

Harnessing enforcement leverage at the border to minimize biological risk from international live species trade¹

Michael R. Springborn², Amanda R. Lindsay³, Rebecca S. Epanchin-Niell⁴

11/6/2015

DRAFT

Abstract

Allocating inspection resources over a diverse set of imports to prevent entry of plant pests and pathogens presents a substantial policy design challenge. We model inspections of live plant imports and producer responses to inspections using a “state-dependent” monitoring and enforcement model. We capture exporter abatement response to a set of feasible inspection policies from the regulator. Conditional on this behavioral response, we solve the regulator’s problem of selecting the parameters for the state-dependent monitoring regime to minimize entry of infested shipment. We account for exporter heterogeneity, fixed penalties for noncompliance, imperfect abatement control and imperfect inspections at the border. Overall, we estimate that state-dependent targeting (based on historical interceptions) cuts the rate of infested shipments that are accepted by one-fifth, relative to the status quo of uniformly allocated inspections.

Key words: state-dependent enforcement, invasive species, international trade, optimal inspections, externalities

JEL codes: Q58, Q57, Q56

¹ This work was funded by USDA-APHIS cooperative agreement 14-8100-1724-CA and by NSF grant 1414374 as part of the joint NSF-NIH-USDA Ecology and Evolution of Infectious Diseases program, and by UK Biotechnology and Biological Sciences Research Council grant BB/M008894/1.

² Department of Environmental Science & Policy, University of California Davis

³ Department of Agricultural & Resource Economics, University of California Davis

⁴ Resources for the Future, Washington, D.C.

1. Introduction

The strongest pathway for the unintentional introduction of non-native insect pests and pathogens to the United States is the importation of live plants (Liebhold et al., 2012). This leading vector is also expanding at a substantial rate: over the past four decades, plant imports to the US have grown at an average of 33% per decade (Liebhold et al., 2012). In order to protect agriculture and natural resources, the US Department of Agriculture's (USDA) Animal and Plant Health Inspection Service (APHIS) is tasked with minimizing the "entry, establishment, and spread of exotic plant pests, diseases, pathogens, and noxious weeds" (USDA-APHIS, 2012). As part of this mission, APHIS inspects shipments of imported plant material at ports of entry across the country, and shipments that are found to be infested are treated, rejected, or destroyed to prevent pest and pathogen entry. Resources for these inspections have not grown at the same rate as imports, prompting APHIS to reexamine the efficiency of the allocation of inspection effort across various sources of imports, differentiated by country-commodity pairs. Historically inspections have been allocated essentially uniformly despite heterogeneous risks. In 2011 APHIS proposed to implement a risk-based inspection (RBI) process to target high risk sources with greater inspection effort (USDA-APHIS, 2011).

While conceptually it is clear that adoption of a risk-based method for targeting inspections should provide gains when risks across shipments are non-uniform, several important policy design and evaluation questions present a challenge. If shipments are to be categorized into groups based on historical risk, how should thresholds determining group membership be set and how should inspection effort differ between groups? How should the policy be set to ensure that available inspection resources are not overburdened? What monitoring structure generates the greatest incentive for offshore producers to engage in phytosanitary abatement efforts? Finally, given that changing inspection policy is costly, what level of improvement in reducing pest and pathogen entry should be expected from a shift to RBI? In this manuscript we develop an integrated model of producer (exporter) behavior and targeted border inspections calibrated to the data on live plant imports into the US in order to provide insight into the questions above.

We model inspections of live plant imports and producer responses to inspections using a "state-dependent" monitoring and enforcement model that was first applied to pollution control problems by Harrington (1988). In Harrington's original model homogeneous firms make the decision to "comply" or "violate" an emissions standard and the regulator sets a policy in order to achieve a predetermined

target compliance rate with the fewest number of inspections. Firms are divided into two groups, each assigned with an inspection frequency and fine for noncompliance, according to an assumed set of transition rules and the outcome of the most recent inspection. Generally in state-dependent regulatory schemes, entities with worse compliance records are subject to some combination of more intense inspection, greater penalties for violations, or tougher standards (Cohen, 1998).

Most state-dependent enforcement models use the group transition policy detailed by Harrington (1988) (e.g. Harford and Harrington, 1991; Harford, 1991 ; Raymond, 1999). A firm stays in or is moved into the low-compliance group when an inspection reveals noncompliance. If an inspection reveals compliance, the firm stays in or moves into the high-compliance group with some assumed probability. In contrast to Harrington's (1988) assumed transition rule, Friesen (2003) mathematically derives the optimal transition policy for moving firms between a targeted and a non-targeted group. The "optimal targeting" policy moves firms into the targeted group at random and out of the targeted group with positive probability if an inspection reveals firm compliance. Optimal targeting out-performs the policy from Harrington (1988) unless the desired compliance rate is sufficiently high.

A direct benefit of a state-dependent or targeted inspection policy is that inspection-driven incentives for cleaner activity are focused on the dirtiest entities. An additional indirect abatement incentive is generated from the threat of moving into the low-compliance group or the prospect of escaping into the high-compliance group; this is known in the literature as "enforcement leverage" (Harrington 1988). Harrington shows abatement is greatest when the fine for the group with high compliance rates is set to the minimum and the fine for the group with low compliance rates is set to the maximum. If the expected cost of a violation is too low, however, firms in the high-compliance group may have a dampened abatement incentive. Raymond (1999) explores how heterogeneous compliance costs may impact the effectiveness of the state-dependent monitoring and enforcement policy, noting that if an industry has a large proportion of firms with low compliance costs, it is optimal to set both groups' fines to the maximum.

Harrington's 1988 model yields a counterintuitive result: it is optimal to assign different inspection intensities to firms with identical marginal abatement costs. Harford and Harrington (1991) show that a static policy, a policy where all firms face the same inspection probability, is superior to state-dependent policy only when the standard and monitoring policy can be set simultaneously. Allowing for inspection measurement error, Harford (1991) finds evidence that state-dependent enforcement is still preferred

over static enforcement once a standard has been selected. A socially optimal state-dependent mechanism treats identical firms differently because the firms' increased abatement costs are less than the reduction in the regulator's monitoring costs (Harford, 1991).

None of the aforementioned models include a regulatory budget constraint despite the fact that a targeted inspection policy is needed precisely because regulatory enforcement budgets are limited. To our knowledge, Liu and Neilson (2013) were the first to incorporate a fixed inspection capacity constraint in their state-dependent enforcement and monitoring model. They define the feasible policy set using inspection probabilities and the proportion of otherwise identical firms assigned to each of the two groups. In order ensure the chosen policy does not over-burden or under-utilize inspection capacity, Liu and Neilson develop new group transition rules. Firms with the highest level of noncompliance are moved into or stay in the targeted group according to a rank order tournament between inspected firms. Their model allows the regulator to minimize the aggregate level of pollution, more accurately capturing the objective of the regulatory agency. However, the rank order tournament requires each firm to consider the best response functions of all other firms, a particularly unrealistic assumption in our application, which considers a large number of globally distributed, heterogeneous exporters.

Previously published work on border inspections for invasive species management has considered policies that are (1) sensitive to inspection history, (2) involve high and low risk groups, (3) account for firm behavioral response, (4) allow for firm heterogeneity, and (5) set policy parameters through optimization. However no previous paper unifies all of these components. Robinson et al. (2011) consider a state-based inspection model of trade pathways to prevent invasive species introduction in which sources of trade are allocated to high and low risk groups that differ in their inspection frequency. In contrast to our application, the risk threshold to determine group membership and inspection frequencies are selected without guidance from an outcome based objective – inspections of the high risk group are set arbitrarily to 100% and the low risk group is inspected at a frequency to verify continued low risk group membership. The approach does not account for exporter behavioral response resulting from enforcement leverage or variation in inspection frequency. They nonetheless find gains relative to uniform sampling.

Ameden et al. (2007) develop a theoretical model of firm response to border enforcement in order to understand the relationship between due-care technology and probability of detection. In their four-stage model, heterogeneous exporters first choose shipment quantities and treatment effort, then regulators conduct inspections and choose enforcement actions. While Ameden et al. (2007) restrict

their attention to uniform inspections, similar to our approach they solve their model using nested optimization and backward induction. Building on this theoretical model of firm behavior, Ameden et al. (2009) develop an agent-based model (ABM) of border enforcement incorporating a spatially explicit damage function as well as an array of exogenous and endogenous port-specific variables. Regulators can choose a port-specific level of inspection and can increase inspection effort for firms with historically high inspection violations. However, this prior violation record is set to zero at the beginning of each year.

We extend and apply a state-dependent monitoring and enforcement modeling approach to APHIS' trade inspections RBI problem in which APHIS seeks to minimize the number of infested shipments that enter the US, given a capacity constraint that limits the number of shipments that can be inspected. The first model component captures exporter abatement response to a set of feasible inspection policies. Conditional on this behavioral response, the second model component captures the regulator's problem of selecting the parameters for the state-dependent monitoring regime. In particular, the model is used to determine (1) the compliance threshold for distinguishing groups and (2) group inspection frequencies. Given inspection capacity constraints, our structure captures a key tradeoff: increasing inspection frequency in the high risk group (similar to Harrington's low-compliance group) must be accompanied by either decreasing inspection frequency in the low risk group (high-compliance group) or reducing the number of entities in the high risk group. Our model accounts for exporter heterogeneity, fixed penalties for noncompliance (interceptions), imperfect abatement (infestation control) by exporters, and imperfect inspections at the border. We incorporate a fixed inspection budget by defining group assignment based on historic exporter compliance, allowing us to relax the assumptions proposed by Liu and Neilson. We characterize policy design tradeoffs, estimate numerical inputs to the model, and assess the expected level of reduced pest and pathogen entry that such a risk-based approach might achieve.

We estimate that a state-dependent monitoring and enforcement policy can reduce the expected number of infested shipments that enter into the US by one-fifth. The optimal policy entails inspecting all shipments from exporters with poor historic compliance records and less than 30% of shipments from exporters with good historic compliance records. We develop and show results for a model with homogenous exporters—to illustrate outcomes in a simple framework—and with heterogeneous exporters to increase realism. We find that the heterogeneous exporter extension is essential for arriving at an optimal compliance threshold governing group designation that is plausible.

2. Model

We begin with an overview of the model components, and then characterize the exporters' optimal response to inspection regimes which subsequently informs the regulator's optimal policy for minimizing accepted infested shipments.

2.1 Overview of model components

Each of the n exporters (offshore producers) represents a unique country-commodity pairing, where commodities are defined by plant genus. Differential treatment of exporters is governed by their historical interception rate, α , as determined by the results from inspection of their shipments. A shipment is considered intercepted if it is identified as infested during an inspection.⁵ The regulator's policy design challenge is determining how to allocate exporters into groups, indexed by g , that differ in inspection frequency. For simplicity we consider two groups, a medium frequency inspection group ($g = M$) and high frequency inspection group ($g = H$).⁶

We structure the problem using a two-stage game theoretic framework. In the first stage, the regulator announces an interception rate cutoff, $z \in [0, 1]$, which determines whether an exporter will fall into the medium inspection rate group ($g = M$ if $\alpha < z$) or the high group ($g = H$ if $\alpha > z$). The regulator also announces each group's inspection rate, ρ_g , defined as the expected proportion of an exporter's shipments targeted for inspection. Group H is undesirable for exporters given its higher expected inspection rate: $\rho_H > \rho_M$. The regulator's inspection policy vector is defined as $\Omega = (\rho_M, \rho_H, z)$ and is announced publicly to exporters. Given the regulator's expectations over the exporters' cost-minimizing response to the policy vector, the regulator selects a policy to minimize expected accepted infested shipments conditional on the available inspections budget, B . The constrained minimization problem driving this cost-effective policy design challenge is specified in detail after we first describe the exporters' problem.

2.2 Exporter policy response

⁵ Rather than "infested", APHIS uses the broader terminology of "actionable".

⁶ The existing approach of APHIS already includes the designation of a low risk group. Country-commodity pairs that are determined to 'pose an extremely low risk' are placed into a Propagative Monitoring and Release Program (PMRP). APHIS periodically inspects exporters in this group, but does not inspect all shipments. This monitoring is used to determine if PMRP status can be maintained. Because of this special inspection procedure, we set aside PMRP shipments from our data and modeling approach.

Exporters choose an infestation abatement effort level on the unit interval, $e \in [0,1]$, after the inspection policy Ω is announced. The exporter's shipment rate is given by s (shipments per period), and exporters face convex costs of abatement effort given by $c(e) = w/(1-e)^2$ for each shipment. Their shipment infestation rate (expected proportion of shipments infested) is decreasing linearly in e according to $\theta(e) = \theta_0(1-e)$, where θ_0 is the base rate of infestation when $e = 0$.

To evaluate the benefit of additional abatement effort, exporters form expectations over shipment interceptions and consider the implications of crossing the interception rate cutoff, z , which determines group membership and associated inspection frequency, ρ_g . Expectations on group membership depend on the dynamics of the observed interception rate, which are a function of the historical interception rate (a), the number of shipments inspected (I), and the number found to be infested (k):

$$a'(a, I, k) = a \left(\frac{\varepsilon}{\varepsilon + I} \right) + \frac{k}{I} \left(\frac{I}{\varepsilon + I} \right), \quad (1)$$

where $\varepsilon \geq 0$ is a parameter determining the relative weight given to the interception rate observed historically (a) and currently (k/I). This functional form is motivated by a streamlined Bayesian approach detailed in Appendix A. Its properties are such that as the value of ε increases, the "memory" of historical interception rate is stronger, and as ε becomes very large ($\varepsilon/(\varepsilon + I)$ approaches one), current outcomes have little effect on the measure. In addition, Equation (1) assigns more weight to the most recently observed interception rate (k/I) when it is based on a larger number of observations (I). When there are no new inspections the interception rate is unchanged.

For a given exporter we assume that the number of inspections, I , is a binomial random variable, conditional on the shipment rate and the inspection frequency, s and ρ_g . The number of infested shipments identified, k , is a binomial random variable, conditional on I and the underlying infestation rate, θ . We also account for the fact that inspections are imperfect. The detection rate, d , is the proportion of inspections of infested shipments that result in successful interceptions. Inspections can be imperfect due either to drawing a noninfested sample from an infested shipment or failing to detect infestation of the sample when inspecting it.⁷ We incorporate imperfect inspections as a deflator on the

⁷ Thus, d is the product of two terms, both on the unit interval: (1) the sample infestation confidence, which is the probability that a particular sample inspected from an infested shipment is itself infested, and (2) the efficiency rate, which is the probability that an infestation in an examined sample is detected.

infestation rate in the process governing interceptions: $d \cdot \theta$. Thus the likelihood of an observation (I, k) is the product of two binomial densities:

$$\begin{aligned} \Pr(I, k) &= \Pr(k | I, d \cdot \theta(e)) * \Pr(I | s, \rho_g) \\ &= \left[\binom{I}{k} (d \cdot \theta(e))^k (1 - d \cdot \theta(e))^{(I-k)} \right] * \left[\binom{s}{I} \rho_g^I (1 - \rho_g)^{(s-I)} \right]. \end{aligned} \quad (2)$$

The density function for state transitions, $\pi(\hat{a}' | a)$, i.e. the probability of transitioning to

$\hat{a}' \equiv a'(a, I, k)$ from a , is given by $\pi(\hat{a}' | a) = \sum_{I \leq s, k \leq I} \mathbf{1}_{\hat{a}'}(a') \cdot \Pr(I, k)$. Here $\mathbf{1}_{\hat{a}'}(a')$ is an indicator function equal to one whenever $a' = \hat{a}'$.

Transitions are bounded from below by $a = 0$ and from above by $a = a_{ban}$, the interception rate at which an exporter is simply banned from shipping. The probability of being in group M next period—i.e. the probability that a' is less than z —is given by the cumulative density function of a' evaluated at the policy interception rate cutoff z : $\Pi(a' = z | a)$. For simplicity of notation, let $\Pi(a' = z | a) = \Pi_{z|a}$.

Membership in group H has probability $\Pi_{a_{ban}|a} - \Pi_{z|a}$. Finally, the likelihood of being banned in the next period is $1 - \Pi_{a_{ban}|a}$. We assume that if an exporter is currently banned, there is a fixed probability of recovery in the next period, at which point their historical interception rate is reset to just below the banishment threshold at $0.95a_{ban}$.

In each period, the exporter chooses an abatement effort level e in order to minimize the expected value of long-run discounted losses. Losses in the current period equal the sum of abatement and inspection costs across the exporter's shipments:

$$L(e | \rho_{g|a}) = \begin{cases} s[c(e) + \rho_g(\delta + \gamma \cdot d \cdot \theta(e))], & \text{if } a < a_{ban} \\ sL_{ban}, & \text{if } a \geq a_{ban}. \end{cases}, \quad (3)$$

where δ is the marginal cost of an inspection (e.g. cost of delay), γ is the cost per intercepted shipment (e.g. treatment, rejection or destruction), and L_{ban} is the per-shipment loss when an exporter is banned.

The subscript on $\rho_{g|a}$ indicates that inspection frequency depends on whether the current historical

interception rate, a , places the exporter in group $g = M$ or H . The Bellman equation for the exporter's loss minimization problem given the policy vector $\Omega = (\rho_M, \rho_H, z)$ is:

$$V_{g|a}(a|\Omega) = \min_e \left\{ L(e|\rho_{g|a}) + \beta \left[\Pi_{z|a} EV_M(a') + (\Pi_{a_{ban}|a} - \Pi_{z|a}) EV_H(a') + (1 - \Pi_{a_{ban}|a}) EV_{ban}(a') \right] \right\}, \quad (4)$$

where β is a discount factor and E is the expectations operator. Using value function iteration (Judd, 1998) we solve the exporter problem above for the optimal response to any given inspection policy announcement, $e^*(a|\Omega)$.

Exporter heterogeneity is represented by dividing our large pool of exporters into $J = 4$ exporter types that differ based on combinations of shipment frequencies, s_j , and abatement costs, as characterized by the cost parameter w_j . These type-specific parameters are detailed further in our numerical application description below. Exporters of the same type may also differ in their historical interception rate, a , since the inspection and interception processes are stochastic.

2.3 Regulator policy selection

The regulator chooses the policy vector $\Omega = (\rho_M, \rho_H, z)$ to minimize the expected number of accepted infested shipments, $E(X|\Omega)$, conditional on an inspection budget, B . The budget reflects the number of shipments that can be inspected in each period (e.g. month). The number of shipments actually inspected in group g per exporter of type j is given by $I_{g,j}$. A priori, in expectation, the inspection budget constraint takes the form

$$B = S \sum_g \lambda_g \left[\sum_j \phi_{g,j} \left(\frac{E(I_{g,j})}{s_j} \right) \right] = S \sum_g \lambda_g \rho_g, \quad (5)$$

where S is the total number of shipments across all exporters, λ_g is the share of shipments in group g , and $\varphi_{g,j}$ is the proportion of shipments in group g from exporter type j . Because we are considering just two groups, let $\lambda_H = \lambda$ and $\lambda_M = 1 - \lambda$.⁸

A challenge is created by the fact that exporters respond to the vector $\Omega = (\rho_M, \rho_H, z)$ while the feasible set of policy vectors determined from the budget constraint features λ rather than z : $\Omega_\lambda = (\rho_M, \rho_H, \lambda)$. It would be essentially meaningless to announce Ω_λ to exporters since λ —the share of exporters in the high group—conveys only a very noisy idea of the threshold level of an exporter’s historical interception rate that will determine group membership. Conversely, a priori, it is unknown whether any given vector Ω will be both feasible and not leave inspection resources unallocated. Thus we require a bridge between the regulator’s need to plan in terms of λ and the exporters’ need to plan with respect to z . We address this problem as follows. For each vector $\Omega_\lambda = (\rho_M, \rho_H, \lambda)$ that is just binding with respect to the budget constraint, we solve the exporter’s problem across a range of z and then identify the particular level of this cutoff, z_λ , that induces λ , thus allowing the regulator to announce $\Omega = (\rho_M, \rho_H, z_\lambda)$.

Ensuring that the share of group H shipments indeed matches λ requires calculation of the stationary distribution of incoming shipments by interception rate conditional on the exporters’ optimal response, $e^*(a|\Omega)$. This stationary distribution is represented by the probability mass function $f_s(a)$. It is constructed in two steps: first we identify the distribution of interception rates where *exporters* are the unit of observation, $f_e(a)$, and then transform to express the distribution by *shipment*, $f_s(a)$. We start by identifying $\pi(a'|a)$ conditional on the exporter response, e^* . Here $\pi(a'|a)$ takes the form of a Markov transition matrix, since inspection observations result in discrete shifts in a . The unit of observation underlying π is the exporter. The stationary distribution of a Markov transition matrix—loosely speaking, the long-run probability of any given state—is given by the left eigenvector with corresponding eigenvalue of one, by the fundamental theorem of Markov chains (Diaconis, 2009). The appropriate eigenvector provides $f_s(a|j)$ for each exporter type j . Accounting for differences in shipping levels between types and the shares of each type, we identify the aggregate distribution, $f_s(a)$. The cumulative mass function $F_s(a)$ can then be identified. Finally we identify the particular cutoff z_λ such that the share of shipments in H matches the intended share, $1 - F_s(a|z_\lambda) = \lambda$.

⁸ While technically it is feasible for there to be a third share of exports that fall into the banned region, we use the simple notation for shares specified here since in equilibrium the threat of banning at a_{ban} provides sufficient incentive for there to be no exports above this threshold.

For a shipment—conditional on the regulator’s policy, the exporter’s type and given the historical interception rate—the expected accepted infested shipment rate is:

$$x(a | j, \Omega) = \left[\underbrace{(1 - \rho_{g|a})}_{\text{Prop. of shipments not inspected}} + \underbrace{\rho_{g|a}(1-d)}_{\text{Prop. of infested shipments inspected but not detected}} \right] \cdot \theta(e^*(a | j, \Omega)). \quad (6)$$

The regulator calculates the expected number of accepted infested shipments according to

$$E(X | \Omega) = \sum_j \sum_a x(a | j, \Omega) \cdot f_s(a | j) . \quad (7)$$

The regulator’s problem can be summarized as

$$\begin{aligned} & \min_{\Omega} E(X | \Omega) \\ & \text{s.t. } B = S \sum_g \lambda_g \rho_g . \end{aligned} \quad (8)$$

If information on how the present value of expected damages of likely invaders from different exporters were broadly available, it would be ideal to account for that heterogeneity. However, since such information is not available, the objective above implicitly treats each infested shipment equally.

2.4 Numerical parameterization

In Table 1 we list and describe parameters and variables of the model. For each parameter we also specify the value chosen and the source of the value. We use APHIS data on imports of live plants and shipment dispositions (i.e., inspection outcomes) from fiscal year 2012 to estimate and calibrate parameters. We focus on 2012 rather than 2013 or 2014 because over those latter two years APHIS’ inspection strategy has been in transition, namely in the protocol determining the size and selection of samples from a shipment selected for inspection. Observations come from APHIS’ Agriculture Quarantine Activity System (AQAS) PPQ 280/264 database. Most of the management parameters were selected based on summary statistics from these data. Several parameters were also selected based on communications with APHIS staff.

Accounting for exporter response is a challenge given that an exporter’s “abatement cost function” is not observable. We calibrate the model to match 2012 outcomes to construct as realistic a

representation of exporter response as possible. Specifically, the cost parameters w_j are chosen such that the interception rate predicted by our model matches the observed rate from 2012. In the calibration exercise we assume the regulator inspects all shipments with equal probability, essentially the approach in place in 2012. We also assume that the number of inspections available in this calibration year is slightly larger than the number available now because the new within-shipment sampling strategy APHIS began implementing after 2012 is somewhat more time intensive.⁹ The parameters used for calibration are detailed in Table 1.

We apply our model both with and without consideration of exporter heterogeneity in shipping rate and abatement cost. We incorporate exporter heterogeneity by dividing our pool of 1,545 exporters into four types. Each of the four exporter types is initially characterized by empirical estimates of interception rate and shipment frequency (detailed in Appendix B). We then calibrate the model (as described above) to identify a unique cost function parameter value, w_j , for each of the four exporter types. These four exporter types represent pairings of high or low abatement cost with a high or low shipping rate.

[Insert Table 1 about here]

3. Results

We first present results for the exporter's optimal abatement response as a function of the exporter's historical interception rate, a . Figure 1 illustrates the exporter's value function and optimal abatement effort function under the homogeneous exporter model for both the uniform inspection policy and a particular RBI policy, specifically the optimal RBI policy, discussed further below. The value function—the minimized present value of expected losses—is increasing in a . Under both policies, the value function is most sensitive to changes in the historical interception rate near the ban threshold ($a = 0.2$). Under the RBI policy, the value function is also sensitive to changes in the historical interception rate near the threshold separating the high and medium group (z_λ).

⁹ In 2013, APHIS began implementing risk-based sampling (RBS) to ensure consistent, random sampling of all inspected shipments. Under RBS, sample size for inspected shipments is selected to achieve a 95% confidence of detecting an infested shipment, assuming that 5% of the sampling units are infested. An on-line tool indicates to inspectors specific, randomly selected units to inspect from each shipment. This randomized selection of sampled units requires that each shipment be fully unloaded to locate identified samples, which requires more handling and inspection time than the ad hoc sampling strategy used previously.

[Insert Figure 1 about here]

The policy function comparison in Figure 1 demonstrates how state-dependent monitoring and enforcement policies change exporter incentives. Recall that given a shift from the uniform to RBI policy, exporters in the high group are inspected more intensively and those in the low group less intensively. However, despite exporters to the left of the threshold being inspected less intensively under RBI, the policy function solution shows that their optimal abatement response still increases, reflecting the incentive to avoid being placed into the high inspection group. Similarly, we observe a dramatic jump in optimal effort just above the interception rate cutoff that is not observed under the uniform inspection policy. When exporters have historical interception rates that are above but close to the interception rate cutoff, increased abatement effort may help return exporters to group M . These two features of exporter response under RBI illustrate ‘enforcement leverage’. A combination of enforcement leverage and the direct effect of higher inspection frequency in group H leads to higher abatement under RBI until, as α continues to increase, the threat of exceeding the banning cutoff ($\alpha = 0.20$) dominates, leading to similar peak effort under both policies. Given that abatement effort increases due to the shift from uniform inspections to RBI for both groups (except when roughly equal at extreme levels of α), exporters’ total abatement costs also go up. However, the value function shows that on balance exporters are better off (in expectation) given a shift to RBI policy when they are under or near the threshold due to reduced costs from the inspection process.

When the model is extended to include exporter heterogeneity, each exporter type has a unique value and policy function. The qualitative policy function results from the homogeneous model hold for the heterogeneous model, although exporter types with particularly low abatement costs have high abatement effort even before the policy change, which leaves only a small margin for improvement.

An exporter’s optimally selected abatement effort results in a particular rate of expected shipment infestation. Infestations occur in some shipments, which have a positive probability of interception if they are inspected. Given randomness in shipment infestation and inspections—as well as variation in abatement cost and shipment rates when accounting for exporter heterogeneity—a distribution of interception rates arises. In Figure 2 we plot the resulting stationary distributions of interception rates by shipment for the uniform inspection and optimal RBI policies from the homogeneous exporter model. While the uniform policy leads to a fairly smooth distribution, the optimal RBI policy generates a bimodal distribution with peaks occurring to the left and right of the interception rate cutoff that

separates the medium and high groups. This bimodality is not surprising given that abatement effort is high near the threshold z_λ , which serves to drive infestation rates down in order escape group H .

Table 2 displays the average interception rates for all cases (policies, groups and types). It also provides the percentage of shipments that fall into groups M and H in equilibrium—i.e. the shipment shares the above and below the vertical line at z_λ in Figure 2. Since under the uniform policy there is no difference in treatment among groups, group designation in this context simply reflects how exporters are distributed relative to z_λ from the RBI policy. Confirming intuition from Figure 2, we find that the share of shipments with interception rates greater than z_λ declines by about one-third under optimal RBI. The average interception rate for shipments in group M increases but this is due to the influx of exporters improving from group H to group M under RBI.

[Insert Figure 2 about here]

In the heterogeneous exporter model each exporter type has a unique stationary distribution of historical interception rates. The two low abatement cost exporter types (1 and 2) display qualitative differences similar to those observed in the homogeneous exporter model. In particular, under the optimal RBI policy the mass of shipments below the interception rate cutoff increases substantially, while above the interception rate cutoff the average interception rate decreases. In both the uniform inspection policy and optimal RBI policy, shipments from high abatement cost exporter types (3 and 4) lie entirely above the interception rate cutoff. Thus, the observed reduction in historical interception rates for high abatement cost types is due entirely to greater inspection frequency under RBI, rather than ‘enforcement leverage’.

[Insert Table 2 about here.]

Combining results for the exporters’ policy function and the resulting likelihood of different interception rate states, we calculate the expected accepted infested shipment (EAIS) rate for a broad set of feasible policies. This is the key statistic that the regulator seeks to minimize. Note that while interception rates (discussed above) are an observable statistic, in contrast the EAIS is determined by infested shipments, which are estimated rather than directly observed.

Under the homogenous model we find that the optimal RBI policy that minimizes EAIS is $\Omega_\lambda = (\rho_M, \rho_H, \lambda) = (0.28, 1.00, 0.57)$. The corresponding equilibrium interception rate cutoff to be announced by the

regulator is given by $z_\lambda = 1.2\%$. To put this threshold value in context we examined the pool of exporters in our dataset and determined that the targeted level of exports in the high group ($\lambda=57\%$) is achieved using a cutoff of $z_\lambda=0.16\%$. However, this cutoff calculation ignores the behavioral response we expect from the implementation of the RBI policy: we have shown that the RBI policy increases optimal abatement and shifts the distribution of expected historical interception rates down towards zero. Thus in our applied setting we would expect that the requisite cutoff would be substantially lower than $z_\lambda=0.16\%$. This implies that the cutoff identified under the simple homogenous model is not appropriate for our setting.

When we increase realism via the heterogeneous exporter model, we find that the interception rate cutoff is in the expected range, $z_\lambda=0.00016\%$. While the equilibrium cutoff is quite different, we find that the optimal policy vector is the same as in the homogeneous model, $\Omega_\lambda = (\rho_M, \rho_H, \lambda) = (0.28, 1.00, 0.57)$. We present the full set of results for the EAIS in Figure 3 for the heterogeneous exporter model (for the analogous figure for homogeneous exporters see Appendix C). The set of feasible policy vectors, $\Omega_\lambda = (\rho_M, \rho_H, \lambda)$, underlying the results in Figure 3 is depicted in the left panel of Figure 4. In the right panel of Figure 4 we present the equilibrium threshold, z_λ , corresponding to each of the policies considered.

The left-most column of cells in Figures 3 and 4 reflects the outcome under uniform inspections ($\rho_M = \rho_H = 0.69$) which results in an EAIS rate of 8.28%. The optimal policy minimizes this rate at 6.66% by inspecting $\rho_H = 100\%$ of shipments in the high group, which contains a little over half of all shipments ($\lambda = 57\%$). Inspections of the remaining 43% of shipments that fall in the medium risk group are much less intense, with $\rho_M = 28.2\%$ of shipments inspected. The general pattern in Figure 3 is quite regular with outcomes generally improving as λ and ρ_H increase. However it is not optimal to set λ at its maximum feasible value (given $\rho_H = 100\%$), as this would reduce ρ_M even further. Overall, relative to uniform inspections, we find that increasing inspection frequency in the high group and decreasing inspection frequency in the medium group—both by roughly 50%—cuts the EAIS rate by one-fifth. This optimal strategy involves targeting the 57% of shipments from exporters with the highest interception rates with approximately 82% of the inspection budget.¹⁰

¹⁰ The share of inspections allocated to group H is given by $\lambda\rho_H / (\lambda\rho_H + (1-\lambda)\rho_M)$.

All cases that fall above the gray line in Figure 3 are feasible but do not completely exhaust the inspection budget. For these cases it was not numerically feasible to identify a level of z_λ such that the equilibrium share of shipments in group H was equal to λ within a tolerance of $\pm 0.5\%$ as was achieved in other cases (formally, $|[1-F_s(a|z)] - \lambda| < 0.005$). For cases above the grey line, the share of exports in group H changes with discrete jumps that are too large to closely exhaust but not exceed the inspections budget. For these cases—which include the policy vector that performed best—we selected the level of z_λ that came closest to exhausting inspections without exceeding the budget.

[Insert Figure 3 about here]

[Insert Figure 4 about here]

4. Conclusion

In this manuscript we have shown how the state-dependent, enforcement leverage model can be adapted to the problem of inspecting international trade for biological risk. A key constraint in this setting is the reality of limited inspection capacity, as previously explored by Liu and Neilson (2013). This setting also generated novel challenges. Each potential target for inspection (exporter) presents the regulator with multiple opportunities to inspect each period (multiple shipments) as opposed to inspection of a single fixed facility. Simply announcing the share of targets that will lie in the most intensely inspected group—as is typical in the standard fixed-polluter setting (e.g. Liu and Neilson, 2013)—will not work in a setting like ours with such diversity and scale that inspected entities are unable to generate expectations about where they stand relative to others. We overcome this challenge with a structure in which the regulator announces a threshold for the exporter’s cumulative compliance metric (historic interception rate) that it *expects* to induce the desired share of shipments in each group.

Many papers in the state-dependent enforcement literature incorporate group-specific penalties and inspection intensities. The models that additionally assume homogeneous firms find that the penalty for the low violation group should be set to a minimum and inspection intensity in the high violation group should be close to unity (e.g. Harrington, 1988; Harrington and Harford, 1991; Friesen, 2003). Similarly, we find that under a homogeneous exporter model, it is optimal to set inspection effort to its maximum

for the high interception group—a result that continues to hold under heterogeneous exporters. The existing literature also suggests that a non-zero penalty for the low violation group is preferred once measurement error is introduced (Harford 1991) or there is firm heterogeneity with a large proportion of low compliance cost firms (Raymond 1999)—both of which also apply in our setting. Furthermore, our costs of violation (interception) are fixed, and differences in group-specific expected cost of noncompliance is generated completely by inspection intensities, as in Liu and Neilson.

The optimal strategy in our setting involves inspecting with certainty the 57% of shipments from exporters with the highest interception rates, using approximately 82% of the inspection budget. We estimate that relative to uniform inspections—which target 69% of shipments from each exporter—the optimal RBI policy cuts the EAIS rate by one-fifth by increasing inspection frequency in the high group and decreasing inspection frequency in the medium group—both by roughly 50%. This improvement is substantial in light of the fact that this EAIS rate under a uniform inspection policy (8.3%) is already quite low in absolute terms—Harrington (1988) and Friesen (2003) show that gains from state-dependent enforcement and monitoring diminish as the targeted compliance rate increases—i.e. as the baseline violation rate becomes small. Friesen (2003) also notes that as a firm’s compliance cost increases relative to the maximum fine, higher inspection frequency is required but may not be sufficient, to induce compliance. Along these lines, in our results, under the optimal policy we find that high cost exporters (types 3 and 4) increase abatement effort but never sufficiently to escape the high violation group.

While our results regarding optimal inspection intensities and shares of exports allocated to the two groups were the same under the homogeneous and heterogeneous exporter model, identification of the appropriate interception rate cutoff delineating the groups changed substantially. While it could be the case that this threshold would change further as more fine scaled heterogeneity is included (i.e., further exporter types) the four types examined here utilized were only just computationally feasible, despite the use of computing resources with high RAM (80 GB). From a practical standpoint, this suggests that a regulator might be well served by approaching the ultimate threshold level from above in a series of declining steps over time so as to approach eventual full usage of inspection resources without generating an unfeasible workload.

References

- Ameden, H.A., Boxall, P.C., Cash, S.B., Vickers, D.A., 2009. An Agent-Based Model of Border Enforcement for Invasive Species Management. *Canadian Journal of Agricultural Economics* 57(4), 481-496.
- Ameden, H.A., Cash, S.B., Zilberman, D., 2007. Border Enforcement and Firm Response in the Management of Invasive Species. *Journal of Agricultural and Applied Economics* 39, 35-36.
- Cohen, M.A., 1998. Monitoring and Enforcement of Environmental Policy. *International Yearbook of Environmental and Resource Economics 2006/2007*. H. Folmer and T. H. Tietenberg. Cheltenham, Edward Elgar Publishing.
- Diaconis, P., 2009. The Markov Chain Monte Carlo Revolution. *American Mathematical Society* 46(2), 179-205.
- Friesen, L., 2003. Targeting enforcement to improve compliance with environmental regulations. *Journal of Environmental Economics and Management* 46(1), 72-85.
- Harford, J.D., 1991. Measurement error and state-dependent pollution control enforcement. *Journal of Environmental Economics and Management* 21(1), 67-81.
- Harford, J.D., Harrington, W., 1991. A reconsideration of enforcement leverage when penalties are restricted. *Journal of Public Economics* 45(3), 391-395.
- Harrington, W., 1988. Enforcement leverage when penalties are restricted. *Journal of Public Economics* 37(1), 29-53.
- Judd, K.L., 1998. *Numerical Methods in Economics*. MIT Press.
- Liebhold, A.M., Brockerhoff, E.G., Garrett, L.J., Parke J.L., Britton, K.O., 2012. Live plant imports: the major pathway for forest insect and pathogen invasions of the US. *Frontiers In Ecology and the Environment* 10(3), 135-143.
- Liu, L., Neilson, W., 2013. Enforcement Leverage with Fixed Inspection Capacity. *Strategic Behavior and the Environment* 3(4), 305-328.
- Raymond, M., 1999. Enforcement leverage when penalties are restricted: a reconsideration under asymmetric information. *Journal of Public Economics* 73(2), 289-295.
- Robinson, A., Burgman, M.A., Cannon, R., 2011. Allocating surveillance resources to reduce ecological invasions: maximizing detections and information about the threat. *Ecological Applications* 21(4), 1410-1417.
- USDA-APHIS, 2011. *Importation of Plants for Planting; Risk-Based Sampling and Inspection Approach and Propagative Monitoring and Release Program Animal and Plant Health Inspection Service, US Department of Agriculture*. Washington, DC, Government Printing Office. 76, 65165.
- USDA-APHIS, 2012. *Questions and Answers: New Risk-Based Sampling Protocol and Propagative Monitoring and Release Program at Plant Inspection Stations. Animal and Plant Health Inspection Service, US Department of Agriculture*. Washington, DC.

Appendix

A. Specification of historical interception rate dynamics.

A simple but problematic specification for historical interception rate dynamics which allows for the contribution of past observations to diminish over time is given by:

$$a'(a, k, I) = \alpha \cdot a + (1 - \alpha) \frac{k}{I}, \quad (9)$$

where $\alpha \in (0,1)$ determines the relative weight given to the interception rate observed historically (a) and currently (k/I). While appealing in its simplicity, this function has the undesirable properties of regressing towards zero automatically if there are no new observations and the most recently observed rate of interception (k/I) is given the same weight regardless of the number of underlying observations. An alternative specification can be derived from a Bayesian perspective, which overcomes these limitations. Let a_o represent the unobserved true expected rate of interception. Let beliefs about a_o be given by a Beta distribution with a mean of μ and concentration parameter of c : $a_o \sim \text{Beta}(\mu, c)$. We replace μ with a since group membership will be determined by the expected interception rate, $E(a_o) = \mu = a$. Updating beliefs using Bayes rule given the number of inspections, I , and interceptions, k , results in updated parameters of $c^* = c + I$ and $a^* = (ac + k)/(c+I)$. This form incorporates no memory loss, i.e. the influence of an observation does not diminish with time, only with the total number of observations. If we introduce memory loss in the concentration parameter before it is used in updating, then we have $c' = c\alpha + I$, with $\alpha \in (0,1)$. This has the effect of increasing the spread of beliefs since c' is smaller than it otherwise would be. Updating of the mean is also affected. Specifically, memory loss ($\alpha < 1$) leads to a shift in relative weight from historical observations to new:

$$a'(a, c, k, I) = \frac{ac\alpha + k}{c\alpha + I} = a \left(\frac{c\alpha}{c\alpha + I} \right) + \frac{k}{I} \left(\frac{I}{c\alpha + I} \right). \quad (10)$$

This equation is of similar form to the simple version in Equation (9)—a weighted combination of a and k/I . However, the Bayesian form here has two advantages: (1) the relative strength of the new information in the ratio k/I is increasing in the number of observations (I), and (2) if there are no observations then there is no change in the historical interception rate: $a' = a$. However, this dynamic equation also introduces substantial additional complexity—a second variable (c) would need to be tracked by the regulator and would also enter the exporters' decision problem. To maintain the two

advantages listed above but avoid the additional complexity we adopt a simplified version of Equation (10) in which c is fixed. Consolidating the two constant terms in to one, $c\alpha = \varepsilon$, results in Equation (1) in the main text. Here we set c equal to the average number of observations for a year (given by the number of annual shipments times the likelihood they are inspected) which approximately reflects the level of confidence we would have if we focused on observations from just the last year. Also consistent with the notion of emphasizing observations from the last year, we set α such that “old” observations are given a weight of 11/12. Thus we have $\varepsilon = c\alpha = (36*0.796)*(11/12) = 26.25$.

B. Characterizing exporter heterogeneity

We consider two aspects of exporter heterogeneity: abatement cost and shipment frequency. While abatement cost is not directly observed in our empirical data, we observe heterogeneity in historical interception rate, which we use to estimate exporter abatement costs. Thus, to represent exporter heterogeneity in our model, we categorize exporters (i.e., commodity - country combinations) into four groups based on shipment frequency and historical interception rates and then use the average of these characteristics for each group to represent four exporter “types.” We then estimate each group’s abatement costs using the calibration approach described in the main text. Our approach requires estimating interception rates and shipment frequencies for each exporter and grouping exporters based on these characteristics.

For each exporter with at least one shipment in FY 2012, we calculated the number of shipments and proportion of inspected shipments that were found infested using data on shipment dispositions (i.e., inspection outcomes) from the AQAS PPQ 280/264 database. Shipments were considered intercepted if they were found to have an “actionable pest” as indicated by a disposition of "Destroyed Actionable Pest on NARP", "Destroyed-Actionable Pest", "Fumigated Actionable Pest on NARP", "Fumigated-Actionable Pest", "Other Action Taken-Actionable Pest", or "Returned-Actionable Pest" (USDA-APHIS, personal communication). The observed interception rate across all 71,651 inspected shipments in FY 2012 was 0.0252 (i.e., actionable pests were found in 2.52% of inspected shipments) (Table A1). These shipments were sent by 3198 different exporters with an average shipment frequency of 22.4 per year in FY 2012.

Many exporters had far too few shipments in FY 2012 to rigorously estimate their interception rate. For example a sample of five shipments provides a poor estimation of an interception rate that likely lies

between 0-10%, and almost certainly below 20%. Thus, to estimate the set of exporter interception rates, some with few observations, we use a Bayesian approach in which a prior distribution describing beliefs about the true interception rate is updated with available observations. Let a_j represent the unobserved true probability that a shipment from exporter j would be found infested if inspected. This probability is not known with certainty. We characterize beliefs about the value of a_j using a Beta distribution, $a_j \sim \text{Beta}(\mu_j, c_j)$, where $\mu_j = E(a_j)$ is the expected value and c_j is the concentration. We use the observed interception rate across all shipments in FY 2012 as the initial value of μ_j for each exporter. We select the initial value of the concentration parameter to reflect diffuse beliefs with a wide spread, specifically $c_j = 2$. The prior distribution is thus given by $(\mu_0, c_0) = (0.0252, 2)$ for all exporters. We then update the prior for each exporter using their total number of inspections in FY 2012, n_j , and total number of shipments intercepted, k_j . Using Bayes rule, the posterior expected value of exporter j 's interception rate is $\mu'_j = (c_0 * \mu_0 + k_j) / (c_0 + n_j)$. With this approach we assigned each exporter (i.e., each commodity– country pair) a historic interception rate estimate, $a_j = \mu'_j$.

In our numerical analyses we focus on a subset of exporters for inclusion in the RBI policy. This excludes exporters with 2 or fewer shipments in FY 2012. These extremely low frequency exporters are excluded because APHIS intends to always inspect these exporters at the maximum intensity, as their low volume of shipments limits learning about their underlying infestation rate. This group includes approximately half of all exporters, but less than 3% of shipments (Table A1). In addition, exporters with more than 150 shipments and a historic interception rate of less than 1% are excluded since these exporters have already been identified as extremely low risk and are subject to less frequent inspections under APHIS's Propagative Materials Release Program. This includes less than 2% of exporters and 20% of all shipments (Table A1). Our analysis focuses on the remaining 55,302 shipments from 1,545 exporters. Summary statistics for all exporters are presented in Table A1.

[Insert Table A1 about here]

We divide the focal set of exporters into 4 categories defined by low and high interception rates and shipment frequencies. Average shipment frequency for a type is simply the mean number of shipments per exporter in the type. The average interception rate for a type is the shipment-weighted mean interception rate across exporters in the type (though note that this simplifies to the interception rate for the pool of all shipments within the type). An interception rate of $a=0.0046$ (or 0.46% of inspected

shipments found infested) was chosen as the cutoff for dividing exporters roughly equally between the low and high interception rate type. The cutoffs for dividing exporters into low and high shipment frequency groups were selected so that the numbers of exporters were relatively even between groups and the average *monthly* shipment frequency of the low shipment group is an integer (e.g. 1 or 2). We impose this constraint because the inspection model operates on a monthly time step and non-integer shipping frequencies adds unnecessary complexity. Monthly shipping rates for high shipment rate groups were rounded to the nearest integer. The resulting shipment frequency cutoffs were 52 and 85 for the low and high interception rate exporters, respectively. The average shipment frequency and interception rates of exporters in each group are used to characterize four representative exporters for parameterizing the model. We subsequently estimate a separate abatement cost parameter for each group, as we assume that heterogeneity in exporter interception rates and underlying infestation rates arises due to differences in abatement costs. The four groups of exporters and their interception rate and shipment frequency characteristics are shown in Table A2.

[Insert Table A2 about here]

C. Additional results.

In the figures below we present the optimal policy results for the homogenous model.

[Insert Figure A1 about here]

[Insert Figure A2 about here]

Tables

Table 1. Descriptions, values and sources for model parameters and variables.

<i>Parameter</i>	<i>Description</i>	<i>Value</i>
Regulator parameters and variables		
B'	Number of shipments inspected per period (month) under pre-RBS ^d inspection policy, for calibration: $B' = [(total\ shipments\ 2012) * \rho] / 12$	3687 ^{a,b}
B	Number of shipments inspected per period under uniform and RBI sampling inspection policies: $B = B' / \mu$	3206 ^{a,b}
μ	Proportional increase in average inspection time required per inspection under uniform and RBI relative to pre-RBS ^d	1.15 ^b
a	Exporter observed interception rate	Determined by inspections
a_{ban}	Interception rate ban level	0.20 ^b
$Pr(recovery)$	Per period probability that an exporter's shipments will be permitted again after an exporter has been banned	0.05 ^b
ρ	Proportion shipments inspected pre-RBS ^d , for calibration	0.80 ^b
ρ_g	Proportion of shipments in group g inspected per period under RBI policy	Regulator optimal choice
λ_g	Proportion of exports in group g under RBI policy $\lambda = \lambda_H = 1 - \lambda_M$	Regulator optimal choice
z_λ	Interception rate cutoff determining group membership that induces the targeted share of exporters in the high group, λ	Regulator announcement
<i>efficiency rate</i>	Percentage of examined infested samples identified as infested	0.40 ^b
<i>sample infestation confidence</i>	Conditional on shipment infestation, the probability that the sample chosen for inspection is infested	0.80 ^b (pre-RBS ^d) 0.95 ^b (RBS ^d)
d	Detection rate, i.e. percentage of infested shipments that are intercepted during inspection. $d = (efficiency\ rate) * (sample\ infestation\ confidence)$	0.32 ^b (pre-RBS ^d) 0.38 ^b (RBS ^d)
Exporter parameters and variables		
θ_0	Base rate of shipment infestation given no abatement effort	0.80 ^b
.	Average value of shipment	\$5000
δ	Marginal cost of inspection (shipment delay) to the exporter. As a proportion of the total value of the shipment	0.01 ^b
γ	Expected cost to exporter per intercepted shipment, associated with treatment, destruction or rejection of intercepted shipments. As a proportion of the total value of shipment.	0.452 ^{a,b}
e	Abatement effort. Endogenously selected by each exporter type	Exporter optimal choice
<i>Homogenous exporter model parameters</i>		
n	Total number of exporters	1,545
s	Average shipment rate per period (month)	3 ^a
w	Phytosanitary effort cost function parameter	0.89 ^c
<i>Heterogeneous exporter model parameters</i>		
n_1, n_2, n_3, n_4	Total number of exporters for each exporter type	[603,202,669,71] ^a
s_1, s_2, s_3, s_4	Average shipment rate per period (month) for each exporter type	[2,7,1,19] ^a
w_1, w_2, w_3, w_4	Phytosanitary effort cost function parameter for each exporter type	[1.7x10 ⁻⁴ , 5x10 ⁻⁶ , 8.26, 46.19] ^c

Other parameters

ε	Parameter of interception rate updating function determining persistency of past observations	11/12
β	Discount factor	1/1.03

^a Calculated from 2012 data, the benchmark year

^b Established in personal communications with APHIS

^c Calibrated to ensure uniform inspection model output matches 2012 data

^d RBS (risk-based sampling) is an amended (post 2012) APHIS protocol for rigorously sampling units within shipments.

Table 2: Average equilibrium interception rate and equilibrium percentage of shipments in groups.

Homogeneous exporter model							
Uniform inspection policy	Average interception rate	Group M	3.22%				
		Group H	3.18%				
		All	3.18%				
	% shipments in group H		87.4%				
Optimal RBI policy	Average interception rate	Group M	1.27%				
		Group H	3.69%				
		All	2.65%				
	% of shipments in group H		57.0%				
Heterogeneous exporter model			Exporter type				
			1	2	3	4	All types
Uniform inspection policy	Average interception rate	Group M	0.20%	0.066%	-	-	
		Group H	0.20%	0.066%	5.84%	6.84%	
		All	0.20%	0.066%	5.84%	6.84%	2.90%
	% of shipments in group H		38.9%	9.8%	100.0%	100.0%	60.8%
Optimal RBI policy	Average interception rate	Group M	0.081%	0.027%	-	-	
		Group H	0.22%	0.056%	7.70%	9.71%	
		All	0.11%	0.028%	7.70%	9.71%	3.97%
	% of shipments in group H		22.1%	3.7%	100.0%	100.0%	49.8%

Table A1. Summary statistics of live plant exports to the US in FY 2012

	Total Exporters	Total Shipments	Average Posterior Interception Rate	Average Shipment Frequency	Average Prior Interception Rate
All exporters	3198	71651	2.50%	22.4	2.52%
Exporters with > 150 shipments & interception rate < 1%	59	14413	0.14%	244.3	0.10%
Exporters with ≤ 2 shipments	1594	1936	3.32%	1.2	4.34%
Focal exporters (>2 shipments, excluding exporters with > 150 shipments & interception rate < 1%)	1545	55302	3.08%	35.8	3.09%

Table A2. Summary statistics for exporters grouped by type.

Intercept. Rate Type	Ship- ment Freq. Type	Avg. inter- cept. Rate	Avg. Shipment Freq. (annual)	Avg. Ship. Freq., (monthly, rounded)	Total Ex- porters	Total Ship- ments	% of Ship- ments	% of Ex- porters
Low	Low	0.19%	24.0	2	603	14469	26.2	39.0
Low	High	0.07%	82.8	7	202	16723	30.2	13.1
High	Low	6.10%	12.0	1	669	7993	14.5	43.3
High	High	7.31%	227.0	19	71	16117	29.1	4.6

Figures

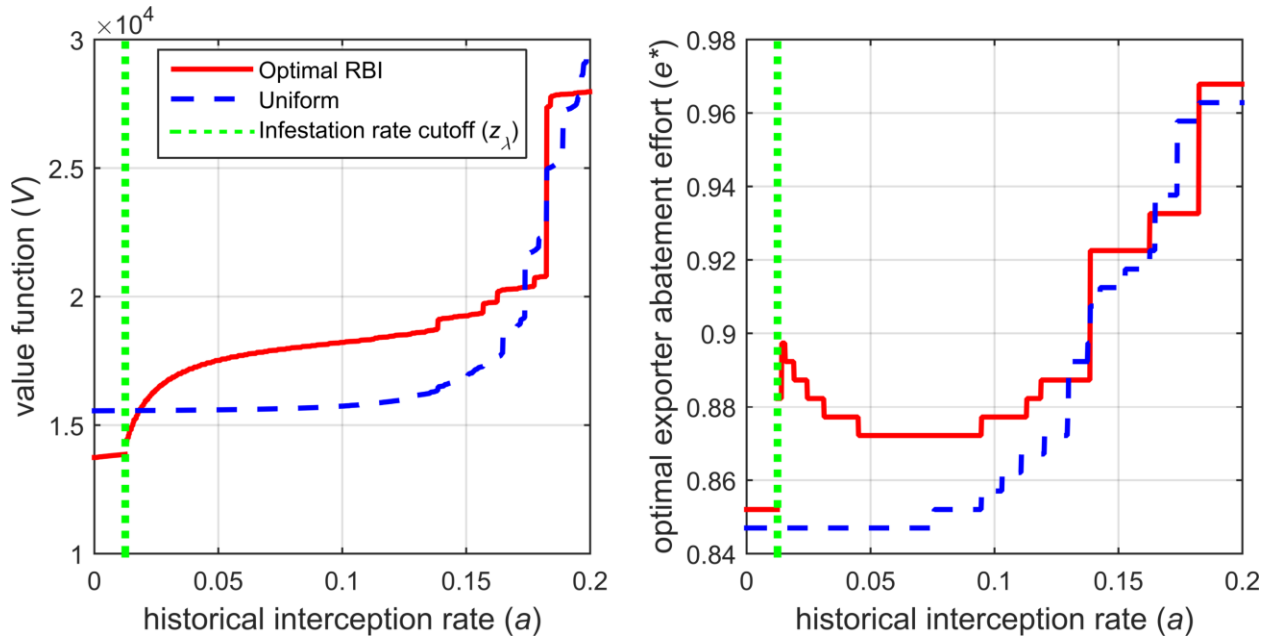


Figure 1. Value function (left) and policy function (right) solutions for homogenous exporter model under the uniform policy and optimal RBI policy, $[\rho_M, \rho_H, z_\lambda, \lambda] = [0.28, 1.00, 0.012, 0.57]$.

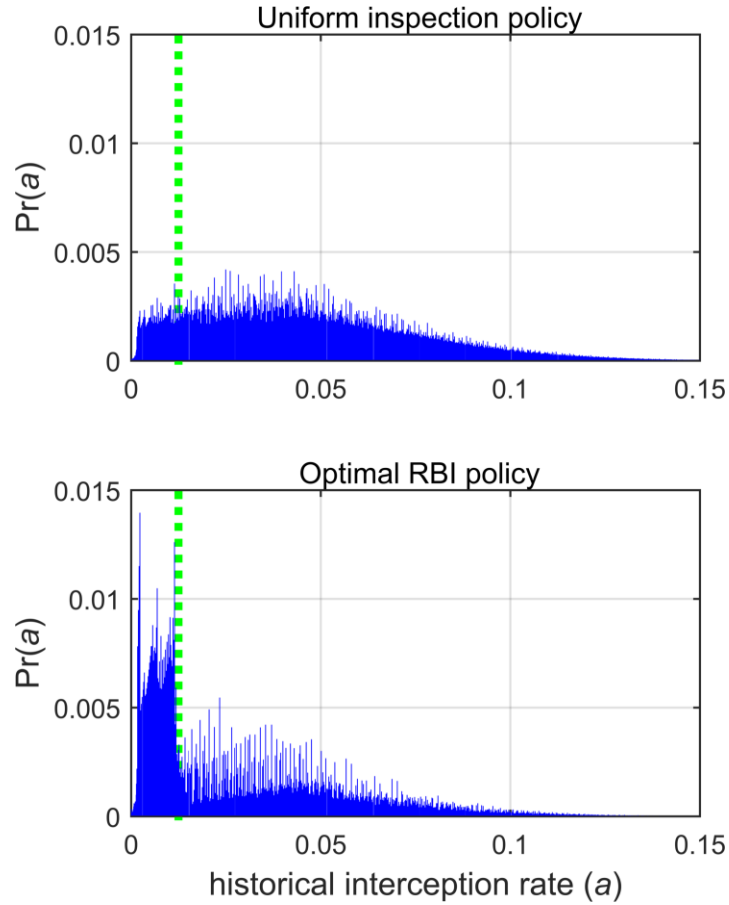


Figure 2. Stationary distribution of historical interception rates by shipment, for the homogeneous exporter model. Graphs display equilibrium probability mass of historical interception rates, $\Pr(a)$, for a uniform inspection policy (top), $[\rho_M, \rho_H] = [0.69, 0.69]$, and the optimal RBI policy (bottom), $[\rho_M, \rho_H, z_\lambda, \lambda] = [0.28, 1.00, 0.012, 0.57]$. The vertical dotted line depicts the interception rate cutoff, z_λ .

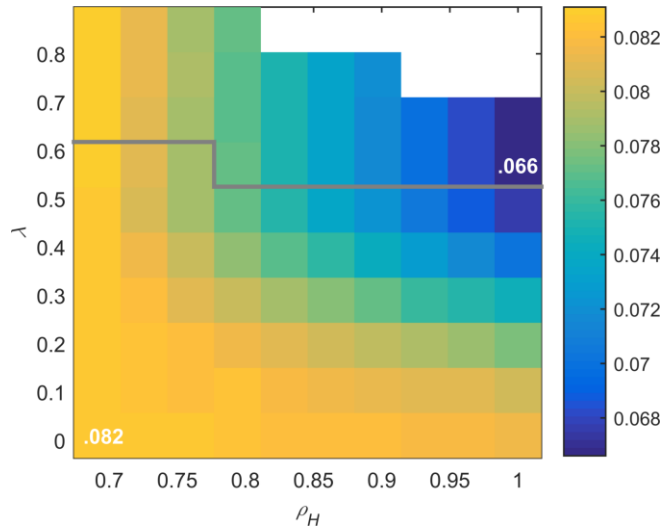


Figure 3. Expected accepted infested shipment (EAIS) rate across all shipments under the heterogeneous model. Numbers in two individual cells indicate values for uniform inspection policy (southwest corner) and the optimal policy (northeast region). Each colored cell corresponds to a feasible policy vector characterized in Figure 4. The left column of cells represents the uniform policy ($\rho_M = \rho_H$). Feasible policies (non-white cells) that fall above the gray line do not completely exhaust the inspection budget.

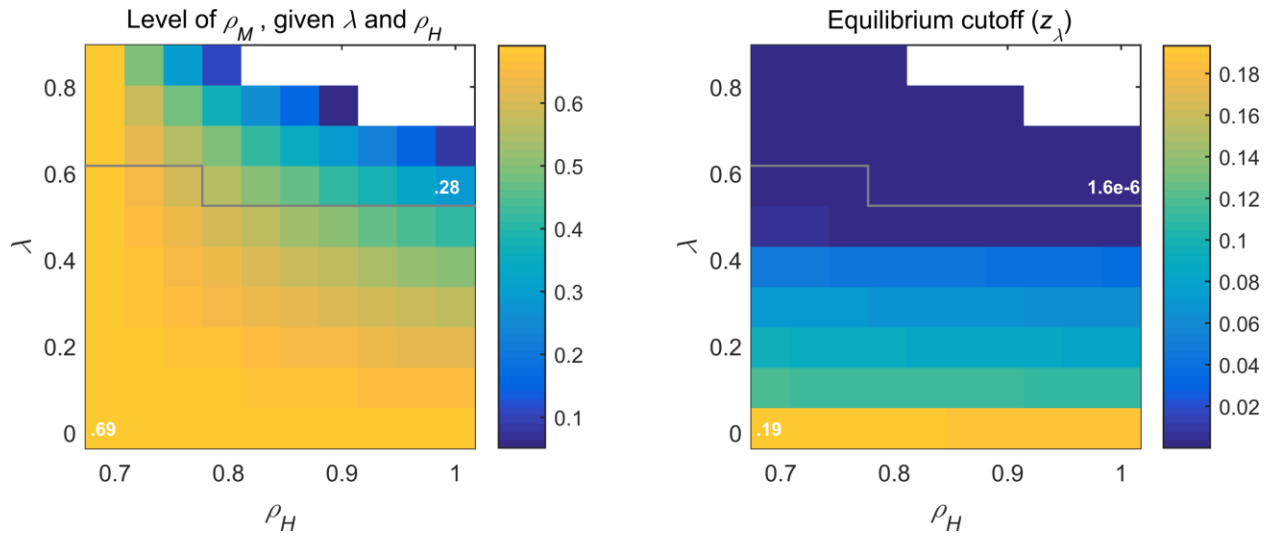


Figure 4. The left image illustrates the set of feasible policy vectors, $\Omega_\lambda = (\rho_M, \rho_H, \lambda)$, considered by the regulator. The right image illustrates each policy's corresponding interception rate cutoff z_λ . Numbers in two individual cells indicate values for uniform inspection policy (southwest corner) and the optimal policy (northeast region) where $z_\lambda = 1.6 \times 10^{-6}$. The left column of cells represents the uniform policy ($\rho_M = \rho_H$). In both images, white cells represent policies that exceed the budget constraint. Feasible policies (non-white cells) that fall above the gray line do not completely exhaust the inspection budget.

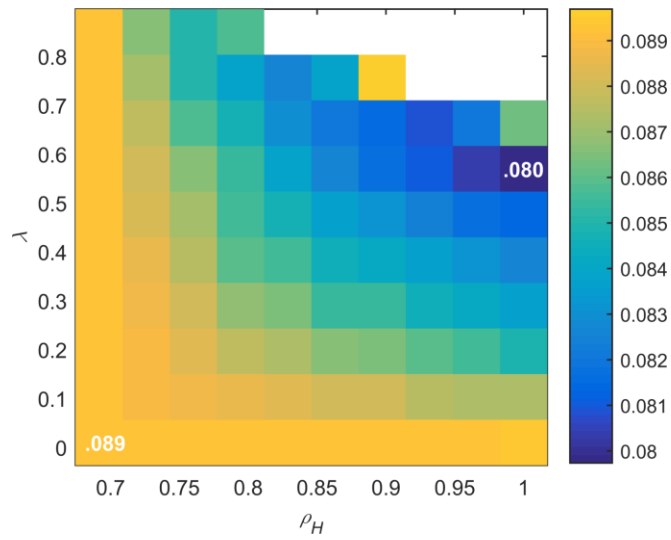


Figure A1. Expected accepted infested shipment rate (per shipment) from homogeneous model results given inspection frequency in high group (horizontal axis, ρ_H) and equilibrium targeted shipment share in high group (vertical axis, λ). Numbers in two individual cells indicate values for uniform policy (bottom-left) and optimal policy (rightmost). Each colored cell corresponds to a feasible policy vector characterized in Figure A2. The left-most column represents the uniform policy.

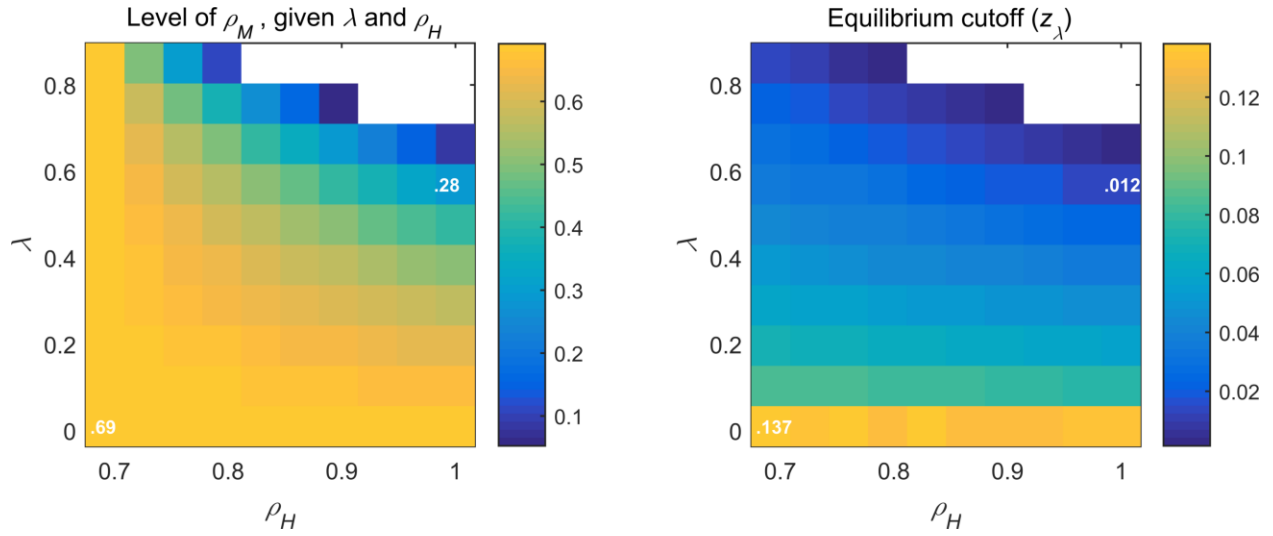


Figure A2. The left image illustrates the set of feasible policy vectors, $\Omega_\lambda = (\rho_M, \rho_H, \lambda)$, considered by the regulator. The right image illustrates each policy's corresponding interception rate cutoff z_λ . In both images, white cells represent policies that exceed the budget constraint. The left-most column represents the uniform policy.