Risk Analysis, Vol. 38, No. 3, 2018 DOI: 10.1111/risa.12880

Optimal Inspection of Imports to Prevent Invasive Pest Introduction

Cuicui Chen,[1,*] Rebecca S. Epanchin-Niell,[2] and Robert G. Haight[3]

The United States imports more than 1 billion live plants annually—an important and growing pathway for introduction of damaging nonnative invertebrates and pathogens. Inspection of imports is one safeguard for reducing pest introductions, but capacity constraints limit inspection effort. We develop an optimal sampling strategy to minimize the costs of pest introductions from trade by posing inspection as an acceptance sampling problem that incorporates key features of the decision context, including (i) simultaneous inspection of many heterogeneous lots, (ii) a lot-specific sampling effort, (iii) a budget constraint that limits total inspection effort, (iv) inspection error, and (v) an objective of minimizing cost from accepted defective units. We derive a formula for expected number of accepted infested units (expected slippage) given lot size, sample size, infestation rate, and detection rate, and we formulate and analyze the inspector's optimization problem of allocating a sampling budget among incoming lots to minimize the cost of slippage. We conduct an empirical analysis of live plant inspection, including estimation of plant infestation rates from historical data, and find that inspections optimally target the largest lots with the highest plant infestation rates, leaving some lots unsampled. We also consider that USDA-APHIS, which administers inspections, may want to continue inspecting all lots at a baseline level; we find that allocating any additional capacity, beyond a comprehensive baseline inspection, to the largest lots with the highest infestation rates allows inspectors to meet the dual goals of minimizing the costs of slippage and maintaining baseline sampling without substantial compromise.

KEY WORDS: Acceptance sampling; biological invasions; biosecurity; inspection; live plant imports

1. INTRODUCTION

Acceptance sampling is a method of statistical quality control that helps decisionmakers accept or reject a lot, based on information obtained in a sample of individuals from the lot.(1) Although acceptance sampling is usually described in the context of manufacturing, it plays an important role in public safety programs by helping administrators attain the program quality standards. Acceptance sampling is used in programs to maintain food safety,(2–4) control grain purity,(5) monitor illegal drug use,(6,7) reduce the spread of animal diseases,(8) and prevent the introduction of damaging pests on agricultural imports.(9)

The literature on acceptance sampling in public safety programs focuses on single sampling plans for individual lots. A single sampling plan involves one sample of items or individuals from a lot. The plan is defined by the sample size and acceptance number, and the decision rule is to accept the lot only if the number of defective items in the sample is equal to or less than the acceptance number. The sampling plan depends on the size of the lot and is determined

[1]Harvard University, Cambridge, MA, USA.
[2]Resources for the Future, Washington, DC, USA.
[3]USDA Forest Service Northern Research Station, St. Paul, MN, USA.
*Address correspondence to Cuicui Chen, Harvard University, Cambridge, MA 02138, USA; cuicuichen@fas.harvard.edu.

to minimize costs,[6,7] attain an acceptable level of risk,[2,3,8–10] or attain desirable levels of cost and risk.[5][4]

Much of this literature implicitly assumes that each lot is inspected at a level specified to achieve the quality control objective. In practice, however, acceptance sampling often takes place in settings where many heterogeneous lots that vary in size and composition are received and inspected simultaneously with budget or capacity constraints that limit the overall level of sampling across all lots. For example, inspection of live plant imports at ports of entry is a prominent part of the system used by the U.S. Department of Agriculture, Animal and Plant Health Inspection Service's (USDA-APHIS's) Plant Protection and Quarantine program to protect U.S. agriculture and natural resources from unwanted and damaging pests.[9] Sometimes, the desired acceptance sampling plans are not employed (i.e., some lots may be sampled at less than desired levels) because the overall number of samples prescribed for lots received during a given day exceeds the capacity of personnel to perform the sampling. Work is needed on this broader problem of allocating resources to inspection sampling based on externally imposed resource constraints.[4,10]

We address the problem of allocating a targeted sampling budget among a set of lots that are received and inspected simultaneously and can differ in their sizes, composition, costs of sampling, and damages from defective items in accepted lots. The objective is to determine the single acceptance sampling plan for each lot to minimize the total damages caused by defective items across accepted lots subject to the budget constraint. We use the term "targeted sampling budget" to represent the portion of a total sampling budget that is available for acceptance sampling.[5] We first derive a formula for the expected number of accepted defective items in a lot—which we call expected slippage—given the lot size, sample size, defect rate, and detection rate (i.e., the likelihood of detecting a defective item if it is inspected). We examine how expected slippage varies with those four parameters. Then, we set up the inspector's optimization problem for allocating the targeted sampling budget among a set of incoming lots and examine the comparative statics of the optimal strategies.

Our analysis extends the theoretical foundation of single sampling inspection based on cost minimization. Hald[13] was among the first to develop and analyze cost functions for single acceptance sampling plans, including the cost of sampling, the cost of the expected number of accepted defective items, and the cost of a rejected lot. He derived a formula to determine the cost-minimizing sampling plan for a single lot given lot size, prior distribution of the number of defective items in the lot, and the cost parameters. Since Hald's paper was published, many studies have expanded our knowledge of the economics of acceptance sampling of individual lots,[14] including the sensitivity of optimal plans to changes in cost parameters and defect rates, the development and performance of optimal multistage sampling plans, and the effects of inspection on producer behavior. In recent work similar to ours, authors allocate a fixed inspection budget among lots or shipment pathways to minimize the expected number or cost of accepted nonconforming lots.[4,15,16] Our work departs from those models by deriving a formula for expected slippage (number of defective items in an accepted lot) and allocating the inspection budget to individual lots to minimize the aggregate cost of slippage across heterogeneous lots. Yamamura *et al.*[10] address a similar problem of allocating a fixed number of samples across lots to minimize the total number of defective items admitted in accepted lots; however, they use simplifying assumptions that both defective rate and sample proportion of each lot are small.

Having analyzed the inspector's optimization problem, we illustrate its application with an empirical analysis of the inspection of live plant imports in the United States. Live plant imports, including rooted plants, bulbs, roots, and unrooted cuttings, are a valuable retail commodity: the U.S. horticulture industry imports more than 1 billion live plants annually, with sales exceeding $600 million.[17] At the same time, live plant imports are an important and growing pathway for damaging nonnative pests of agriculture and natural resources, including

[4]While the literature on acceptance sampling in public safety programs focuses on single sampling plans, other types of sampling plans, including double sampling, multiple sampling, sequential sampling, and skip-lot sampling, have been developed and standardized (e.g., see *NIST/SEMATECH e-Handbook of Statistical Methods,* http://www.itl.nist.gov/div898/handbook/pmc/section2/pmc22.htm). For example, in the context of import inspection sampling for invasive species, Batabyal and Yoo[11] study a random inspection scheme based on the "continuous sampling" plan first formulated by Dodge.[12]

[5]This recognizes that sampling may also be conducted for other purposes, such as estimating defect rates of populations of imported commodities, which is an important input to the optimization model.

invertebrates and pathogens.[17–19] For example, damage from nonnative forest insects, many of which arrived on imported plants, costs U.S. taxpayers billions of dollars annually.[20,21] The nonnative forest pathogen, *Phytophthora ramorum*, has caused substantial mortality in coastal live oak and other tree species on the U.S. Pacific Coast.[22] Discovered in 2000, the pathogen likely arrived on imported nursery stock in the 1990s,[23,24] and projections suggest that an expanding infestation will warrant treatment, removal, and replacement of thousands of oak trees on residential property with annual costs exceeding $10 million.[25] The Australian light brown apple moth (*Epiphyas postvittana*) was discovered in wholesale nurseries in California in 2006 with evidence suggesting that it arrived on imported live plants.[26] An economic risk analysis suggested that damages to California's four main fruit crops, as well as quarantine and other costs, could reach $105 million annually.[27]

One means for reducing the introduction of pests on live plants is the inspection of shipments at ports of entry. Most imported plants enter the United States through one of 16 plant inspection stations, which are located near major international airports and seaports and managed by USDA-APHIS. If there is sufficient capacity, inspectors examine a number of selected sample plant units from each incoming lot. If a regulated pest or pathogen is detected in the sample units, inspectors may require that the lot be treated, returned, or destroyed. The process for determining the sample size of each lot is evolving.[6] Through 2012, inspections of live plant imports involved sampling approximately 2% of the items in each lot. These inspections were inconsistent and nonrandom. In 2013, a new, random sampling method, called risk-based sampling (RBS), was rolled out, and a version was fully implemented by late 2014. With RBS, lots are inspected according to hypergeometric sampling, and the sample size for each lot is determined to achieve a specified detection level, where the detection level is the minimum proportion of infested items in a lot for which the probability of detection is 95%. Currently, the detection level is 5% for all lots. Eventually, lots will be assigned to risk categories based on history of pest interceptions, potential damage of pests associated with the commodity, and other information. Lots with higher risk are anticipated to be assigned lower detection levels.[7]

The goal of our empirical analysis is to investigate the performance of sampling plans that minimize slippage given inspection capacity constraints relative to the performance of risk-based and proportional sampling plans. Minimizing slippage is an important objective because slippage is related to propagule pressure, and hence to the likelihood of establishment of pests. Invasion biology has shown that the more unwanted propagules released into the environment, the greater the chance of establishment.[28–32] While the number of infested plant units is not an exact measure of propagule pressure, it provides a useful proxy—all else equal about a lot, a greater number of infested plant units is likely to comprise a higher pest or pathogen load.

In our empirical analysis, we first introduce a maximum likelihood estimation procedure to estimate the infestation rate of a population of imported plants using historical inspection data and illustrate the procedure using data for 91 plant genera imported from Costa Rica via freight to Miami from 2010 through 2012. Second, for a set of lots arriving during a representative single day at the Miami plant inspection station, we calculate sample strategies that minimize total expected slippage across lots subject to a range of inspection capacity constraints, and we compare the performance of the slippage minimizing strategy with RBS and proportional (2%) sampling strategies, using comparable capacity constraints. We demonstrate the robustness of the sampling strategy's performance improvements to uncertainties in the daily sample composition and the infestation rate estimates and also consider an alternative optimization procedure that directly accounts for uncertainty in the infestation rate estimates. Third, recognizing that a minimum level of sampling is required to generate inputs to the optimization model and may be desired for information collection but may not be optimal under the slippage minimization goal, we analyze a sampling strategy where the decision-maker first allocates inspection resources according to the RBS rule and then allocates any unused inspection resources among lots to minimize total slippage.

[6] Robert L. Griffin, AQI Coordinator, USDA-APHIS, PPQ, 1730 Varsity Drive, Suite 400, Raleigh, NC 27606. Personal communication.

[7] Indeed, a small subset of live plant imports have already been categorized as very low risk and are subject to skip-lot sampling under the Propagative Monitoring and Release Program.

2. EXPECTED SLIPPAGE

In simple random acceptance sampling, a population is composed of countable individuals with a defect rate $\gamma \in (0, 1)$. The defect rate is the probability that a randomly drawn individual from the population is defective. A lot contains N individuals randomly drawn from the population. The number of defective individuals in a lot is a random variable that follows a binomial distribution parametrized by the defect rate and the lot size. The inspector selects a random sample of size $n \in \{1, 2, \ldots, N\}$ from the lot to inspect. Each individual in the sample is inspected with detection rate e, which is the probability of detecting a defect when the individual is in fact defective. The inspector adopts the zero-acceptance rule, which means she rejects the lot if and only if at least one individual in the randomly drawn sample is found to be defective. Her goal is to minimize the expected number of accepted defects in a given set of lots by choosing an optimal sample size for each lot, where lots may be drawn from different populations, subject to a capacity constraint that puts an upper limit on the total number of individuals that she can inspect.

We define *expected slippage* as the expected number of accepted defects in a lot. Let K be the number of defects in the lot and ξ be a binary variable indicating whether the lot is accepted ($\xi = 1$) or not ($\xi = 0$). Given a lot size N, sample size n, defect rate $\gamma > 0$, and detection rate e, the expected slippage is:

$$ES(N, n, \gamma, e) = \Pr(\xi = 1 | N, n, \gamma, e) \mathrm{E}[K | N, n, \gamma, e, \xi = 1] + \Pr(\xi = 0 | N, n, \gamma, e) \times 0, \tag{1}$$

which is the probability that a lot is accepted times the expected number of defects in the lot given that it was accepted. If the lot is rejected, no defects are accepted.

In Appendix A.2, we show that the formula for expected slippage is:

$$\begin{aligned} ES &= (1 - \gamma e)^n \cdot \left\{ \gamma(N - n) + \frac{1 - e}{1 - \gamma e} \cdot \gamma n \right\}, \\ &\equiv P \cdot \{D_1 + D_2\} \end{aligned} \tag{2}$$

where P is the probability of accepting the lot (i.e., none of the inspected samples were detected as defective), D_1 is the expected number of defects in the sublot that is not inspected, and D_2 is the expected number of defects in the sample conditional on acceptance. If $e = 0$, or detection rate is zero, the conditioning factor $\frac{1-e}{1-\gamma e} = 1$, and conditioning on acceptance does not affect the inspector's posterior about the number of defects in the sample, which is always accepted. When $e = 1$, or detection is perfect, $\frac{1-e}{1-\gamma e} = 0$, and the inspector's posterior about the number of defects in the accepted sample equals zero (because she has inspected every individual in the sample and detection is perfect). When $e \in (0, 1)$, the conditioning factor is decreasing in e and increasing in γ. Intuitively, the higher the detection rate, the more optimistic the inspector is about the number of defects in the sample she has accepted. The higher the defect rate, the less optimistic the inspector is because both the conditioning factor and γn are increasing in γ.

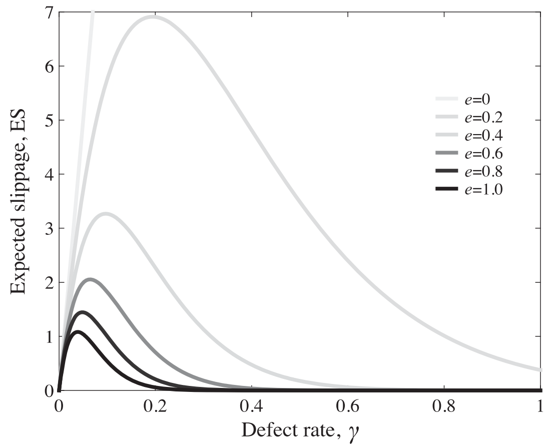

Fig. 1. Expected slippage versus defect rate under various detection rates. The lot size N is 100 and the sample size n is 25.

We can make several observations about the behavior of expected slippage as a function of its parameters. Expected slippage is increasing in lot size N, which only affects expected slippage via D_1, which is increasing in N. Expected slippage is not necessarily monotone in defect rate γ, which affects P (decreasing), D_1 (increasing), and D_2 (increasing) (Figs. 1 and 2). Expected slippage is decreasing in detection rate e, which affects expected slippage via P and the conditioning factor, both of which are decreasing in e (Figs. 1 and 3). Finally, expected slippage is decreasing in sample size n, which affects expected slippage via P and $D_1 + D_2$, both being decreasing in n (the latter is decreasing in n because the decrease in D_1 offsets the decrease in D_2) (Figs. 2 and 3).

3. OPTIMAL SAMPLING PROBLEM

Suppose the inspector receives a finite set of lots, J, for inspection. For each lot $j \in J$, let N_j, γ_j, and e_j be its size, defect rate, and detection rate,

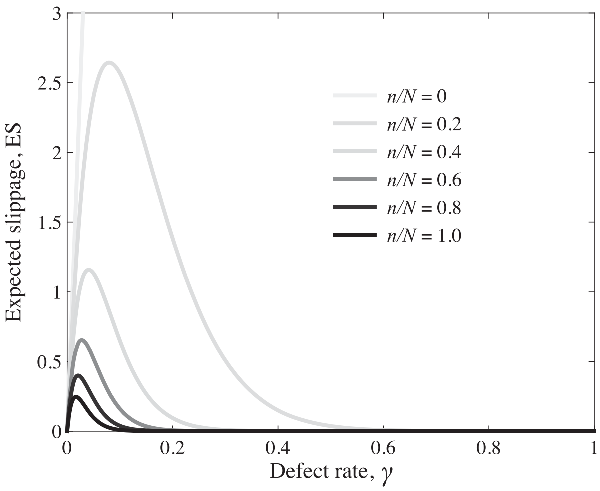

Fig. 2. Expected slippage versus defect rate under various sampling intensities. The lot size N is 100 and the detection rate e is 0.6.

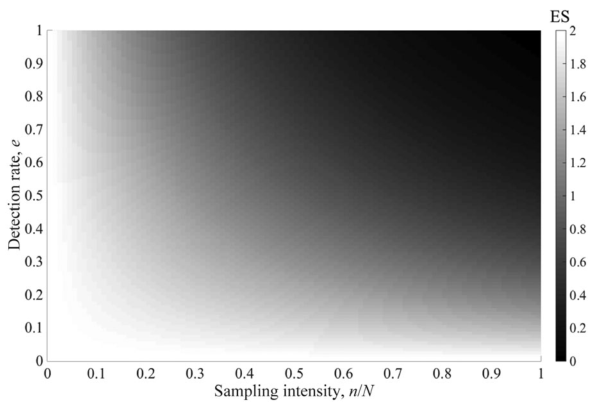

Fig. 3. Expected slippage under various sampling intensities and detection rates. The defect rate γ is 0.02, and the lot size N is 100.

respectively. Her goal is to minimize the expected damage from accepted defects (expected slippage) from J by choosing a sample size n_j for each lot j subject to an overall capacity constraint. We formulate the objective function as a weighted sum of the expected slippage of each lot. We interpret the weights $\kappa_j > 0$ for all $j \in J$ as measures of harm or cost associated with an accepted defective individual. We model the capacity constraint as restricting a cost-adjusted sum of the sample sizes of each lot, with the weights c_j for all $j \in J$ being the unit cost of inspecting an individual from lot j.[8] Let $\bar{c} > 0$ be the upper limit of the total cost, or the budget or capacity constraint.

The inspector's problem is as follows:

$$\min_{(n_j)\, j\in J} \sum_{j\in J} \kappa_j\, ES(N_j, n_j, \gamma_j, e_j)$$
$$\text{subject to}: \sum_{j\in J} c_j n_j \le \bar{c} \tag{3}$$
$$N_j \ge n_j \ge 0, \forall j \in J.$$

The Lagrangian for this constraint minimization problem is:

$$\sum_{j\in J} \kappa_j\, ES(N_j, n_j, \gamma_j, e_j) - \lambda\left(\bar{c} - \sum_{j\in J} c_j n_j\right) - \sum_{j\in J} \mu_j n_j - \sum_{j\in J} \nu_j (N_j - n_j). \tag{4}$$

The associated Kuhn–Tucker necessary conditions are:

$$\forall j \in J, \lambda = \frac{1}{c_j}\left[\mu_j - \nu_j - \kappa_j (1-\gamma_j e_j)^{n_j}\left\{ ln(1-\gamma_j e_j)\left[\gamma_j(N_j - n_j) + \frac{1-e_j}{1-\gamma_j e_j}\gamma_j n_j\right] - \frac{1-\gamma_j}{1-\gamma_j e_j}\gamma_j e_j\right\}\right]$$
$$\lambda \ge 0, \bar{c} \ge \sum_{j\in J} c_j n_j, \lambda\left(\bar{c} - \sum_{j\in J} c_j n_j\right) = 0 \tag{5}$$
$$\forall j \in J, \mu_j \ge 0, n_j \ge 0, \mu_j n_j = 0$$
$$\forall j \in J, \nu_j \ge 0, N_j - n_j \ge 0, \nu_j (N_j - n_j) = 0.$$

These conditions are also sufficient for optimality because the objective function is convex and the constraint is quasi-convex.[9]

We provide some intuition for the Kuhn–Tucker necessary conditions by considering an optimal interior solution (i.e., $N_j > n_j > 0, \forall j \in J$), such that for all $j \in J$, $\mu_j = \nu_j = 0$, and therefore:

[8] We note that the cost function could be expanded to include a fixed cost for each lot sampled, which would account for the costs of delivering and unloading selected lots. We do not include these fixed inspection costs in our model because in practice shipments of live plants are all delivered and unloaded for processing.

[9] The Hessian of the objective function, a diagonal matrix, has only positive eigenvalues, which are the diagonal elements, implying that the Hessian is positive semi-definite. The constraint is linear, which is a special case of quasi-convexity.

Table I. Comparative Statics for the Case with Two Shipments, Size N_1 and N_2, and Assuming an Interior Solution n_1* and n_2*

Parameter		Change in n_1*
Budget	$\bar{c}$	+
Cost of sampling lot 1	c_1	−
Cost of sampling lot 2	c_2	±
Cost of damage lot 1	κ_1	+
Cost of damage lot 2	κ_2	−
Size of lot 1	N_1	+
Size of lot 2	N_2	−
Infestation rate lot 1	γ_1	±
Infestation rate lot 2	γ_2	±
Detection rate lot 1	e_1	±
Detection rate lot 2	e_2	±

$$\forall j \in J, \lambda = \frac{1}{c_j}\left[\kappa_j(1-\gamma_j e_j)^{n_j}\left\{ln(1-\gamma_j e_j)\left[\gamma_j(N_j-n_j)+\frac{1-e_j}{1-\gamma_j e_j}\gamma_j n_j\right]-\frac{1-\gamma_j}{1-\gamma_j e_j}\gamma_j e_j\right\}\right] = \frac{1}{c_j}\kappa_j\frac{\partial ES(N_j, n_j, \gamma_j, e_j)}{\partial n_j}. \quad (6)$$

That is, at an optimal interior solution, sample sizes are selected to equalize the marginal damage- and cost-weighted expected slippage reduction across lots, so that gains cannot be made by shifting effort between lots. Higher per unit damage costs or lower marginal inspection costs justify lower marginal reductions in the expected slippage at the optimal interior solution. If marginal expected slippage reduction cannot be equated across all lots (i.e., the optimal solution is not interior), inspection capacity is allocated to the lots with higher marginal reductions, leaving some lot(s) inspected 100% and/or some lot(s) not inspected at all.

Additional intuition can be gained from the Kuhn–Tucker conditions by examining the comparative statics for sample allocation among two lots of size N_1 and N_2, assuming an interior solution n_1* and n_2*. Table I shows how n_1* changes with an increase in the parameters of each shipment (proofs are given in Appendix B). Many results are immediately intuitive, such as a higher budget leads to higher sampling intensity. Also, higher inspection cost or lower damage cost of a lot leads to lower sampling intensity of that lot. Sampling intensity also increases with the lot's size because expected slippage increases with lot size, thereby increasing the benefits of sampling. Similarly, the sampling intensity of a lot decreases as the other lot's damage cost and size increase because the benefits of sampling the other lot increase.

Note, however, that n_1* does not change monotonically with changes in γ, e, or c_2; instead, the direction of change of n_1* depends on the levels of these parameters. The intuition for the nonmonotonic response to a change in defect rate γ can be gained from Figs. 1 and 2, which show that expected slippage, given a fixed sampling intensity, tends to increase with increasing defect rate when the defect rate is low and decrease with increasing defect rate when the rate is high. As such, changes in a lot's defect rate can increase or decrease the marginal benefits of sampling, affecting optimal sampling intensity in two alternative directions.

The intuition behind the effects of detection rate e and other lot sampling costs c_2 on sampling intensity is different. An increase in detection rate e always decreases the expected slippage of a lot for a given sampling intensity, thereby decreasing the amount of sampling required to achieve the same expected slippage level. Thus, an increase in detection rate can be viewed as effectively freeing up some of the sampling budget, which can be used to further decrease expected slippage for whichever lot has the higher ratio of benefits to costs at the given levels of sampling and parameters.

Similarly, for some sets of parameter values, the sampling effort for lot 1 can increase or decrease as the cost of sampling the other lot (c_2) decreases. A decrease in sampling cost enables a larger total number of samples to be drawn for any given budget constraint. Thus, both lots could be sampled more intensively. However, the decrease in lot 2's costs could also increase the favorability of its ratio of marginal benefits to marginal costs from sampling, such that sampling of lot 1 could optimally decrease, while sampling of lot 2 increased.

4. DEFECT RATE ESTIMATION

The defect rate of a population can be estimated using maximum likelihood estimation based on random sampling inspection data of lots, assumed to have been drawn from a single population with defect rate γ (we discuss in Section 5.1 how one might define a population based on the variability of the defect rate and the richness of data). For each record i, we have sample size n_i and an indicator variable r_i for whether lot i was rejected (i.e., $r_i = 1$ if one or more defects are found in the sample).

The probability of accepting a lot is:

$$\Pr(r_i = 0|n_i, e, \gamma) = (1 - \gamma e)^{n_i}, \tag{7}$$

where γ is the defect rate to be estimated. Appendix A derives this formula by using the compound probability of two random events: first, some number of individuals are defective in the sample; and second, for each potential number of defects, the inspector fails to detect any defects. Alternatively, the following intuitive argument would do the job. The detection rate, e, is a scalar modifier on the defect rate, such that γe is the probability that a sampled individual is found to be defective. A lot is accepted if and only if all of the sampled individuals drawn are found not to be defective. Because each sample is an independent Bernoulli trial with the probability of being found not defective 1- γe, the probability of no samples being found defective is 1- γe raised to the power of the number of samples.

A consequence of the detection rate entering as a scalar modifier on the defect rate for the estimation is that the detection rate cannot be separately identified from the defect rate. For this reason, we assume certain values for the detection rate in our empirical analysis.

The likelihood of each observation is:

$$f(r_i|n_i, e, \gamma) = \begin{cases} (1 - \gamma e)^{n_i} & \text{if } r_i = 0 \\ 1 - (1 - \gamma e)^{n_i} & \text{if } r_i = 1 \end{cases} \tag{8}$$

It follows that the maximum likelihood estimation maximizes the log likelihood function:

$$\begin{aligned} L(\gamma) &\equiv \sum_{i \in I} \log f(r_i|n_i, e, \gamma) \\ &= \sum_{i \in I} \left[(1 - r_i) n_i \log(1 - \gamma e) \right. \\ &\quad \left. + r_i \log(1 - (1 - \gamma e)^{n_i}) \right]. \end{aligned} \tag{9}$$

5. APPLICATION TO LIVE PLANT IMPORT INSPECTION

In applying our framework to the inspection of live plant imports, we treat each combination of plant genera and exporting country as a unique commodity (we discuss in Section 5.1 ways to work with a finer definition of commodity to account for within-commodity variations in the infestation rates, given rich enough data sets). All lots of the same commodity are assumed to have been drawn from the same population with infestation (defect) rate γ, where the infestation rate is the probability that a plant unit will be infested with at least one pest or pathogen. We first estimate the plant infestation rates for each commodity in our application, and then apply our optimization framework to minimize expected slippage across imported lots.

5.1. Estimating Plant Infestation Rates

We use the maximum likelihood estimation procedure described in Section 4 to estimate infestation (defect) rates of plants of 91 genera imported from Costa Rica based on available inspection data (see Appendix C). The 91 genera represent live plants imported from Costa Rica via freight at the Miami plant inspection station in September 2010, focusing on genera that had at least one interception in Miami 2010 through 2012 (maximum likelihood estimation is feasible only if there is at least one "success" realization in the data set, which in this context is interception). The data for estimating the infestation rate for each genus include information for all shipments of the genus from Costa Rica to Miami for 2010 through 2012. These shipments were inspected following the 2% sampling guideline, which recommends that inspectors inspect roughly 2% of the shipment. Under this methodology, inspectors may inspect more or less than 2% of a shipment, and the specific sample is at the discretion of the inspector. From USDA-APHIS inspection records, for each lot, the lot size and whether or not the lot was found to be infested are known. We calculated the sample size of each lot assuming 2% inspection intensity. A detection rate of 80% is assumed, which is the rate that has been used by USDA-APHIS in its internal calculations of sample size under RBS and that it refers to as inspection efficiency. Under these assumptions, the estimated plant infestation rates vary from 0.888% for *Dendrobium* to 0.0002% for *Petunia* (Appendix C, Table C.I).

Because the actual sample size for each lot inspected is not known but estimated based only on the 2% sampling guideline, we conduct robustness analyses of the estimation procedure using Monte Carlo experiments (Appendix D). We find that our estimation approach is quite robust to random deviations in sample size relative to a fixed 2% proportional sample, but such deviations may lead to a slight downward bias in the estimated infestation levels (Appendix D, Table D.I). Yet, there are other data challenges that are not addressed by our robustness checks. Specifically, inspections under USDA-APHIS's 2% sampling guideline are inconsistent and nonrandom, with sample selection and deviations in sampled proportions determined nonrandomly by inspectors.

Another constraint to our empirical application is that we assume that all lots of a particular commodity are drawn from a single population with a fixed underlying infestation level. In reality, the underlying infestation rate of a commodity is likely to vary across years, seasons, and producers. If one had a rich enough inspection data set, the estimation method could be applied to subsets of data to obtain heterogeneous infestation rate estimates across these different dimensions. For example, to capture within-commodity year-to-year variability, the estimation method could be applied to each year's data to obtain yearly infestation rate estimates, though these likely would still be subject to out-of-sample prediction error. In our application, we estimate an average (across year, season, and producer) infestation rate for each commodity due to the limited number of years in the data and lack of producer information. As such, our empirical estimates obscure some of the interlot variability in underlying infestation rate within a commodity and should be considered illustrative.

5.2. Minimizing Expected Slippage with a Capacity Constraint

We now solve for sampling strategies that minimize slippage under various capacity constraints for a hypothetical set of shipments that arrive during one day at the Miami plant inspection station. In solving for the optimal sampling strategy, we set all the weights in the optimization problem (3) equal to one (i.e., $\kappa_j = c_j = 1$ for all lots $j \in J$). We use the AMPL-Knitro solver[(33)] and impose the additional constraint that the solution is integer-valued. We then compare these counterfactual sampling strategies and values with a proportional sampling strategy and the RBS scheme currently in use.

To construct the hypothetical set of shipments, we draw, with replacement, a random sample of lots received from Costa Rica and inspected in Miami in September 2010. The size of the sample is 39 lots, which is the average daily number of such lots received. The attributes of the lots are shown in Table E.I in Appendix E. The number of plants per lot varies from 10 to 240,000, with a total of 756,762 plants in 23 genera received. Estimated plant infestation rates vary from 0.148% for *Codiaeum* to 0.0004% for *Euphorbia.*

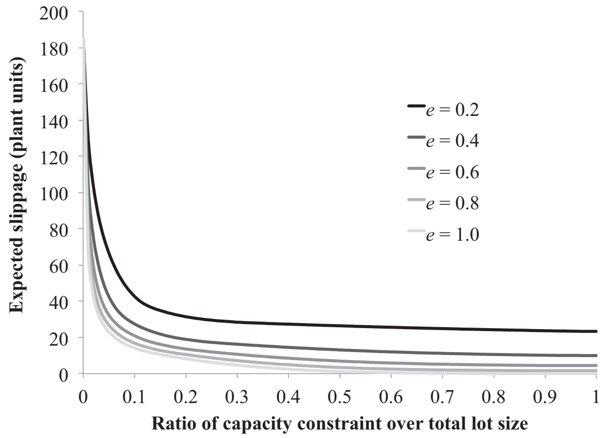

Fig. 4. Expected slippage across a range of capacity constraints for various detection rates under optimal sampling of the hypothetical set of lots received from Costa Rica and inspected in Miami (Appendix E, Table E.I). The expected slippage is calculated as the value of the objective function in the expected slippage minimization problem. The capacity constraint is expressed as the ratio of the total number of plant units inspected to the total number of plant units across all lots.

We solved the optimization problem for capacity constraints varying between 0 and the total number of imported plants (756,762) (Fig. 4). Minimized expected slippage decreases with the capacity constraint rapidly at first and then slowly, after the inspection capacity constraint exceeds about 10% of the total imported plant units. For any given capacity constraint, expected slippage decreases with increasing detection rate *e*.

We are interested in the performance of these slippage-minimizing sampling plans relative to the RBS plan and the proportional sampling plan. To be comparable, we computed the slippage-minimizing sampling plans using capacity constraints equal to the total number of samples taken under RBS (2,193 plants) and proportional (2%) sampling (15,143 plants). The RBS sample size for each lot detects a 5% infestation level with 95% confidence.[10] The

[10]In mathematics, the sample size for a size-N lot is the smallest integer n such that $\binom{n}{N - \lceil 0.05N \rceil} / \binom{n}{N} \leq 1 - 0.95$. We use MATLAB to generate the RBS sample size. In practice, PIS inspectors determine the RBS sample size by manually entering lot information into the plant inspection station (PIS) sampling tool at http://ports.cphst.org. The current protocol requires the PIS to select the medium button under part D of the sampling variable input section, where medium corresponds to a 5% infestation level.

proportional sampling sample size for each lot is the ceiling of 2% of the lot size.

The slippage-minimizing sampling strategies yield substantial reductions in expected slippage relative to the other sampling plans: up to 36% relative to the RBS plan and up to 61% relative to the proportional sampling plan (Table II). The reduction in expected slippage under optimal sampling increases with increasing detection rates.

The substantial reduction in expected slippage by our optimal sampling strategy is robust to alternative random samples of lots received from Costa Rica and inspected in Miami in September 2010. In addition, the reduction in expected slippage also holds when shipments are drawn from a distribution of smaller sized lots representative of shipments to non-Miami USDA-APHIS plant inspection stations in September 2010 from any country. We discuss the approach and results of these robustness checks in Appendix F.

How are the samples allocated among lots? Table III shows lot-specific sample sizes and expected slippage for the slippage-minimizing sampling plan and RBS. With an inspection capacity of 2,193 plants (the RBS total sample size), the slippage-minimizing plan prescribes allocating samples across only four of the largest lots with the highest plant infestation rates, whereas RBS prescribes sampling every lot. These differences in inspection allocation among lots reflect the differences in objectives of RBS and slippage minimization. Under RBS, all lots are inspected at the intensity required to detect an infested shipment with 95% probability, assuming that 5% of the plants are infested. In contrast, the slippage-minimizing sampling plan allocates inspections to minimize expected slippage, which is higher for larger lots with higher plant infestation rates. For example, under RBS, assuming a detection rate of 0.80, the *Codiaeum* lot with 36,800 plant units has a sample size of 74 units with a corresponding expected slippage of 49.81 units. For comparison, under the slippage-minimizing plan, the same lot has a sample size of 1,300 units with a corresponding expected slippage of 11.34 units, a reduction of 38 units. Under RBS, a *Codiaeum* lot with 4,000 units has a sample size of 73 units and corresponding slippage of 5.35 units. For comparison, under the slippage-minimizing plan, the same lot is not sampled and the expected slippage is only 0.57 units higher than under RBS.

The results are similar for an inspection capacity of 15,143 plants (the proportional sampling total sample size). The slippage-minimizing plan prescribes sampling only nine of the largest lots with the highest plant infestation rates, whereas the proportional sampling plan prescribes sampling 2% from all lots (Table III). We note that, as the targeted sampling budget increases, more lots are sampled with higher intensity under optimal sampling plans. For example, as the sampling budget increases from 2,193 (0.36% of the total number of plant units) to 15,143 (2% of the total number of plant units), the number of lots sampled increases from four to nine under the optimal sampling plans (Table III); Appendix F.1 shows that this pattern holds for the distributions as well.

While a highly concentrated sampling portfolio might seem intuitively undesirable, the concentrated sampling in Table III represents the optimal, slippage-minimizing sampling plan when the inspector is risk neutral and given our empirical assumptions. If the inspector were risk averse (caring about variance as well as expectation) or if the analysis accounted for additional uncertainty (e.g., from annual, seasonal, or producer-by-producer variation in infestation rate), then a more diversified inspection portfolio would be expected.

5.3. Sensitivity to the Uncertainty in the Infestation Rate Estimates

Our slippage minimization problem takes infestation rates as parameters. We do not observe the true underlying infestation rates, but rather use the infestation rate estimates from maximum likelihood estimation. In this section, we first investigate the extent to which the uncertainty in these estimates induces a *distribution* of minimized slippages. Second, we consider an alternative objective function that incorporates this uncertainty directly.

To evaluate the sensitivity of minimized slippage to the uncertain infestation rates, we simulate independent draws of infestation rates from their (asymptotic) sampling distributions obtained from the maximum likelihood estimation, calculate for each draw the slippage achieved by the original optimal sampling plan laid out in Table III, and take the average. To avoid drawing a negative infestation rate, we reparametrize the log likelihood function (9) by substituting $\frac{\exp(\theta)}{1+\exp(\theta)}$ for the infestation rate γ and obtain the maximum likelihood estimates of θ and their sampling distributions. Indeed, as γ ranges between 0 and 1, θ ranges from negative to positive

Table II. Comparison of Expected Slippage for Sampling Strategies that Minimize Expected Slippage and Strategies Based on Risk-Based Sampling and Proportional (2%) Sampling; Slippage-Minimizing Sampling Strategies are Computed with Capacity Constraints Consistent with Risk-Based Sampling (2,193 Plant Units) or Proportional (2%) Sampling (15,143 Plant Units)

	Inspection Capacity = 2,193		Inspection Capacity = 15,143	
Detection Rate	Slippage Minimization	Risk-Based Sampling	Slippage Minimization	Proportional (2%) Sampling
0.2	158.23	182.04	103.56	164.13
0.4	143.91	179.68	75.40	147.00
0.6	133.21	177.36	59.86	132.48
0.8	124.07	175.06	49.61	120.15
1.0	116.15	172.81	42.65	109.62

infinity. We then obtain independent draws of θ from their sampling distributions and convert each draw back to γ. Table IV reports the standard errors of the expected slippage under various detection rates with capacity constraints consistent with RBS sampling (2,193 plant units) and proportional (2%) sampling (15,143 plant units), respectively, using 100 independent draws of the reparametrized infestation rates θ. Table IV shows that the standard error of the expected slippage (across draws of θ) is only a small fraction of the minimized slippage, and thus the variation in θ has only a small impact on expected slippage. This indicates that, at least in our context, the performance of the optimal sampling plan derived using *point* estimates of infestation rates is quite robust to the distribution of infestation rates.

Next, we consider an alternative optimization procedure that directly takes the uncertainty into account; it minimizes the *expected* value of the expected slippage across a large number of infestation rate scenarios derived from the sampling distributions of the infestation rate estimates:

$$\min_{(n_j)j\in J} \sum_{s\in S} p_s \sum_{j\in J} ES_{js}(N_j, n_j, \gamma_{js}, e_j)$$
$$\text{subject to} : \sum_{j\in J} c_j n_j \leq \bar{c} \qquad (10)$$
$$N_j \geq n_j \geq 0, \forall j \in J,$$

where p_s is the probability of scenario s of plant infestation rate estimates, or the vector $(\gamma_{js})_{j\in J}$.

Thus, the inspector takes into account the uncertainty in the infestation rate estimates when he or she designs the sampling strategies. We report the values of this expectation-based objective function under various detection rates in Table IV. The associated sampling strategies can be found in Appendix G. The expected slippage and the optimal sampling plans are almost identical to those derived for the original optimization problem (3). The similarity in solutions is explained by the relatively small variances in the estimated means of the infestation rates. If the estimated means of the infestation rates had higher variances, we would have greater variation in the alternative realizations of the infestation rates, and the optimal sampling plan that directly accounts for this variability would have a more diversified allocation of limited sampling resources across lots. Similarly, accounting for out-of-sample prediction error or uncertainty arising from variation in infestation rates across years, seasons, and producers would lead to a more diversified portfolio. Nonetheless, sampling would be targeted at larger lots with higher underlying infestation rate estimates.

5.4. Optimization with Information-Gathering Constraint

In practice, inspection of shipments is important for not only intercepting infested shipments, but also for gathering information about plant infestation rates and providing incentives to importers to ship clean stock. These latter objectives are in tension with the prescription of the slippage-minimizing sampling plan, which focuses resources on only a few large shipments with high plant infestation rates. To incorporate these latter objectives, we consider a scenario in which all shipments are inspected with at least the RBS sample size, and any remaining sampling capacity is allocated to minimize expected slippage. We consider a total capacity constraint equal to the proportional (2%) sampling sample size (15,143). Results show that once the RBS inspection levels are satisfied, additional inspection capacity is allocated to the larger lots with higher plant infestation rates under the constrained expected-slippage

Table III. Comparison of Lot-Specific Sample Sizes and Expected Slippage for the Optimal Sampling and the Risk-Based Sampling Plans (Inspection Capacity of 2,193 Plant Units), and the Optimal Sampling and Proportional (2%) Sampling Plans (Inspection Capacity of 15,143 Plant Units)

Lot Attributes				Inspection Capacity = 2,193				Inspection Capacity = 15,143			
				Optimal Sampling		Risk-Based Sampling		Optimal Sampling		Proportional (2%) Sampling	
Plant Genus	Infestation Rate Percentage	Lot Size	Expected Slippage Without Inspection	Sample Size	Expected Slippage	Sample Size	Expected Slippage	Sample Size	Expected Slippage	Sample Size	Expected Slippage
Codiaeum	0.148	36,800	54.30	1,197	12.86	59	50.58	2,817	1.83	736	22.40
Codiaeum	0.148	7,506	11.08	–	11.08	59	10.27	1,452	1.68	151	9.12
Codiaeum	0.148	4,000	5.90	–	5.90	58	5.45	956	1.54	80	5.28
Codiaeum	0.148	1,250	1.84	–	1.84	57	1.66	260	1.13	25	1.76
Codiaeum	0.148	504	0.74	–	0.74	54	0.64	–	0.74	11	0.72
Dracaena	0.104	28,697	29.79	586	18.01	59	28.32	2,855	2.56	574	18.20
Dracaena	0.104	5,860	6.08	–	6.08	59	5.75	1,028	2.23	118	5.43
Dracaena	0.104	4,900	5.09	–	5.09	59	4.80	851	2.16	98	4.61
Dracaena	0.104	1,125	1.17	–	1.17	57	1.07	–	1.17	23	1.13
Dracaena	0.104	956	0.99	–	0.99	57	0.90	–	0.99	20	0.96
Dracaena	0.104	193	0.20	–	0.20	49	0.15	–	0.20	4	0.20
Schefflera	0.081	1,850	1.50	–	1.50	58	1.41	–	1.50	37	1.44
Cordyline	0.069	49,200	34.19	410	27.04	59	33.05	3,815	3.84	984	19.46
Cordyline	0.069	10,020	6.96	–	6.96	59	6.71	1,109	3.43	201	6.13
…	…	…	…	…	…	…	…	…		…	…
Total		756,762	184.44	2,193	124.07	2,193	175.06	15,143	49.61	15,143	120.15

Note: The detection rate is assumed to be 0.80. The "baseline expected slippage" in the "Lot Attributes" panel is the expect slippage without any inspection. The lots (rows) are organized in descending order of the estimates of plant infestation rates. All omitted rows have sample sizes of zero under optimal sampling but positive sample sizes under the risk-based sampling and proportional sampling plans.

Table IV. Standard Errors of the Minimized Slippage, and the Values of the Expectation-Based Minimized Slippage, Under Capacity Constraints Consistent with RBS (2,193 Plant Units) and Proportional (2%) Sampling (15,143 Plant Units), Respectively

	Inspection Capacity = 2,193			Inspection Capacity = 15,143		
Detection Rate	Slippage Minimization	Standard Error	Expectation-Based Minimization	Slippage Minimization	Standard Error	Expectation-Based Minimization
0.2	158.23	5.07	160.47	103.56	2.74	105.19
0.4	143.91	4.74	146.12	75.40	2.44	76.86
0.6	133.21	4.18	135.27	59.86	2.49	61.25
0.8	124.07	3.58	125.96	49.61	2.56	50.96
1.0	116.15	3.18	117.92	42.65	2.59	44.01

Note: Calculated using 100 independent draws of the reparametrized infestation rates based on their (asymptotic) sampling distributions from maximum likelihood estimation.

Table V. Comparison of Lot-Specific Sample Size and Expected Slippage for Plans Developed for Slippage Minimization and Slippage Minimization with a Risk-Based Sampling (RBS) Constraint

			Inspection Capacity = 15,143			
	Lot Attributes		Slippage Minimization		Slippage Minimization with RBS Constraint	
Plant Genus	Infestation Rate Percentage	Lot Size	Sample Size	Expected Slippage	Sample Size	Expected Slippage
Codiaeum	0.148	36,800	2,817	1.83	2,671	2.18
Codiaeum	0.148	7,506	1,452	1.68	1,316	2.01
Codiaeum	0.148	4,000	956	1.54	830	1.85
Codiaeum	0.148	1,250	260	1.13	153	1.39
Codiaeum	0.148	504	–	0.74	54	0.64
Dracaena	0.104	28,697	2,855	2.56	2,650	3.05
Dracaena	0.104	5,860	1,028	2.23	845	2.67
Dracaena	0.104	4,900	851	2.16	673	2.59
Dracaena	0.104	1,125	–	1.17	57	1.07
Dracaena	0.104	956	–	0.99	57	0.90
Dracaena	0.104	193	–	0.20	49	0.15
Schefflera	0.081	1,850	–	1.50	58	1.41
Cordyline	0.069	49,200	3,815	3.84	3,509	4.58
Cordyline	0.069	10,020	1,109	3.43	831	4.10
…	…	…	…	…	…	…
Total		756,762	15,143	49.61	15,143	52.90

Note: The inspection capacity is 15,143 plant units and the detection rate is 0.80. The lots (rows) are organized in descending order of the estimates of plant infestation rates. All omitted rows have sample sizes of zero under slippage minimization but positive sample sizes under slippage minimization with an RBS constraint.

minimization strategy (Table V). Although the constrained optimization produces a slightly larger expected slippage than if we do not impose the minimum RBS sample size, it still provides a substantial improvement over proportional (2%) sampling.

6. CONCLUSION

This article addresses an acceptance sampling problem that is commonly encountered in public and environmental safety programs such as reducing the introduction of damaging pests on agricultural imports. The problem is to allocate a limited inspection budget to multiple, heterogeneous lots to minimize the cost from damage caused by defective items in accepted lots. We derive a formula for expected slippage (i.e., number of accepted defects) under single acceptance sampling, as a function of the lot size, sample size, defect rate, and detection rate. Using the formula for accepted slippage, we formulate an optimization problem to allocate a fixed targeted sampling budget among a set of lots that can differ in their size, costs of sampling, damages from slippage, and the populations they are drawn from, with the objective of minimizing the total damages from slippage across all lots. From comparative statics analysis, we find that, for an interior solution, the optimal sample intensity of a given lot increases with larger lot size and higher damage cost and decreases with higher inspection cost. Further, sample intensity for a given lot decreases with higher other damage cost and larger other lot size. Sample intensity, however, is nonmonotonic with respect to defect rate, detection rate, and other lot sampling cost.

We apply our optimization formulation to the problem of sampling shipments of live plant imports to the United States with the goal of reducing the introduction of plant pests. First, we derive and implement a procedure to estimate the infestation rate of plant populations based on historical sampling outcomes. Using the estimated plant infestation rates and an example of daily shipments received at the Miami plant inspection station, we find that slippage-minimizing sampling strategies target the subset of largest lots with highest plant infestation rates (i.e., lots with the highest expected slippage). Further, those strategies yield substantial reductions in expected slippage relative to strategies in which a fixed proportion of plant units in each lot is sampled or the sample size for each lot is determined to achieve a specified detection level, where the detection level is the minimum proportion of infested items in a lot (e.g., 5%) for which the probability of detection is 95%—two sampling strategies representative of those that have been or are currently being used by USDA-APHIS, which conducts the inspections. These improvements are robust to alternative examples of daily shipments, as well as alternative infestation rate scenarios based on sampling from infestation rate distribution estimates. We also propose an alternative optimization procedure that directly accounts for the uncertainty in the infestation rate estimates when minimizing slippage. Recognizing that inspections provide benefits beyond interception of infested shipments, including information about infestation rates and incentives to importers to ship clean stock, we also consider an inspection scenario in which all shipments are sampled at least at a baseline level. We identify how any additional sampling capacity, beyond a baseline survey effort, could be allocated to lots to minimize total expected slippage. We find that allocating additional sampling capacity to the largest shipments with highest plant infestation rates can substantially reduce expected slippage relative to the baseline sampling effort. This result suggests that the dual goals of minimizing the number of accepted infested plants and maintaining a comprehensive baseline sampling effort could be met without substantial compromise when additional sampling capacity is available. This approach would also allow continued data collection for updating estimates of infestation rates for implementing our methodology.

Several caveats affect our numerical estimates of defect rate, as described in Sections 4 and 5.1. First, the detection rate and the infestation rate are not separately identified by the maximum likelihood estimation. Therefore, our estimates are obtained under an assumed value of the detection rate e. Second, because USDA-APHIS's inspection data do not record the actual number of samples inspected, our estimated infestation rate estimates may be biased slightly, as suggested by the Monte Carlo experiments in Appendix D. Third, under the 2% sampling guideline formerly used by USDA-APHIS and the source of data for our estimations, inspectors could select their samples nonrandomly based on their knowledge of the commodity and other factors. Detection rates may also vary across inspectors and lots. Finally, there can be annual, seasonal, and cross-producer variation in infestation rates that could be accounted for in estimation given a rich enough data set. For these reasons, the numerical estimates of defect rate presented in this article are largely illustrative. It is worth noting, however, that the data currently being collected under USDA-APHIS's RBS strategy are based on random sampling and follow a detailed protocol. Thus, these data should be more amenable to estimating plant infestation rates using our maximum likelihood estimation technique, particularly if they record the number of plant units inspected. Unfortunately, these data were not available at the time of these analyses. Future work should also seek to develop empirical methods that better represent the nonreducible uncertainties in infestation rates that result from natural variability in pest populations across time and space.

Our empirical estimation and optimization analyses highlight the value of inspection data for informing risks and targeting inspections across commodities. However, these data for live plant imports are rare outside the United States and incomplete within the United States.[34] Our findings support the value of collecting data on both random sampling inspection outcomes and sample sizes for estimating infestation rates, as collection of such data would improve overall efforts to reduce accidental importation of pests.

Our goal of minimizing expected slippage, or the expected number of accepted defective *individuals*, is an alternative to that of minimizing the expected number of accepted defective *lots*, which is the objective of some applications of acceptance sampling in public safety programs.[4,15,16] Even though our approach instructs inspectors to focus on larger lots with higher plant infestation rates (Table III), we expect that solutions that minimize the number of accepted defective lots will allocate sampling effort to lots with higher infestation rates with less regard to lot size. For example, in our case with an inspection capacity of 2,193 (Table III), only the two largest among five *Codiaeum* lots (which has the highest estimated infestation rate) are sampled (and rather intensively) under the objective of minimizing expected slippage. In contrast, all five *Codiaeum* lots are sampled almost evenly under an objective of minimizing the expected number of accepted defective lots. This makes intuitive sense because conditional on the same defect rate, lot size has a larger effect on the number of expected infested individuals than on the probability that a lot is infested, across most lot sizes relevant to live plant imports.

The preference to minimize accepted defective *lots* versus damages from accepted defective *individuals* may vary across public policy problems. For live plant import inspections, optimizing based on the number of accepted defective (i.e., infested) plants is most relevant because the number of introductions of a pest into the environment is a key predictor of establishment.[28–32] We note that the number of infested plant units is not an exact measure of propagule pressure because the number of pests may vary from plant to plant within a lot and across lots. If estimates of infestation rates and number of propagules per plant are available for each pest and plant taxa, our formulation could be refined to calculate a separate expected slippage estimate for each pest and lot, and the objective function would choose the optimal sample size for each lot to minimize the total damage from expected slippage across all pests and lots.

Our optimal sampling problem also allows for refinement of the damage cost, κ_i, of accepted infested individuals across commodities dependent on the level of damages expected from establishment of pests associated with a commodity. For example, pests on a plant genus that is more closely related to a commercially valuable plant in the importing country may receive a higher damage weight. Alternatively, the damage cost parameters could be replaced by a function that captures nonlinear relationships between the number of accepted infested plants and the likelihood of pest establishment. Parameterization of these damage cost parameters or functions is an important future research area and will require drawing on expertise of entomologists, plant pathologists, and economists. Finally, if the inspector were also concerned with welfare losses associated with intercepting an infested shipment, the objective function could be expanded to account for these values as well.

Several other extensions of this acceptance sampling problem are worthy of future research. First, the inspection problem could be posed in a dynamic framework to allow for the termination of sampling of a lot once a defective individual is found or for learning about infestation rates as lots are sampled. Optimal sampling with learning involves allocating inspection effort among lots to maximize the expected benefits from slippage reduction, accounting explicitly for the benefits from inspection for reducing uncertainty about risks across commodities and hence dynamically improving inspection targeting. Springborn[16] studies a dynamic optimization problem of inspection where the probability of a lot's being infested is uniform across all lots from a single producer and the goal of sampling is to maximize the discounted sum of the expected number of averted infested lots under a capacity constraint. The probability of infestation of a lot from a given producer is updated over time as lots are sampled. In contrast to Springborn's[16] assumption that lots from a given producer have an equal probability of being infested and binary decision of whether or not to sample a given lot, our model treats the population infestation rate as the primitive from which to derive a specific expected level of infestation for each lot and allows a different number of individuals to be sampled from each lot. Combining the approach of our static model with a dynamic optimization approach and Bayesian learning regarding population infestation rates could further improve inspection efficacy. The Bayesian learning approach also overcomes the inability of maximum likelihood estimation to estimate the infestation rates of commodities for which there has been no interception; a Bayesian inspector would have some prior belief on the possible infestation rate, and update such belief given any observation, with or without interception.

Finally, accounting for strategic interactions between the inspector and the producers of inspected lots would enhance the realism of optimal acceptance sampling design. In the context of U.S. live plant imports, USDA-APHIS has communicated to us that live plant exporters adapt to changes in inspection policies in various innovative ways (e.g., mixing infested material with clean material in response to a more stringent inspection policy). Similarly, exporters could adjust infestation rates or lot sizes in response to inspection policies in order to minimize their losses from detected infested shipments. Ameden *et al.*[35,36] and Springborn *et al.*[37] have considered exporter responses in the contexts of uniform inspections and risk-based inspections. Anticipating and accounting for the response of regulated firms in the design of optimal acceptance sampling will be an important focus for future research and policy design and could be incorporated into the slippage minimization framework developed here.

ACKNOWLEDGMENTS

This work is the product of the "Globalization of the Live Plant Trade: Informing efficient strategies for reducing nonnative pest invasion risk" working group (Co-PIs Epanchin-Niell and Liebhold) that was supported by the National Socio-Environmental Synthesis Center (SESYNC), NSF award DBI-1052875. The authors gratefully acknowledge working group participants for their ideas and feedback; Bob Griffin and Cathy Sue Katsar for providing invaluable data support, knowledge of APHIS-PPQ inspections, and article feedback; Mark Powell, Mike Springborn, René Eschen, Eric Lichtenberg, Lynne Westphal, and an anonymous reviewer for their very helpful comments on the article; and Che-Lin Su for guidance and discussions on the optimization tool. Any and all errors and perspectives are the responsibility of the authors.

REFERENCES

1. Schilling EG, Neubauer DV. Acceptance Sampling in Quality Control, 2nd ed. Boca Raton, FL: Chapman & Hall/CRC Press, 2009.
2. Starbird SA. Moral hazard, inspection policy, and food safety. American Journal of Agricultural Economics, 2005; 87(1): 15–27.
3. Whiting RC, Rainosek A, Buchanan RL, Miliotis M, LaBarre D, Long W, Ruple A, Schaub S. Determining the microbiological criteria for lot rejection from the performance objective or food safety objective. International Journal of Food Microbiology, 2006; 110(3):263–267.
4. Powell MR. Optimal food safety sampling under a budget constraint. Risk Analysis, 2014; 34(1):93–100.
5. Kobilinsky A, Bertheau Y. Minimum cost acceptance sampling plans for grain control, with application to GMO detection. Chemometrics and Intelligent Laboratory Systems, 2005; 75(2):189–200.
6. Baker JR, Lattimore PK, Matheson LA. Constructing optimal drug-testing plans using a Bayesian acceptance sampling model. Mathematical and Computer Modelling, 1993; 17(2):77–88.
7. Lattimore PK, Baker JR, Matheson LA. Monitoring drug use using Bayesian acceptance sampling: The Illinois experiment. Operational Research, 1996; 44(2):274–285.
8. Christensen J, Gardner IA. Herd-level interpretation of test results for epidemiologic studies of animal diseases. Preventive Veterinary Medicine, 2000; 45(1):83–106.
9. Venette RC, Moon RD, Hutchison WD. Strategies and statistics of sampling for rare individuals. Annual Review of Entomology, 2002; 47(1):143–174.
10. Yamamura K, Katsumata H, Yoshioka J, Yuda T, Kasugai K. Sampling inspection to prevent the invasion of alien pests: Statistical theory of import plant quarantine systems in Japan. Population Ecology, 2016; 58:63–80.
11. Batabyal AA, Yoo SJ. A theoretical analysis of random inspections and fines in invasive species management. Economic Bulletin, 2008; 17(15):1–9.
12. Dodge HF. A sampling inspection plan for continuous production. Annals of Mathematical Statistics, 1943; 14:264–279.
13. Hald A. The compound hypergeometric distribution and a system of single sampling inspection plans based on prior distributions and costs. Technometrics, 1960; 2(3):275–340.
14. Wetherill GB, Chiu WK. A review of acceptance sampling schemes with emphasis on the economic aspect. International Statistical Review, 1975; 43(2):191–210.
15. Surkov IV, Alfons GJM, Lansink O, Van der Werf W. The optimal amount and allocation of sampling effort for plant health inspection. European Review of Agricultural Economics, 2009; 36(3):295–320.
16. Springborn MR. Risk aversion and adaptive management: Insights from a multi-armed bandit model of invasive species risk. Journal of Environmental Economics and Management, 2014; 68(2):226–242.
17. Liebhold AM, Brockerhoff EG, Garrett LJ, Parke JL, Britton KO. Live plant imports: The major pathway for forest insect and pathogen invasions of the US. Frontiers in Ecology and the Environment, 2012; 10(30):135–143.
18. Kenis M, Rabitsch W, Auger-Rozenberg M-A, Roques A. How can alien species inventories and interception data help us prevent insect invasions? Bulletin of Entomological Research, 2007; 97(5):489–502.
19. Smith RM, Baker RHA, Malumphy CP, Hockland S, Hammon RP, Ostojá-Starzewski JC, Collins DW. Recent nonnative invertebrate plant pest establishments in Great Britain: Origins, pathways, and trends. Agricultural and Forest Entomology, 2007; 9(4):307–326.
20. Aukema JE, Leung B, Kovacs K, Chivers C, Britton KO, Englin J, Frankel SJ, Haight RG, Holmes TP, Liebhold AM, McCullough DG, Von Holle B. Economic impacts of nonnative forest insects in the continental United States. PLoS One, 2011; 6(9):e24587.
21. Epanchin-Niell RS, Liebhold AM. Benefits of invasion prevention: Effect of time lags, spread rates, and damage persistence. Ecological Economics, 2015; 116:146–153.
22. Meentemeyer RK, Rank N, Shoemaker DA, Oneal CB, Wickland AC, Frangioso KM, Rizzo DM. Impact of sudden oak death on tree mortality in the Big Sur ecoregion of California. Biological Invasions, 2008; 10:1243–1255.
23. Ivors K, Garbelotto M, Vries ID, Ruyter-Spira C, Te Hekkert B, Rosenzweig N, Bonants P. Microsatellite markers identify

three lineages of *Phytophthora ramorum* in US nurseries, yet single lineages in US forest and European nursery populations. Molecular Ecology, 2006; 15:1493–1505.
24. Mascheretti S, Croucher PJP, Vettraino A, Prospero S, Garbelotto M. Reconstruction of the sudden oak death epidemic in California through microsatellite analysis of the pathogen *Phytophthora ramorum*. Molecular Ecology, 2008; 17:2755–2768.
25. Kovacs K, Vaclavik T, Haight RG, Pang A, Cunniffe NF, Gilligan CA, Meentemeyer RK. Predicting the economic costs and property value losses attributed to sudden oak death damage in California (2010–2020). Journal of Environment Management, 2011; 92:1292–1302.
26. Suckling DM, Brockerhoff EG. 2010. Invasion biology, ecology, and management of the light brown apple moth (Tortricidae). Annual Review of Entomology, 2011; 55:285–306.
27. USDA-APHIS-PPQ. [Fowler G, Garrett L, Neeley A, Magarey R, Borchert D, Spears B.] Economic Analysis: Risk to US Apple, Grape, Orange and Pear Production from the Light Brown Apple Moth, *Epiphyas postvittana* (Walker). Raleigh, NC: US Department of Agriculture Animal and Plant Health Inspections Service, 2009. Available at: www.aphis.usda.gov/plant_health/plant_pest_info/lba_moth/downloads/lbameconomicanalysis.pdf, Accessed on March 2, 2016.
28. Bartell SM, Nair SK. Establishment risks for invasive species. Risk Analysis, 2004; 24(4):833–845.
29. Brockerhoff EG, Kimberley M, Liebhold AM, Haack RA, Cavey JF. Predicting how altering propagule pressure changes establishment rates of biological invaders across species pools. Ecology, 2014; 95(3):594–601.
30. Leung B, Drake JM, Lodge DM. Predicting invasions: Propagule pressure and the gravity of Allee effects. Ecology, 2004; 85(6):1651–1660.
31. Lockwood JL, Cassey P, Blackburn T. The role of propagule pressure in explaining species invasions. Trends in Ecology and Evolution, 2005; 20(5):223–228.
32. Simberloff D. The role of propagule pressure in biological invasions. Annual Review of Ecology, Evolution and Systematics, 2009; 40:81–102.
33. Byrd RH, Nocedal J, Waltz RA. KNITRO: An integrated package for nonlinear optimization. Pp. 35–59 in Pillo G, Roma M (eds). Large-Scale Nonlinear Optimization. Boston, MA: Springer, 2006.
34. Eschen R, Britton K, Brockerhoff E, Burgess T, Dalley V, Epanchin-Niell RS, Gupta K, Hardy G, Huang Y, Kenis M, Kimani E. International variation in phytosanitary legislation and regulations governing importation of plants for planting. Environmental Science & Policy, 2015; 51:228–237.
35. Ameden HA, Boxall PC, Cash SB, Vickers DA. An agent-based model of border enforcement for invasive species management. Canadian Journal of Agricultural Economics, 2009; 57(4):481–496.
36. Ameden HA, Cash SB, Zilberman D. Border enforcement and firm response in the management of invasive species. Journal of Agricultural and Applied Economics, 2007; 39:35–46.
37. Springborn M, Lindsay AR, Epanchin-Niell RS. Harnessing enforcement leverage at the border to minimize biological risk from international live species trade. Journal of Economic Behavior & Organization, 2016; 132(B):98–112.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website:

APPENDICES

Appendix A. Derivation of the probability of accepting a lot and expected slippage.

Appendix B. Comparative statics.

Appendix C. Estimates of the infestation rates of 91 plant genera.

Appendix D. Evaluating the maximum likelihood estimator with various data structures.

Appendix E. The subsample used in expected slippage minimization.

Appendix F. Robustness checks.

Appendix G. Sampling strategies under expectation-based expected slippage minimization.

Table C.I. Maximum likelihood estimates of the infestation rates of 91 plant genera imported from Costa Rica via freight to the Miami PIS in September 2009 with estimable infestation rates, based on observations from FY 2010–2012, assuming a 2% inspection intensity and a detection rate of 80%.

Table D.I. Monte Carlo results for the maximum likelihood estimator with various data structures.

Table E.I. The subsample used in expected slippage minimization.

Figure F.1. Distribution of the reduction in expected slippage by the optimal sampling strategy relative to RBS, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010.

Figure F.2. Distribution of the reduction in expected slippage by the optimal sampling strategy relative to proportional (2%) sampling, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010.

Figure F.3. Distribution of the proportion of lots sampled by the optimal sampling strategy under a capacity constraint consistent with RBS, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010.

Figure F.4. Distribution of the proportion of lots sampled by the optimal sampling strategy under a capacity constraint consistent with 2% sampling, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010.

Table F.I. Comparison of expected slippage for sampling strategies that minimize expected slippage and strategies based on risk-based sampling and proportional (2%) sampling, based on a daily representative subsample of the subset of the shipments from Costa

Rica inspected at the Miami PIS in September 2010 whose size ranges from 92 and 4,479.

Figure F.5. Distribution of the reduction in expected slippage by the optimal sampling strategy relative to RBS, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010 whose sizes range from 92 and 4,479.

Figure F.6. Distribution of the reduction in expected slippage by the optimal sampling strategy relative to proportional (2%) sampling, for a range of detection rates *e*, based on 100 daily representative subsamples of shipments from Costa Rica inspected at the Miami PIS in September 2010 whose sizes range from 92 and 4,479.

Table G.I. Comparison of optimal sampling strategies and expectation-based sampling strategies, under capacity constraints consistent with RBS (2,193 plant units) and proportional (2%) sampling (15,143 plant units) and with a detection rate of 0.8.