63*(2), 260–270.

Food and Agriculture Organization. (2021). *Climate change fans spread of pests and threatens plants and crops: new FAO study.* Retrieved from `https://www.fao.org/newsroom/detail/Climate-change-fans-spread-of-pests-and-threatens-plants-and-crops-new-FAO-study/zh`.

Gale, F., Kee, J., & Huang, J. (2023). *How China's African swine fever outbreaks affected global pork markets* (Technical Report No. ERR-326). U.S. Department of Agriculture, Economic Research Service.

Jacobson, S. H., Kobza, J. E., & Nakayama, M. K. (2000). A sampling procedure to estimate risk probabilities in access-control security systems. *European Journal of Operational Research*, *122*(1), 123–132.

Kıbış, E. Y., & Büyüktahtakın, İ. E. (2017). Optimizing invasive species management: A mixed-integer linear programming approach. *European Journal of Operational Research*, *259*(1), 308–321.

Kim, B., Hong, S. C., Egger, D., Katsar, C. S., & Griffin, R. L. (2019). Predictive modeling and categorizing likelihoods of quarantine pest introduction of imported propagative commodities from different countries. *Risk Analysis*, *39*(6), 1382–1396.

Kovacs, K. F., Haight, R. G., Mercader, R. J., & McCullough, D. G. (2014). A bioeconomic analysis of an emerald ash borer invasion of an urban forest with multiple jurisdictions. *Resource and Energy Economics*, *36*(1), 270–289.

Makgoba, M., Tshikhudo, P., Nnzeru, L., & Makhado, R. (2021). Impact of fall armyworm (*spodoptera frugiperda*) (*j.e. smith*) on small-scale maize farmers and its control strategies in the Limpopo province, South Africa. *Jamba*, *13*(1), 1016.

Montgomery, K., Petras, V., Takeuchi, Y., & Katsar, C. S. (2023). Contaminated consignment simulation to support risk-based inspection design. *Risk Analysis*, *43*, 709–723.

Moore, J. L., Rout, T. M., Hauser, C. E., Moro, D., Jones, M., Wilcox, C., & Possingham, H. P. (2010). Protecting islands from pest invasion: Optimal allocation of biosecurity resources between quarantine and surveillance. *Biological Conservation*, *143*(5), 1068–1078.

National Honey Bee Health Stakeholder Conference Steering Committee. (2012). *Report on the National Stakeholders Conference on Honey Bee Health* (Tech. Rep.). Department of Agriculture (USDA), Environmental Protection Agency (EPA), United States Government (USG).

Olson, L. J. (2006). The economics of terrestrial invasive species: A review of the literature. *Agricultural and Resource Economics Review*, *35*(1), 178–194.

Port of Long Beach. (2023). *Port Facts & FAQs.* Retrieved from `https://polb.com/port-info/port-facts-faqs#facts-at-a-glance`

Sharov, A. A., & Liebhold, A. M. (1998). Model of slowing the spread of gypsy moth (Lepidoptera: Lymantriidae) with a barrier zone. *Ecological Applications*, *8*(4), 1170–1179.

Thompson, S. K. (2012). *Sampling* (Third ed.). Hoboken, New Jersey, U.S.: Wiley.

Trouvé, R., & Robinson, A. P. (2021). Estimating consignment-level infestation rates from the proportion of consignment that failed border inspection: Possibilities and limitations in the presence of overdispersed data. *Risk Analysis*, *41*(6), 992–1003.

Trouvé, R., & Robinson, A. P. (2024). Efficient border biosecurity inspection leverages superspreading to reduce biological invasion risk. *Risk Analysis*, *forthcoming*.

U.S. Department of Agriculture. (2019). *Qualitative assessment of the likelihood of African swine fever virus entry to the United States: Entry assessment* (Tech. Rep.). Author.

U.S. Department of Agriculture. (2021). Agricultural Quarantine Inspection Monitoring (AQIM) Handbook.

Yamamura, K., Katsumata, H., Yoshioka, J., Yuda, T., & Kasugai, K. (2016). Sampling inspection to prevent the invasion of alien pests: Statistical theory of import plant quarantine systems in Japan. *Population Ecology*, *58*, 63–80.