MANAGEMENT SCIENCE
Vol. 57, No. 1, January 2011, pp. 1–20
ISSN 0025-1909 | EISSN 1526-5501 | 11 | 5701 | 0001

informs®
DOI 10.1287/mnsc.1100.1252


# Estimating the Operational Impact of Container Inspections at International Ports

Nitin Bakshi
London Business School, Regent's Park, London NW1 4SA, United Kingdom, nbakshi@london.edu

Stephen E. Flynn
Center for National Policy, Washington, DC 20001, sflynn@centerfornationalpolicy.org

Noah Gans
The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104, gans@wharton.upenn.edu

A U.S. law mandating nonintrusive imaging and radiation detection for 100% of U.S.-bound containers at international ports has provoked widespread concern that the resulting congestion would hinder trade significantly. Using detailed data on container movements, gathered from two large international terminals, we simulate the impact of the two most important inspection policies that are being considered. We find that the current inspection regime being advanced by the U.S. Department of Homeland Security can only handle a small percentage of the total load. An alternate inspection protocol that emphasizes screening—a rapid primary scan of all containers, followed by a more careful secondary scan of only a few containers that fail the primary test—holds promise as a feasible solution for meeting the 100% scanning requirement.

*Key words*: homeland security; container inspections; queueing simulation
*History*: Received February 13, 2009; accepted August 8, 2010, by Yossi Aviv, operations management. Published online in *Articles in Advance* November 19, 2010, and updated December 3, 2010.

## 1. Introduction

The consensus among security experts is that the most probable way that Americans would be targeted by a nuclear weapon would be for al-Qaeda or a future adversary to smuggle it into the United States (Flynn 2008). The millions of shipping containers that are used to transport goods in ocean-going vessels provide terrorists with one promising way to hide a nuclear device destined for U.S. shores.[1] By using a container, terrorists can potentially achieve mass disruption of global supply chains: widespread public anxiety that other containers may contain nuclear devices would result in stepped-up inspections that would cause congestion throughout the global intermodal transportation system. Abt (2003) estimates that the detonation of a nuclear device in a port may lead to losses in the range of $55–$220 billion.

### 1.1. U.S. Security Initiatives in Place at International Ports

To counter the threat of nuclear terrorism, the United States has initiated various security measures, both at domestic and foreign ports. These measures can require the cooperation of foreign nations, trading companies, terminal operators, customs brokers, trucking companies, ocean carriers, and other participants in the maritime supply chain. In this paper, we focus on security initiatives implemented at international ports, namely, the Container Security Initiative (CSI) and the Secure Freight Initiative (SFI). These constitute only 2 out of nearly 25 to 30 U.S. and international initiatives and legislations directed at enhancing maritime security. We refer the interested reader to Boske (2006), which provides a nice overview of the various initiatives in place.

CSI is a security program administered by U.S. Customs and Border Protection (CBP), an agency that falls within the Department of Homeland Security. The program, announced in January 2002, uses an "automated targeting system" (ATS) that employs rules-based software to identify containers bound for the United States that are at risk of being tampered with by terrorists. A key input to this system is the container's shipping manifest, which contains information about the container's sender, recipient, and contents (CBP 2004). CBP's "24-hour rule" mandates that an ocean carrier transporting a container to the United States forward manifest information to CSI officials at least 24 hours prior to the container's lading onto a vessel that will call on a U.S. port. Once

[1] A shipping container is a metallic box that is typically $20' \times 8' \times 8'$ or $40' \times 8' \times 8'$ in size.

**Table 1 Differences Between CSI and SFI Protocol**

| | CSI | SFI |
|---|---|---|
| Process flow | 24 hours before scheduled departure containers moved from stack[a] to inspection facility | Containers inspected upon arrival to terminal and then taken to stack |
| Technology | Handheld detectors for passive radiation detection; high energy x-ray scanners for nonintrusive imaging | Radiation portal monitors for passive radiation detection; gamma-ray radiography for nonintrusive imaging; optical character recognition to record container ID |
| Decision rules | Targets 5% of U.S.-bound containers, based on manifest information | Inspects every U.S.-bound container arriving at the terminal |
| Equipment location | Somewhere in the interior of the terminal or even at an off-site location | At terminal entrance |

[a]The area where containers are stacked and stored as they wait to be transported.

transmitted, manifests are analyzed at CBP's National Targeting Center in Arlington, Virginia, and containers that are identified as suspect are flagged to be inspected by the local customs authority at the port of origin before they are shipped to U.S. ports. These customs officials use high-energy x-ray radiography (a form of nonintrusive inspection) and handheld, mobile, or stationary radiation detection technology to screen the high-risk containers and ensure that they do not contain a nuclear weapon or radiation dispersal device.[2]

Today about 58 international ports are members of the CSI. According to CBP, about 5%–6% of containers pose a potential risk that warrants a closer review of the associated documentation or a physical examination (Marine Link 2004, McClure 2007). Because of logistical and jurisdiction-related challenges, however, the actual number of containers inspected at international ports is much lower (U.S. Government Accountability Office 2005, 2008a).

Launched in 2007, the SFI is a joint initiative of CBP, the U.S. Department of Energy, and the U.S. Department of State. It is meant to leverage learning from other port security initiatives, such as Operation Safe Commerce, and to serve as a pilot for a system that is capable of scanning 100% of U.S.-bound containers (U.S. Government Accountability Office 2008b).

Under the SFI all U.S.-bound containers arriving at participating overseas seaports are scanned with both nonintrusive radiographic imaging and passive radiation detection equipment placed at terminal entrance gates. Optical Character Recognition is used to identify containers and classify them by destination. Sensor and image data gathered through this "primary" inspection are then transmitted in near real time to the National Targeting Center in Virginia. There, CBP officials incorporate these data into their overall scoring of the risk posed by containers and target high-risk containers for further scrutiny overseas. Any container that triggers an alarm during primary inspection is automatically deemed to be high risk.[3] High-risk containers then undergo a more sensitive "secondary inspection."

The SFI program has been piloted for one year on a full scale in three small, international ports: Karachi (Pakistan), Puerto Cortes (Honduras), and Southampton (United Kingdom). It has also been implemented on a limited capacity basis in four additional ports, including Hong Kong (U.S. Government Accountability Office 2008b).

We observe that the CSI and SFI protocols differ along multiple dimensions, including the pool of containers targeted for inspection, the timing and location of inspections, and the equipment used to perform the inspection. We summarize these differences in Table 1.

CSI and SFI are the two security schemes that, to date, have already undergone extensive field testing, and they are the most promising approaches that the U.S. government is likely to support going forward. Therefore, our analysis focuses on them and their likely variants.

### 1.2. One-Hundred Percent Scanning Requirement

The immediate need for this study arises from a U.S. law enacted in August 2007, "Implementing Recommendations of the 9/11 Commission Act of 2007" (Pub. L. No. 110-53), popularly called the 9/11 Commission Act. The law requires that, before *any* cargo bound for the United States is loaded onto a ship at an international port, it must be scanned using nonintrusive imaging (NII) and radiation detection technology to detect radiological contraband. The deadline for compliance with this law is July 1, 2012, unless the Secretary of Homeland Security grants extensions, which can be offered in two-year increments (U.S. Congress 2007).

This law is a significant deviation from CBP's CSI approach of scanning only cargo it identifies as being

[2] Some ports may be using gamma-ray radiography as an alternative to high-energy x-ray scanning.

[3] Throughout this paper, we use the word "alarm" to mean an outcome that triggers a decision to further inspect a container. In some cases, there may be a physical alarm system and in others, not.

high risk, and the operational feasibility of 100% NII scanning has been questioned by a wide range of participants in the maritime supply chain: CBP and European customs officials, trade associations such as the U.S. Chamber of Commerce and the National Association of Manufacturers (NAM), and corporate leaders. The most common concern is that the congestion that would result from this security requirement will substantially increase the cost of doing business and hurt commerce (Field 2009).

**1.2.1. Objectives of 100% Scanning.** An obvious goal of 100% container scanning is to detect and neutralize any nuclear weapons and to curb the illegal movement of radiological material that may be present in a U.S.-bound container. It does not necessarily aim to replace the risk-based targeting approach followed under regimes currently in practice, but supplement the targeting algorithm with data gathered through the scans. A stringent security regime also serves to deter terrorists from attempting to infiltrate the maritime supply chain in the first place. A less obvious benefit of 100% scanning is associated with disaster recovery. In the case that an unfortunate event were to occur, to contain losses and resume port operations quickly, it would be imperative to identify the stage in the global supply chain at which the security breach occurred. The images and scan information gathered through 100% scanning would provide vital information to facilitate this task (CBP 2008).

**1.2.2. Potential Supply Chain Impacts.** There are essentially three broad ways in which the 100% scanning requirement could be potentially detrimental to trade.

First, if there is limited scanning and radiation detection capacity, the delays resulting from waiting in inspection queues could require containers to sit idle at ports for durations that are longer than required in the absence of inspections. These extra delays would lead to increases in transportation lead times, resulting in higher inventory levels in supply chains, and ultimately in higher cost for consumers.

Second, there could be an adequate level of scanning and radiation capacity but if the NII equipment generates more alarms than there is human inspection capacity to resolve, then the result would again be delays as containers wait in inspection queues.

Finally, the need to divert containers from their usual movements within terminals, redirecting them through a centrally managed government inspection facility, has the potential to engender significant terminal congestion; that is, even if more-than-adequate investments are made in NII equipment, the disruption of the process by which containers are moved into and out of terminals can, itself, lead to significant increases in the time and space required for terminals to process the containers that pass through their gates and quays. These delays could adversely impact the annual throughput of terminals, turnaround times for ocean-going vessels, and turnaround times for inland transportation (trucks, trains, etc.), thereby hurting their efficiency. Again, these decreases in efficiency, along with increased lead times, would lead to higher consumer costs (see Policy Research Corporation 2009).

### 1.3. Evaluating the Impact of 100% Scanning on Terminal Operations

If not carefully considered, efforts to satisfy the requirement to scan 100% of U.S.-bound containers have the potential to significantly degrade the efficiency of container terminals and, more broadly, maritime supply chains. Given the economic importance of maritime trade, a rigorous quantitative analysis of the impact of 100% scanning on container terminal operations would be of great value to policy makers, as well as to companies with an economic interest in the efficient movement of containers within the international supply chain. In this paper we perform just such an analysis.

Our study is based on detailed data on the movement of individual containers, collected from two of the world's largest international container terminals. Among other features, these data sets mark the entry and exit times of every container passing through each of the terminals over the course of one month, along with an indication of whether or not the container is bound for the United States. Between the two ports, we have movement records for more than 900,000 containers.

We use these historical records as the basis for a simulation analysis that estimates the effect of a number of inspection protocols on terminal operations. More specifically, during the time over which the data were collected, inspections had no material effect on container movements, and we utilize the historical records of entry and exit times as a baseline for the timing of container movements. Using discrete-event simulation (Law and Kelton 2007), we then overlay simulated inspection processes on top of this historical record.

To estimate the effect of inspections on the flow of containers through the two terminals, we compare the output of the simulated inspection system to the historical records. For each container that undergoes an inspection, we compare the time it completes the simulated inspection process with the time it was loaded onto the vessel bound for the United States. If the simulated completion time falls beyond the actual loading time, then we mark the container as being delayed and record the difference in times. If the completion time falls before the actual loading time, then we mark the container as not delayed.

The simulations also provide us with insight into the impact the inspection process may have on space requirements for container terminals. Numbers of containers waiting to be inspected can be translated into square feet or meters required to stage the containers, and in each simulation we track the average, as well as peak, staging area required over the course of the period of simulation.

Finally, each simulation run makes explicit assumptions concerning the numbers and types of equipment involved. We estimate these equipment costs, as well as associated personnel costs, and we calculate the handling cost per container of the inspection schemes we consider.

### 1.4. Results and Implications

Our simulation results suggest that a variant of the SFI inspection scheme, which we refer to as an "industry-centric" inspection scheme, is capable of being scaled up to satisfy the scanning and radiation detection requirement mandated by the 2007 U.S. law.[4] Its use of rapid screening by relatively low-cost drive-though portals allows it to handle 100% of *all* container traffic—that is bound for the United States, as well as other destinations—on a cost-effective basis. In turn, the relatively small percentage of containers that fail this rapid primary inspection can be scanned in a cost-effective manner by more sensitive drive-though equipment.

In contrast, the current CSI protocol, which relies on more sensitive equipment to scan high-risk containers in a centrally located, government-operated inspection facility, would face significant hurdles were it to be scaled up to scan more than a small fraction of U.S.-bound container traffic. Given the capacity of the scanning equipment currently available at CSI ports, our simulations show that, for the ports we study, it is possible to support passive radiation detection and NII of no more than 5% of U.S.-bound traffic at the smaller port, and no more than 1.5% of U.S.-bound traffic at the larger one. Even if the capacity of scanning equipment were to be scaled up—by a factor of 20 or 67—to accommodate 100% scanning, the associated per-container costs would be an order of magnitude higher than those incurred under the industry-centric scheme.

The economy and robustness with which the industry-centric scheme operates follows, in large measure, from the type of equipment used. The current CSI protocol relies on highly sensitive high-energy x-ray radiography to scan containers that are thought to pose a potential threat. This is a time-consuming procedure. In contrast, the industry-centric inspection scheme performs a rapid initial scan of 100% of inbound traffic with lower-cost drive-through radiation and medium-energy x-ray radiographic portals. Although this equipment is less sensitive than that used under CSI, it is precise enough to verify the safety of the vast majority of containers, thereby reducing the demand on more sensitive inspection equipment. A more detailed discussion regarding the relative security effectiveness of the two regimes is provided in §7.3. Our simulation results clearly imply that the equipment and inspection protocol used in the industry-centric scheme are relevant in guiding the choice of the appropriate inspection regime for international ports.

Furthermore, a qualitative analysis of the two schemes' logistical requirements also suggests that disruptions to terminal operations would be much more severe under the CSI than the industry-centric approach. Under the CSI scheme, containers targeted for inspection must be pulled from a terminal's storage stacks only hours before the time at which they normally would be retrieved for their vessel loadings. This disrupts the highly optimized sequence in which terminals order yard cranes' movements within the stacks. Under the industry-centric scheme, in contrast, targeted containers undergo inspection upon arrival to the terminal, before they are placed in the stacks. Thus, the industry-centric inspection regime avoids the disruptions and delays that would follow from the early removal of even a small fraction of containers from the terminal's stacks.

The remainder of this paper is organized as follows. Section 2 reviews literature and data sources relevant to our study. In §3 we describe the steps in the container flow through a terminal when there are no inspections. In §4 we describe our research methodology, including a description of the data set and of the design of the simulation study. Section 5 discusses the model used for the simulation of the CSI regime, along with an analysis of the results and costs associated with its implementation. Section 6 describes the simulations for the industry-centric regime. Finally, in §7 we discuss our results and present our conclusions.

## 2. Literature Review

Questions related to the streamlining of port and terminal operations have received a fair amount of attention in the academic literature (see Steenken et al. 2004 for a literature review). However, issues pertaining to maritime and port security have only recently started to generate interest. Harrald et al. (2004) and Willis and Ortiz (2004) provide a risk-management framework for securing maritime infrastructure. Boske (2006) reviews the major U.S. domestic and international initiatives in this regard.

[4] We call this scheme "industry-centric" because, as we will discuss in §7, its appeal to terminal operators may be great enough that the maritime industry itself would fund it.

Two books by Flynn (2004, 2007) address this topic in depth. The former highlights the United States' overall vulnerability to maritime terrorism, whereas the latter emphasizes the inadequacy of its present set of security initiatives.

In this paper, we specifically look at the trade-off between the security generated by inspections and the resulting system congestion. Previous work on this question has been largely numerical. Wein et al. (2006) characterizes the optimal investment in security infrastructure across foreign and domestic ports. Wein et al. (2007) considers the optimal spatial deployment of radiation detection equipment at international ports. Martonosi et al. (2006) evaluates the feasibility of 100% container scanning at American ports. An analytical treatment of the container inspection policies followed at U.S.-domestic ports can be found in Bakshi and Gans (2010). A shortcoming of all of these studies is that they lack actual data on container movement at ports. Instead, they rely on broad assumptions concerning the probability distributions and summary statistics associated with the arrival of containers to the port and their departure from it.

Our work is closest in spirit to Bennet and Chin (2008), who also aim to understand the policy implications of 100% container inspection at international ports. Like the work above, this paper does not use container-movement data, and it therefore restricts its focus to that of a stylized analysis. In the absence of port-related data, this paper cannot model the dynamics of terminal operations and is compelled to decouple the analyses of the primary and secondary inspection processes. Thus, although the paper is able to verify the potential cost effectiveness of 100% inspection using a stylized SFI inspection system, it cannot consider the specific process requirements of the SFI or CSI protocols. Similarly, it cannot provide an empirically driven view of the performance of CSI or SFI schemes.

The Government Accountability Office periodically reviews the maritime security programs of the U.S. government. Examples of such reports, focusing on container inspections at international ports, include U.S. Government Accountability Office (2005, 2008a). The latest in this series are U.S. Government Accountability Office (2008b, c) reports that highlight the difficulties involved with rolling out a 100% inspection regime. In particular, the U.S. Government Accountability Office (2008c) emphasizes the challenges associated with adapting the risk management approach of allocating resources toward the thorough inspection of only high-risk containers, as embodied by the CSI regime, to the new paradigm of inspecting every single container.

## 3. Container Flow Without Inspections

In this section, we first provide a high-level description of the flow of containers through a typical terminal at a port when there are no inspections. We then discuss how low-level container-placement decisions can affect terminal operating costs.

Generally speaking, a terminal comprises entrance and exit gates, a container yard, and the quayside. Entrance and exit gates provide access to inland transportation for delivery of containers to the terminal and for pickup of containers from the terminal. The quayside provides similar water access for large and smaller vessels. The container yard is the place where containers are stored during their stay in the terminal. In many terminals, stacks of laden containers may be two or three high; in land-constrained facilities they may be five or six high. Empty containers may be stacked even higher.

### 3.1. Baseline Container Flow Without Inspections

No matter what the mode of transportation, the flow of a container through a terminal follows a predictable pattern. A container enters the terminal, sits in the terminal for some period of time, and then leaves. Small numbers of containers are so-called "hot boxes" which, because of time constraints, exit the terminal as soon they arrive. Figure 1 depicts a high-level schematic of the process.

The arrival process varies slightly by mode of transport. For containers that arrive on an external truck, the relevant paperwork is checked at the entrance gates. The truck's driver then drives the container to an assigned location in the yard. A yard crane picks up the container from the truck and places it in its position in the yard stack. Arrivals by barge or large vessel are first unloaded at quayside onto an internal truck, a truck owned and operated by the terminal itself. This internal truck then carries the container to its assigned yard location, and the yard crane puts the container into the stack as before. The process flow and logistics associated with departures is similar to arrivals: it begins with a yard crane depositing a container from the stack onto a truck and ends with the container leaving the terminal from the gate on an external truck or from the quay on a vessel.

**Figure 1** **Baseline Container Flow**

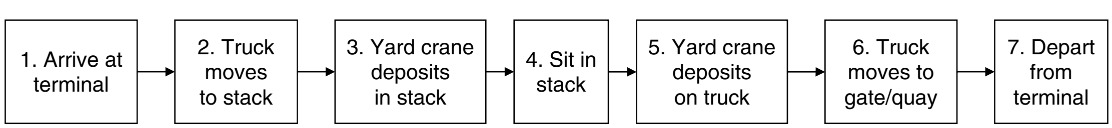

### 3.2. Container Positioning and Terminal Efficiency

At a high level, the typical movement of a container through a terminal is quite simple: it arrives and is placed in the stack, it sits in the stack for some time, and it is pulled from the stack and departs. Nevertheless, because large terminals handle thousands of these containers each day, low-level decisions concerning where and when specific containers are placed in and pulled from the stack can have a significant effect on a terminal's handling costs.

For the purposes of our analysis, two interrelated decisions regarding container movements are worth noting. The first pertains to the number of crane moves required to retrieve a container from the stack. The more containers that sit atop of the container to be moved, the more crane moves it takes to access the container to be moved and (potentially) reposition the containers above. The second regards the grouping of containers within the stack. Containers that are grouped together on a vessel tend to be grouped together within the stacks. That way yard cranes can work in a local area and not waste time moving back and forth across the stacks.

Thus, the positioning of containers within a terminal's stacks has a significant bearing on the labor and equipment time required to load a vessel. Inspections that require pulling containers from the stack on an untimely basis or disrupt the grouping of containers within the stack negatively affect the terminal's costs, as well as the time required for the vessel to stay at the port. The less time required at port, the more efficient the ocean carrier operations will be.

The complexity these container-placement and equipment-scheduling decisions motivates large container terminals to use sophisticated computer algorithms and simulations to help them make these decisions (Steenken et al. 2004). For example, the team that handles terminal design and capacity planning for the company that operates Terminal A consists of about 160 employees.

## 4. Research Methodology

To assess the operational impact of the inspection policies, we create a simulation model of each inspection process that may be followed at a terminal. The model has two elements: historical data on container movements are used to mark times at which containers enter and leave the terminal, and discrete-event simulation is used to track containers as they make their way though the simulated inspection process. In this section, we describe the historical data, as well as how we use those data to drive the simulation models. We also describe how we estimate the per-container costs associated with each inspection regime.

### 4.1. Data Description

In acquiring historical data, we are fortunate in having the support of two of the largest container terminals in the world, which we call Terminal A and Terminal B. The data sets of the two terminals are similar to each other and record the flow of every container that enters or leaves each terminal over the course of one month in the autumn of 2006.

For each container that flows through a terminal, we have a set of time stamps that correspond to the boxes in Figure 1 (except "sit in stack"). For the purposes of this study we concentrate on two: the arrival and exit times. Arrivals by truck are marked at the time the truck enters the terminal's in gate, and arrivals by barge or ocean vessel are marked at the time the quay crane places the container on an internal truck at quayside. Departures by truck are marked as the truck passes through the terminal's exit gate, and those by barge or ocean vessel are marked at the moment the quay crane that will place the container in the ship has latched onto the container at quayside.

Over the course of one month, we have roughly 400,000 to 500,000 records of containers entering and/or leaving each of the terminals. Of these, roughly 40,000 to 85,000 records at each terminal are for containers that arrived early, i.e., before the start of the month during which container movements are tracked. Similarly, about 75,000 to 77,000 records at each terminal are for containers that had not departed by the end of the month.

We retain these records in the data set as they are relevant in certain scenarios that we analyze. Whereas we store the early arrivals "as is" in the data set, we assign the records with departure times after the end of the month a proxy departure time that is beyond the end-of-month horizon for our analysis.

In addition to these time stamps, each container's history also records its destination once it leaves the terminal. This destination field allows us to distinguish U.S.-bound containers from those that are not. In our database, roughly 13% of the traffic at Terminal A is U.S. bound, whereas the corresponding figure for Terminal B is 31%.

Our data capture two important drivers of system performance that are worth noting here: the rate at which containers arrive over time, and the elapsed times between containers' arrivals to and departures from the terminal. We abuse traditional terminal nomenclature and call the latter the container's *dwell time*.[5]

Arrival rates measure the amount of work flowing into the system over time, work that drives the need

[5] Traditionally, a container's dwell time is the time the container sits in the stack. It excludes the time between arrival and placement in the stack, as well as the time between retrieval from the stack and departure.

**Figure 2 Arrival Rates and Dwell Times at Terminal A**

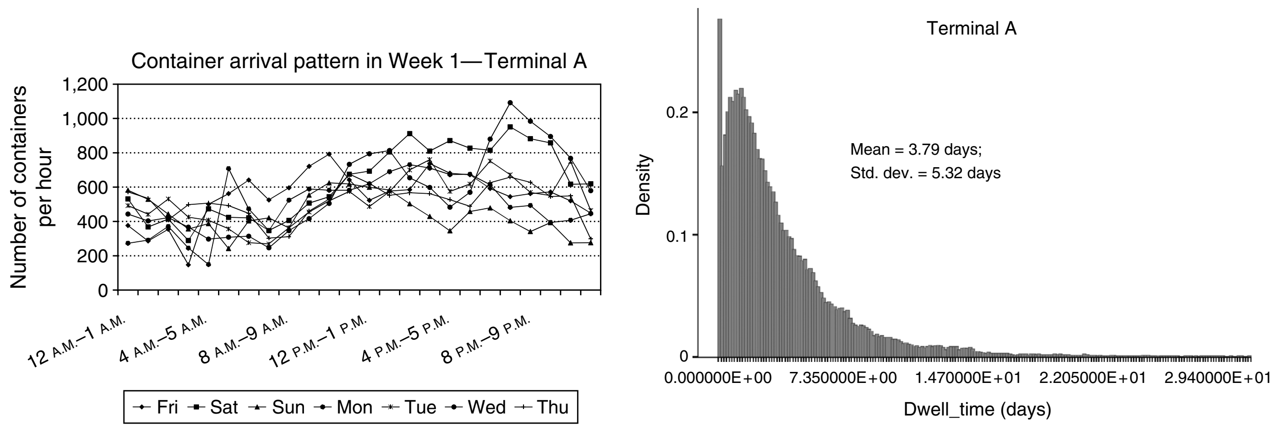


for terminal as well as potential inspection capacity. Although the long-run average arrival rates at the terminals are 500 to 600 per hour, differences in rates can be up to five-fold from one hour to the next. The left panel of Figure 2 shows that, at Terminal A, these rates vary significantly by time of day and day of week. (Terminal B experiences similar fluctuations.) A rigorous analysis of the Terminal A data establishes that arrival process is a time-inhomogeneous Poisson process.[6]

The dwell time provides a measure of how much slack there is in the system, the more slack, the more easily inspections may be completed on a timely basis. The right panel of Figure 2 shows the distribution of dwell times (in days) for containers at Terminal A. The distribution is calculated over all departures that occurred during that month.[7] The spike close to zero reflects a significant number of containers that arrive by vessel and are destined for a location within the country where the port is located. These make their way out of the terminal immediately upon arrival.

### 4.2. Simulation Models

Frequent and significant variation in container arrival rates over time implies that inspection systems that may provide adequate capacity over the course of a month will fluctuate between phases in which the offered load greatly exceeds system capacity and those in which it falls far below. As §5 of Green et al. (2007) notes, these dynamics suggest that traditional, steady-state expressions for queues with constant arrival rates—such as those used in Bennet and Chin (2008), and Wein et al. (2006, 2007)—may not be adequate to capture the terminals' congestion.

Discrete-event simulation provides a simple and robust means of modeling system performance in these settings (Law and Kelton 2007). For each inspection scheme under consideration, we construct a separate simulation model that incorporates the following four elements: (1) the specific actions (such as crane and truck movements and inspection scans) to be taken; (2) the number of "servers" (pieces of equipment or people or combinations of the two) that are dedicated to the performance of each action; (3) the (possibly random) time required for a server to complete a desired action; and (4) the flow of containers from one action to another.

The scope and complexity of terminal operations make it (practically) infeasible for us to replicate them in detail, however.[8] We therefore base the distribution of times required to complete logistics processes, such as internal truck and crane operations, on estimates of typical times provided by managers at the terminals we studied. We note that some distributional variation reflects the effects of terminal congestion: time spent waiting for equipment to become available. The fact that these time distributions are exogenously defined inputs to the simulation (rather than outputs from the simulation) implicitly assumes that the addition of a container inspection process does not (positively or negatively) impact current levels of terminal congestion. We return to this point in §§5.1 and 7.2.

Because the inspection processes we simulate have not yet been implemented comprehensively, we do not have historical data from which we can construct probability distributions for process times. Therefore, we have similarly worked with a security expert,

[6] The Poisson hypothesis was tested, and not rejected, by statistical tests similar to those described in Brown et al. (2005). Subsequently, the hypothesis that the Poisson rate is time homogeneous was rejected using tests similar to those described in Brown and Zhao (2002).

[7] They exclude departures that occurred after the end of the month.

[8] The proprietary nature of terminal operators' detailed simulation models also makes them unavailable for us to use.

Dr. Charles Massey, to develop process-time and outcome distributions for inspection tasks.[9]

Our approach to developing process-time distributions has been similar for the logistics and inspection tasks. In both cases, we have worked with resident experts—the terminal managers or Dr. Massey—to determine mean process times, along with a qualitative judgment concerning whether or not, for a given task, individual times may vary significantly from that mean. If variation is judged to be negligible, then we simulate the process time as deterministic; if it is judged to be significant, then we simulate the process time as lognormally distributed, with a standard deviation is equal to the mean.

We choose this form of distribution for two reasons. First, the lognormal distribution has an appealing form, with a modal response strictly above zero and a long tail that captures infrequent cases of very long process times.[10] Second, by choosing a standard deviation equal to the mean, we set the coefficient of variation of the processing time to be equal to one, a common assumption for relatively highly variable process times.

We note that, although a modification in the specifics of these assumptions may change the numbers that come out of our simulations a bit, they do not affect the qualitative insights our simulations generate. In fact, in addition to the simulation results reported in this paper, we have run simulations with deterministic process times, and the results have remained essentially the same.

The only way to avoid container congestion and queueing is to provide enough inspection capacity to handle the peak arrival rate. But given the significant variability in the arrival rates shown in Figure 2, this so-called "peak-load capacity" would be about twice that required to handle the long-run average load, and the cost of little-used inspection capacity could be significant.

With less than peak-load capacity on hand, however, some queueing will occur. Typically, the lower the capacity, the longer the queues and the longer the average time spent in queue. If too little capacity is available, excessive time in queue can push the time required for containers to complete the inspection process to fall beyond the time of the container's scheduled departure from the terminal. In this case, either the pickup vehicle waits for the delayed container or the container misses its departure. If the means of transportation is an ocean-going vessel, the associated waiting time can be prohibitively expensive, and the outcome is more likely to be a missed voyage.

The use of historical arrival and departure records allows us to show what might happen to actual container flows if the various inspection schemes were to be imposed.[11]

The queueing of containers that are waiting to be processed also poses an internal, logistical burden on terminals. Each 40-foot container takes up 320 square feet of area, and as the number of waiting containers grows, so does the amount of area that must be reserved to handle them. The operators of Terminals A and B estimate that, for containers that are stacked two high, the staging area required for inspections would accommodate about 150 containers per acre. For terminals with constraints on real estate, the required staging area can become prohibitively expensive.

Thus, for each of the inspection schemes that we simulate, we track three metrics: (1) the fraction of inspected containers whose simulated inspection time exceeds its historical departure time; (2) the average amount of time by which containers are delayed beyond their scheduled departure time; and (3) for each capacity-constrained step of the inspection process, statistics regarding queue length.

Furthermore, because large fluctuations in the arrival rates of containers to the terminals drive queues vary dramatically over time, we report three statistics concerning the queue-length distribution over time: the time average, the 99th percentile, and the maximum of the queue length over the course of the month. More specifically, let $q(t)$ be the queue length at time $t$. Then the time average over $[0, T]$ is $1/T \int_0^T q(t)\, dt$, the maximum is $\max_{t\in[0,T]}\{q(t)\}$ and, if $F(q) = 1/T \int_0^T \mathbb{1}\{q(t) \le q\}\, dt$ is the fraction of time the queue length is less than or equal to $q$, then the 99th percentile of the queue length is $\{q \mid F(q) = 0.99\}$.

### 4.3. Interpretation of Simulation Results

We describe two important technical details concerning the simulation analysis and the interpretation of its results. The first concerns the pool of containers over which we calculate statistics. The second pertains to the termination condition for determining the number of simulation runs. For background concerning both of these issues, see Law and Kelton (2007).

[9] Dr. Massey is an independent security consultant. He was formerly the manager of the International Borders and Maritime Security Program at Sandia National Laboratories in Albuquerque, New Mexico. He has extensive experience in the development and implementation of security and transportation programs on the national and international level. He served as the Sandia manager responsible for supporting the Department of Energy/National Nuclear Security Administration's Second Line of Defense Program and its Megaports Initiative.

[10] For example, the dwell times shown in Figure 2 are roughly lognormally distributed.

[11] This approach allows us to include all of the idiosyncracies of the actual container arrival processes at the terminals and, in turn, to more concretely evaluate the feasibility of the inspection schemes under consideration.

First, even though our historical data cover month-long periods, we calculate metrics, such as fraction of containers delayed and average queue lengths, over only a 22- or 23-day period. In real life, the inspection of containers would be an ongoing process, and even at the start of the month there would be containers distributed throughout the inspection process. In our simulations, however, the inspection process starts out empty at the beginning of the month, and simulated monthly averages that included these starting days would understate fractions of delayed containers and average queue lengths.

Therefore, we designate the first seven days of each simulation as a warm-up period, during which the inspection process ramps up. Similarly, to avoid end-of-period effects related to the application of the CSI's 24-hour rule, we exclude containers processed on the last day of each month. Because autumn months have 30 or 31 days, this leaves 22 or 23 days worth of data for the calculation of performance statistics.

Second, for each simulated system we perform multiple simulation runs. More specifically, when we simulate a system for a given month, we calculate multiple performance metrics: the fraction of containers delayed beyond load time, the average delay beyond loading of delayed container, and the average, 99th percentile, and maximum queue length attained over the course of a simulated month. In turn, we perform repeated runs to gather repeated samples of these metrics and report their sample average across runs.

We determine the number of simulation runs as follows. We bound the number of runs below by 25 and above by 100. Within this range, we stop simulating once the half-widths of the confidence intervals of the mean standard errors fall at or below 2.5% of the sample means. Note that in the body of the paper we simply report the estimated averages and do not report half-widths of confidence intervals. We do report these data for the relevant simulations in a separate appendix available with the authors.

### 4.4. Cost Calculations

Each inspection regime requires equipment as well as labor to operate the equipment, each of which carries with it associated costs. An inspection protocol may also generate operational overhead: yard crane moves required to retrieve containers from and replace containers into the stack; truck moves to ferry containers to and from an inspection facility; tophandler moves required to load and unload containers from the trucks at the inspection facility.[12]

To understand how the use of these resources affects supply chain costs on a per-container basis, we use annuities to allocate their costs over a set of containers. The cost of inspection equipment, whenever reported, includes the cost of deployment and initial testing. When evaluating costs accrued by the U.S. government, we use a risk-free rate of $r = 3.7\%$, and for costs borne by private companies, namely, terminal operators, we use an estimated cost of capital for the maritime transportation sector of $r = 7.2\%$ per annum. Changes in discount rate assumptions, such as costing both regimes at the same rate, do not materially affect the results.

## 5. Model and Results for CSI Regime

The CSI inspection regime applies only to U.S.-bound containers and begins 24 hours before a container is to be loaded onto the vessel that will carry it to the United States. According to CBP's "24-hour rule" the responsible ocean carrier must transmit the container's manifest information to CBP so that the risk posed by the container can be evaluated before the container's departure. Note that the 24-hour rule implies that U.S.-bound containers cannot be hot boxes.

Thus, under CSI, there exists a 24-hour window within which a U.S.-bound container has to be inspected, if supply chain lead times are to remain relatively unaffected by security operations. If containers are delayed beyond their scheduled departure time then there may be potentially large cost implications, as explained in §1.2.

Figure 3 provides a high-level schematic of container flow under the CSI regime. Note that the figure's boxes 1–4 correspond to those of the base case, shown in Figure 1. Similarly, boxes 15–18 correspond to boxes 4–7 in the base case. The CSI inspection protocol is outlined in boxes 5–14, with the key inspection step occupying box 11.

At the inspection facility, each container undergoes two forms of inspection: scans with radiation isotope identification devices (RIIDs) are used to detect radioactive emissions; and radiographic imaging, using high-energy (9 MeV) x-ray equipment, is meant to detect the use of lead shielding used to conceal illicit material that emits dangerous isotopes. To date, large CSI ports have operated a single inspection station that uses two handheld RIIDs and one high-energy x-ray radiographic device. Our simulations include results for two variants of this setup, one with a single inspection station and another with two.

Recall that one measure of system performance is the fraction of inspected containers that complete the inspection process before the time stamps of their actual departures. In our simulation models, we define the total inspection time to be the elapsed time between the transmission of the container's manifest

[12] A tophandler is a small mobile crane.

**Figure 3** **Typical Container Flow for CSI Regime**

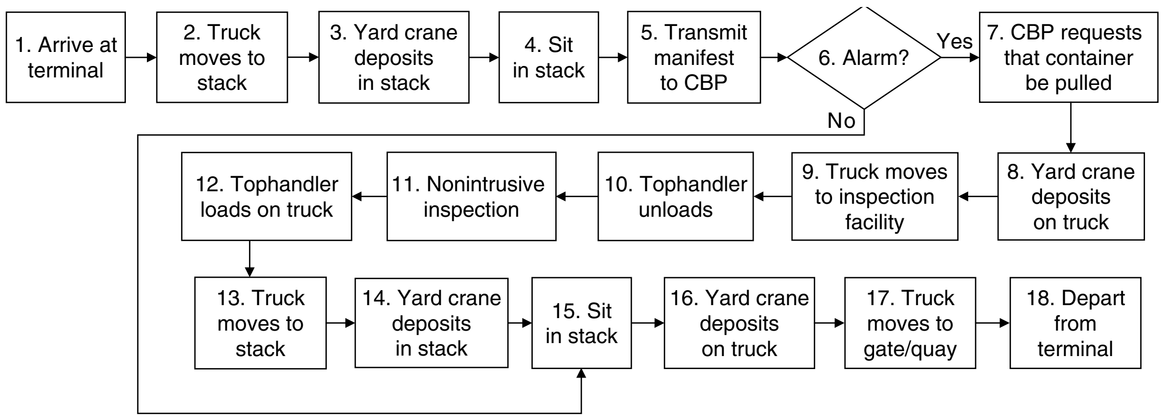


(box 5) and the time at which the container completes its nonintrusive inspection (box 11).[13]

Another measure of system performance concerns the numbers of containers that queue, waiting to be processed at critical process steps. In the CSI simulation, queueing occurs at the nonintrusive inspection step (box 11). At this step we track numbers of containers waiting to be inspected over the course of a simulated month.

### 5.1. Model Details

In this section, we describe the critical statistics associated with our simulation calculations. Statistics regarding inspection steps (box 11 in Figure 3) are developed in collaboration Dr. Massey, and those regarding logistical process steps in consultation with the terminals' managers. We summarize these statistics in Table 2. In a separate appendix, available from the authors, we provide a detailed description of each process step, along with the rationale behind detailed assumptions.

We note the following regarding the process-time statistics reported in Table 2. Variability in the time it takes for a yard crane to pick up a container from the stack and deposit it on a truck (box 8 in Figure 3) arises from uncertainty in the location and/or availability of the crane and the truck.[14] A mean inspection time of 20 minutes (box 11 in Figure 3) is the key source of system congestion and reflects a number of process steps: (1) time for manual scan with handheld RIIDs; (2) time for a truck driver to position the container with respect to the high-energy radiographic equipment, exit the cab, and move away from the inspection station; (3) time to scan the container; and finally, (4) time for the driver to reenter the cab and drive the container away. Variability in the inspection time stems not only from the inspection test itself, but also from the associated logistics of appropriately positioning the container in the inspection station.[15]

Finally, we note that we have implicitly assumed that the false-positive rate due to NII inspection is negligible. We believe that this is indeed the case based on our conversations with security-policy experts and terminal operators. Moreover, the false-positive rate associated with less sensitive equipment

**Table 2** **Process-Time Assumptions Used in Simulations of CSI Protocol**

| Box | Process step | Process time distribution | Mean | Standard deviation | Source |
|---|---|---|---|---|---|
| 5 | Transmit manifest to CBP | Fixed | 60 min | 0 min | TO |
| 6 | Alarm? | | | | |
| 7 | CBP requests container be pulled | | | | |
| 8 | Yard crane deposits on truck | Lognormal | 15 min | 15 min | TO |
| 9 | Truck moves to inspection facility | Fixed | 40 min | 0 min | TO |
| 10 | Tophandle unloads truck | | | | |
| 11 | NII inspection | Lognormal | 20 min | 20 min | CM |
| 12 | Tophandler loads truck | | Omitted | | |
| 13 | Truck moves to stack | | Omitted | | |
| 14 | Yard crane deposits in stack | | Omitted | | |

*Note.* TO, terminal operators; CM, Dr. Charles Massey.

[13] If, in the simulation, a container completes its inspection shortly before its historical departure time, then we count it as not delayed. At this point it might as easily be brought directly to the quay on a timely basis.

[14] As we noted in §4.2, the fact that this distribution is an input to, rather than result of, our simulation follows from the high-level nature of our model, and it reflects an implicit assumption that the inspection process does not affect the distribution of time it takes to complete these other terminal activities.

[15] Alternative assumptions regarding standard deviations of logistics and inspection times have little effect on resulting system performance.

**Figure 4** Delay Statistics for Terminal A Under CSI Regime

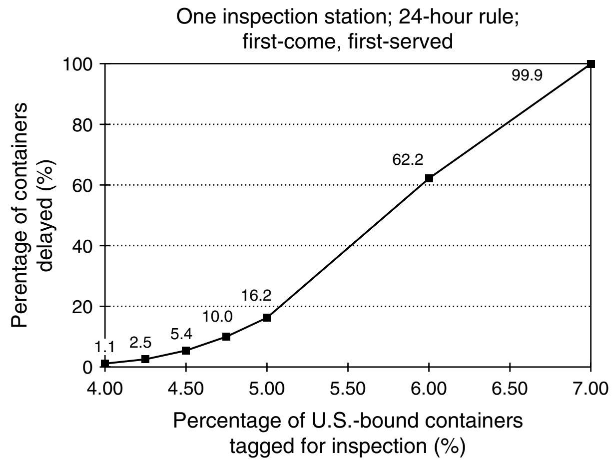


used in SFI pilots has been found to be negligible as well.[16]

## 5.2. Simulation Results

In this section, we present our simulation results for the CSI inspection regime. We vary the percentage of U.S.-bound containers that are inspected and then track the key performance metrics described in §4.3. For each percentage, we assume that every U.S.-bound container is tagged for inspection with a probability equal to the inspection rate, independently of all other containers, and we run two simulations: one with one NII station and another with two. The figures below summarize our main findings for Terminal A. The authors' appendix reports the 95% confidence intervals for all relevant statistics.

Figure 4 plots the fraction of inspected containers that miss their scheduled departure, along with the average time by which the departure time is exceeded (average excess delay). The left panel shows that, for one NII station at Terminal A, a 4% inspection rate generates nearly no delayed containers, and the right plot shows that the few containers that miss their load times experience minimal average delays. With a 5% inspection rate—the rate at which CBP targets containers for inspection—16.2% of tagged containers would have missed their historical vessel loading times, and the average delay beyond loading would have been nearly seven hours. As the inspection rate rises above 5%, the fraction of inspected containers that are delayed explodes, and by 7%, the fraction of inspected containers that are delayed climbs above 99%. In fact, at a 7% inspection rate, the NII station reaches 100% utilization over the month, and further increases in the inspection rate have no impact.

**Figure 5** Staging Area Required: Terminal A Under CSI Regime

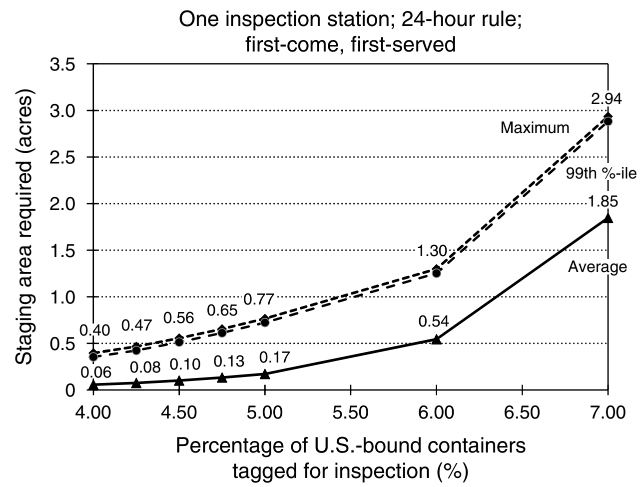


With two NII stations, the figures double. An 8% inspection rate can be supported with nearly no delays, and as the rate increases to 13%, the two stations' utilization again hits 100%, and the fraction of delayed containers explodes.

Figure 5 plots the staging area, in acres, required to handle queues of containers waiting to be inspected. To translate acres into numbers of 40′ containers, one can multiply by 150, an industry rule of thumb for estimating the average number of containers that can be accommodated per acre when container stacks are a maximum of two high.[17] At a 4% inspection rate, the average queue size is negligible, and the maximum is 0.4 acres. At a 5% inspection rate, the maximum nearly doubles, and at a 6% rate it more than triples.

[16] According to Fairnie (2008), the number of containers that originated in Southampton and were physically inspected in U.S. ports (during the SFI pilot at Southampton, UK) went down from 1,000 to 7 over equivalent time periods (nearly five months).

[17] There are 43,560 square feet to an acre, and a 40-foot container has a footprint of 320 square feet. This translates to 136 containers per acre stacked one high, but it neglects square footage between containers required for lanes in which trucks and tophandlers move as they ferry containers into and out of the waiting area.

**Table 3 Cost per Inspected Container at Terminal A, 5% Inspection Rate, One NII Station**

| | | Equipment | | | Labor | | | Cost per container | |
|---|---|---|---|---|---|---|---|---|---|
| Box | Process step | Unit cost ($000) | Life (years) | Source | FTEs | Comp. ($000) | Source | Equip. ($) | Labor ($) |
| 5 | Transmit manifest | n/a | | n/a | n/a | | n/a | n/a | |
| 6 | Alarm? | n/a | | n/a | n/a | | n/a | n/a | |
| 7 | CBP requests | n/a | | n/a | n/a | | n/a | n/a | |
| 8 | Yard crane deposits | 1,200 | 25 | TO | 1.2 | 50 | CM | 1.04 | 2.18 |
| 9 | Truck moves | 20 | 10 | TO | 6.42 | 15 | CM | 0.18 | 3.49 |
| 10 | Tophandler unloads | 450 | 7 | TO | 0.8 | 25 | CM | 0.69 | 0.73 |
| 11 | NII inspection | | | | 20 | 20 | | | 32.59 |
| | 2 RIID units | 75 | 5 | CM | | | CM | 1.19 | |
| | 1 9MeV x-ray | 7,500 | 5 | CM | | | CM | 58.31 | |
| 12 | Tophandler loads | 450 | 7 | TO | 0.8 | 25 | CM | 0.69 | 0.73 |
| 13 | Truck moves | 20 | 10 | TO | 6.42 | 15 | CM | 0.18 | 3.49 |
| 14 | Yard crane deposits | 1,200 | 25 | TO | 1.2 | 50 | CM | 1.04 | 2.18 |
| | Total | | | | | | | 63.34 | 45.38 |

*Note.* TO, terminal operators; CM, Dr. Charles Massey.

For Terminal B, the situation is more extreme. With one NII station, a 1% inspection rate generates few delays, but a 2% rate drives NII utilization to 100%, with nearly all containers being delayed. At the intermediate inspection rate of 1.5%, the fraction of delayed containers is 13.15%. With two NII stations these figures double. That Terminal B's thresholds are lower than those of Terminal A is accounted for by two facts: first, the fraction of U.S.-bound containers is higher at B (31% versus 13%); and second, the overall volume of container traffic is a bit greater at B. Thus, the total number of U.S.-bound containers is more than 138% higher at B, and each NII station can handle a proportionately lower fraction of the offered traffic.

### 5.3. Cost Estimate

Under the CSI protocol, each U.S.-bound container that is inspected passes through steps 5–14, and we can estimate the costs associated with each of these process steps. We recall from §4.4 that, given equipment costs and lifetimes, as well as annual cost per full-time employee (FTE) for labor costs, the essential elements required to allocate costs on a per-container basis are process times and numbers of pieces of equipment or FTEs.

Table 3 provides this information. The terminals have provided estimates for the process times, costs, and lifetimes of the logistics equipment involved in these steps. We have estimated personnel costs for the logistical steps in collaboration with Dr. Massey.[18] For the inspection processes, Dr. Massey has provided relevant equipment and personnel costs, as well as estimates of equipment lifetimes.

In most cases, labor costs are straightforward: they represent the number of FTEs along with cost per FTE of the operators needed to run the associated pieces of equipment. One important exception is the labor required to operate the NII inspection station. We assume that a team of five people per shift is required to man the inspection station. With three shifts a day, and a backup team of five people, the total staffing requirement at the inspection facility is 20.

The final two pieces of data needed to allocate inspection and incremental logistics costs to containers are the number of containers inspected and the discount rate. Here, we assume that U.S.-bound containers are inspected at a 5% rate at Terminal A and at a 1.5% rate at Terminal B, the respective maximum rates that can be handled. We also assume that U.S. government money is used to finance the inspection scheme and apply an annual risk-free rate of 3.7%.[19]

The far right two columns Table 3 show the amortized cost per inspected container for Terminal A, given a 5% inspection rate. Here, the total is $108.71 per inspected container, with $63.34 representing equipment costs and $45.38 representing labor. If we allocate these costs over all U.S.-bound containers, rather than just the 5% that are inspected, then per-container cost drops to $5.44. With two servers and a 10% inspection rate, the analogous cost numbers are $92.42 and $9.24 per container, respectively.[20]

[18] The management of Terminals A and B viewed their personnel costs as proprietary and have declined to provide us with estimates for them.

[19] If the inspection scheme were financed by members of the maritime industry, then a 7.2% discount rate would be applicable. This change in discount rate does not significantly change the numerical result.

[20] As the fraction of containers inspected increases, the unit cost reduction obtained by allocating costs among all U.S.-bound containers naturally decreases. With 100% inspection there would be no difference between the two figures.

**Figure 6** Typical Container Flow for Industry-Centric Regime

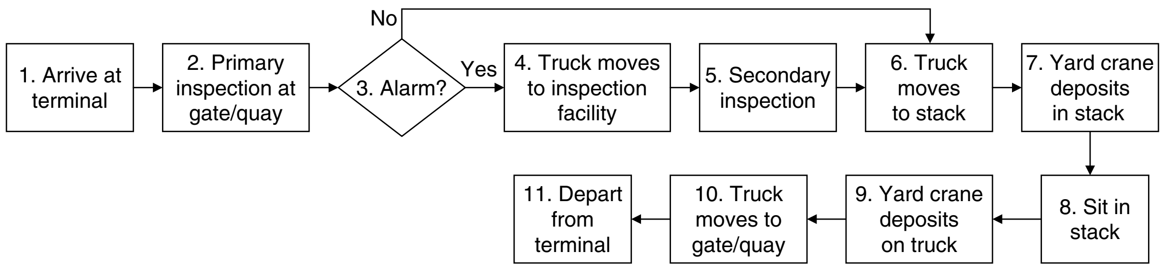


For Terminal B, the amortized inspection costs are $131.18 per inspected container when 1.5% of the containers are inspected with one server. When cost is shared by all U.S.-bound containers, the expense drops to $1.97 per U.S.-bound container. With two servers and 3% inspection rate, the cost figures are $110.91 and $3.33, respectively.

## 6. Model and Results for the Industry-Centric Regime

In the industry-centric protocol, containers undergo a primary inspection upon their arrival to the terminal: at the in gate for containers entering via truck and at quayside for those arriving via vessel. This implies that inspections do not interfere with containers' movements into and out of the stack, and therefore impose minimal additional burden on terminal equipment and operations.

Because inspections are performed on arrival, the entire dwell time of a container is available to complete the inspection process, as opposed to just 24 hours in the case of the CSI regime. This provides the maximum possible slack for completing the inspection of containers without delaying their onward journey.

Figure 6 provides a high-level schematic of the container flow in the industry-centric regime. The inspection protocol is outlined in boxes 2–5. Boxes 6–11 correspond to boxes 13–18 in the CSI protocol shown in Figure 3.

Primary inspection (box 2) is comprised of a drive-through scan using a radiation portal monitor (RPM) followed by another drive-through NII scan using medium-energy x-ray radiography, and we refer to a pair of RPM and x-ray portals as a single inspection station. The number of these stations is a critical determinant of system congestion, and in our simulations we vary the number of primary inspection stations, from two to four at the gate and from three to five at the quay.

Containers that are scanned without incident proceed to the stack (box 6), whereas those that trigger alarms are pulled for secondary inspection (box 5). During secondary inspection a container is subject to a drive-through scan by an Advanced Spectroscopic Portal (ASP), followed by another medium-energy x-ray drive-through scan. We assume that one set of inspection portals (ASP and x-ray) is available for secondary inspection. This capacity level proves to be more than sufficient, because the secondary-inspection workload turns out to be quite small in the industry-centric regime.[21]

The system performance measures used in this context are the same as those used for the CSI regime: delay beyond historic load time and queue length. In the industry-centric simulation, queueing occurs at the primary inspection step (box 2) and the secondary inspection step (box 5). In each simulation, we track the queue lengths attained at both of these steps.

### 6.1. Model Details

As with the CSI simulation, we have estimated statistics related to industry-centric inspections in collaboration with Dr. Massey. Table 4 summarizes the critical inspection-time statistics.

Several characteristics of Table 4's estimates are worth noting. First, the mean inspection times are much lower in the industry-centric regime, when compared to that under CSI: less than half a minute for primary inspection and less than two minutes for secondary. These durations reflect the time it takes to drive a truck through a set of portals. Second, because there is relatively little variation in the drive-through associated with primary inspections, we assume a standard deviation of zero. Third, the mean time required for secondary inspection is double that of primary inspections, and the standard deviation of secondary-inspection time is significant as well. Here, more careful environmental control and slower drive-through time are needed to resolve any uncertainty emerging from the primary scans.[22]

Under the industry-centric regime, a large part of the primary inspection's anomaly detection process

[21] Any alarms that remain unresolved after secondary inspection can then be handled by the backup high-energy x-ray scanner, or manual inspection, if required.

[22] Alternative assumptions regarding standard deviations of inspection times have little effect resulting system performance.

**Table 4 Process-Time Assumptions Used in Simulations of Industry-Centric Protocol**

| Box | Process step | Process time distribution | Mean[a] (sec) | Standard deviation[a] (sec) |
|---|---|---|---|---|
| 2 | Primary inspection at gate/quay | Fixed | 25 | 0 |
| 3 | Alarm? | | | |
| 4 | Truck moves to inspection facility | | omitted/ negligible | |
| 5 | Secondary inspection | | | |
| | ASP scan | Lognormal | 45 | 45 |
| | Medium-energy x-ray scan | Lognormal | 45 | 45 |

[a]Estimates are based on inputs from Dr. Massey.

must be automated. It is worth noting that, in initial SFI pilots, it took approximately three to five minutes to scan containers (CBP 2008). In contrast, we have assumed that primary scans take 25 seconds. This discrepancy reflects differences in the interpretation of scans: in the pilots, CBP officers manually interpreted scans, whereas in our analysis we have assumed that alarm generation is an automated process. We note that CBP (2008) states that there has been good progress on this front, although the technology is still in pilot testing. Furthermore, the 9/11 Commission Act specifies the absence of "automated notification" of alarms as one of the possible reasons for extending the 2012 deadline for implementing 100% scanning at international ports. Given this background, we consider it appropriate to conduct our analysis assuming the availability of automated alarm generation. However, continued progress must be monitored closely.

Finally, we assume that the chance that a primary inspections triggers a secondary inspections (box 3) is 5.95% and is independent and identically distributed across containers. This reflects a 5% alarm rate for the RPM, a 1% alarm rate for the medium-energy x-ray, and an assumption that the triggering of one alarm is independent of another. Although in practice alarm rates are driven by the mix of goods that flow through a terminal and can vary from one port to another, these estimates appear to be quite conservative (CBP 2008, Policy Research Corporation 2009). We implicitly assume that the alarm rate at the secondary inspection stage is small enough that it negates the need to explicitly model congestion at the "backup" high-energy x-ray scanner. Data from the SFI pilot at Southampton Container Terminal (United Kingdom) supports this assumption (Fairnie 2008).

### 6.2. Simulation Results

In this section, we present our simulation results for the industry-centric inspection regime. For the simulations we fix the number of secondary service stations at one and vary the number of primary servers from two to four at the gate, and from three to five at the quay. As before, the performance measures of interest are the number and degree of container delays, as well as the sizes of the queues of containers waiting to be inspected.

We find that, in all cases, the number of containers that miss their scheduled departure is steady at approximately 2,500. These containers are "hot boxes" that have very short dwell times, under six hours. Moreover, none of these containers is U.S.-bound. It is reasonable to expect that, if the industry-centric regime were operationalized, then these short-dwell-time containers would not, in fact, be delayed; that is, knowing that a container inspection process may require a few hours of time, logistics companies would move up scheduled drop-off times so that these containers would not be delayed. We conclude that the delay of containers beyond their historical dwell times does not appear to be a significant problem.

Under the industry-centric regime, there are three points at which containers may wait for inspection: primary inspection stations at the in gate (for trucks) and at the quay (for vessels), as well as a secondary station; hence, there are three sets of queues. We report queue-length statistics for each.

Figure 7 plots the time average, 99th percentile, and maximum of the queue lengths attained during primary inspection, as a function of the number of inspection stations deployed. The figure's left panel shows the results for the in gates at Terminal A, and the right panel those for the quay at Terminal A. We find that, using nine sets of portals as primary servers (four at the gate and five at the quay), time-average queue lengths are quite acceptable: 0.6 containers at the gates and 0.1 containers at the quay. Because of extreme arrival-rate peaks, however, the maximum queue lengths over the course of the month can be large, even with nine servers: 145 at the gate and 34 at the quay.[23]

Whereas a maximum of 34 containers at the quay may be manageable, that of 145 at the in gate is more significant, and this can have serious economic consequences. More specifically, although there is a natural staging area for containers in the drive-up to the terminal entrance, the tolerance for congestion beyond a certain threshold can be quite low at ports. Incoming trucks backed up at a terminal's entrance can impede the flow of city traffic and lead to complications for the local administration. Furthermore, because the queues are comprised of containers on

[23] For the industry-centric simulations, the primary inspection process inspects every container and assumes deterministic process times, and hence confidence intervals are not relevant.

**Figure 7** **Queue-Length Statistics for 1° Inspection at Terminal A, Industry-Centric Regime**

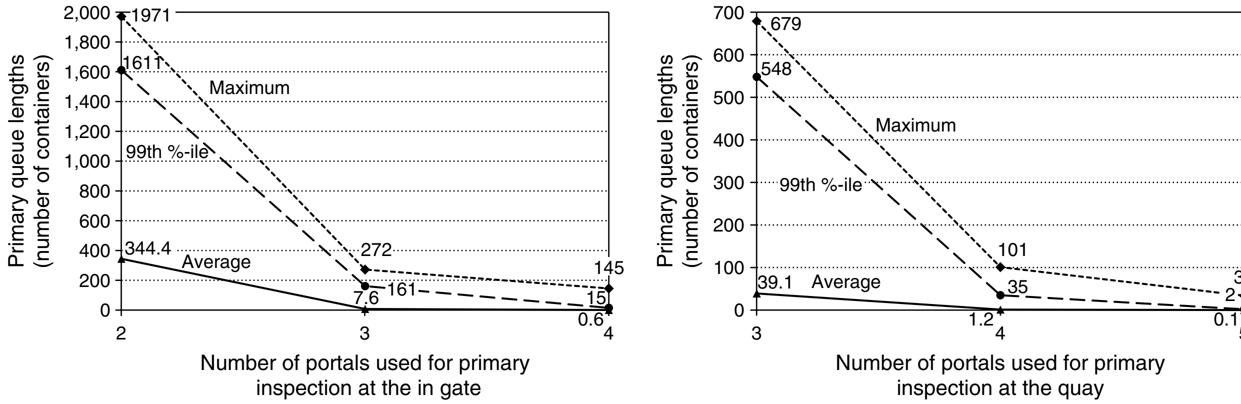


trucks, the containers cannot be stacked, and the associated real estate requirements can be large.

At the same time, we note that the 99th percentiles of the queue lengths are much smaller than the maxima, an indication that the maxima are driven by a short-term spike in arrivals. In contrast, Figure 5 shows that, for the CSI, the 99th percentiles of the queue lengths are quite close to the maxima. We conjecture that the CSI's use of the 24-hour rule effectively smooths the effective arrival rate.

The gap between the 99th percentile and the maximum suggests that active management of in-gate traffic—by the use of call-ahead and appointment systems for inbound trucks—has the potential to significantly moderate these arrival spikes in the industry-centric scheme. To the extent that this is the case, then the 99th percentile figures of 15 at the gate and 2 at the quay may be more indicative of the waiting space required.

In contrast, even with only one inspection station, the secondary inspection queue is not significant. To understand why this is the case, recall that the average secondary inspection time is 90 seconds, so that a single station can handle about 40 inspections per hour. Because fewer than 6% of arriving containers are inspected, a single secondary station can handle an average arrival rate of more than 600 containers per hour ($40 \div 0.06 = 667$), or more than 430,000 containers per month.

Of course, over the course of a month, the arrival rate varies considerably from hour to hour, and Figure 8 shows how the length of the secondary queue varies with the total number of *primary* inspection stations used at the in gate plus quay: $2+3$, $3+4$, and $4+5$. Here, the average queue length holds steady at around two containers, the 99th percentile varies between 17 and 25, and the maximum between 30 and 43. In all cases, a single station is adequate to prevent significant delays.

Our simulation results for Terminal B are similar. With nine sets of primary portals—four at the gate

**Figure 8** **Queue-Length Statistics for 2° Inspection at Terminal A, Industry-Centric Regime**

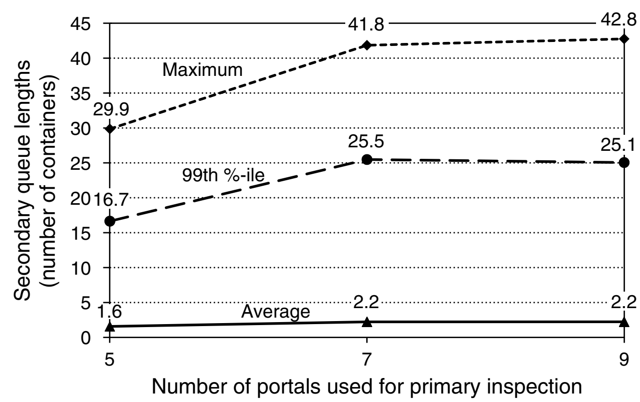


and five at the quay—the time-average queue lengths are 1.06 at the gate and 0.24 at the quay, the 99th percentiles are 19 and 5, and the maxima are 103 and 52, respectively. With nine primary stations, the secondary-inspection queue has an time average of 0.7, a 99th percentile of 8, and a maximum of 17.

Finally, we note that above analysis also provides us with a means to trade off the cost of additional inspection capacity against the real-estate cost associated with the staging area outside the primary inspection facilities. For instance, a move from three to four primary servers at the in gates in Terminal A cuts the 99th percentile of the queue length from 161 down to 15. Suppose a waiting area can house 100 trucks per acre.[24] Then, roughly speaking, a cost of \$4.25 million in additional inspection equipment affords a savings of $\approx$1.5 acres required for trucks to queue as they wait to carry the containers through the terminal in gate.

[24] With 43,560 square feet per acre, there are roughly 136 40-foot by 8-foot containers—without chassis—per acre. For containers that are mounted on trucks, this density would drop further.

**Table 5 Cost per Inspected Container at Terminal A: Nine Primary Stations, One Secondary Station**

| | | | Equipment | | Labor | | Cost per container | |
|---|---|---|---|---|---|---|---|---|
| Box | Process step | Average time (sec) | Unit cost ($000) | Life (years) | FTEs | Comp. ($000) | Equip. ($) | Labor ($) |
| 2 | Primary inspection | 25 | | | 4 | 10 | 2.09 | 0.01 |
| | 9 RPMs | | 250 | 5 | | | | |
| | 9 medium x-rays | | 4,000 | 5 | | | | |
| 3 | Alarm? | n/a | n/a | | n/a | | n/a | |
| 4 | Truck moves | n/a | n/a | | n/a | | n/a | |
| 5 | Secondary inspection | | | | 4 | 20 | 11.10 | 1.40 |
| | 1 ASP | 45 | 600 | 5 | | | | |
| | 1 medium-energy x-ray | 45 | 4,000 | 5 | | | | |
| | 1 9 MeV x-ray (backup) | | 7,500 | 5 | 5 | 45 | | |
| | Total | | | | | | 13.19 | 1.41 |

*Notes.* All estimates are based on inputs from Dr. Massey. Comp., compensation.

### 6.3. Cost Estimate

As with the CSI, the industry-centric regime gives rise to equipment and labor costs that we allocate over the set of inspected containers. In Table 5 we provide the details of the cost breakdown mapped onto the process steps outlined in Figure 6.

We note that, in addition to the equipment costs associated with the primary and secondary process steps required in the industry-centric protocol, we include the equipment cost of the same 9 MeV x-ray station that is currently used in the CSI regime, along with the labor cost of one associated five-member inspection team. Although, in the vast majority of cases, any questions that arise in the context of rapid primary scans can be resolved by a more sensitive secondary scan, in isolated cases a high-energy scan may also be deemed necessary. Thus, even though it is likely to be little used, we include its associated equipment and labor costs. In §7.3, we discuss the security benefit associated with this approach.

Besides differences in inspection equipment, the logistics of the process for the industry-centric regime differs from that for the CSI in the following ways. For the industry-centric regime, we assume that the external or internal truck that brings the container to the terminal stack would carry the container to the secondary inspection facility, if required. Hence, the terminal need not retrieve the container from the yard or transport it by truck to the inspection facility, and it incurs no operational overhead on account of the inspection process. Because the secondary inspection is conducted in a drive-through fashion, there is also no need for tophandlers to unload and load the containers from and onto trucks.

A last difference between the costs found in Table 5 and those calculated for the CSI concerns the discount rate. For the CSI, we assumed that the U.S. government would purchase and run the necessary equipment. For the industry-centric regime, however, we assume that the terminal operator buys and operates the equipment. Therefore, the discount rate we use for the industry-centric regime is 7.2%, rather than 3.7%, which reflects the terminal operator's higher cost of capital.

Table 5's results show that, even with a backup 9 MeV x-ray station and a higher discount rate, the per-container costs under the industry-centric regime remain much lower than those under the CSI. At Terminal A, primary inspection costs roughly $2.10 per container, whereas the total cost is about $14.61 for each container that undergoes both primary and secondary inspection. The per-container cost declines to just $2.97 if secondary inspection costs are allocated across all container traffic, as opposed to just inspected containers. At Terminal B, primary inspection costs $1.80 per container. For containers undergoing both primary and secondary inspection, the per-container cost is $12.52, and if the secondary inspection cost is allocated across all container traffic, then the per-container cost turns out to be $2.54.

## 7. Discussion and Conclusions

Through our simulations we have contrasted the operational performance of the CSI and the proposed industry-centric regime. Although the differences between the CSI and SFI are summarized in Table 1, the industry-centric regime is an adaptation of the SFI regime in the following key ways. The scope of SFI is typically limited to U.S.-bound containers, whereas the industry-centric regime targets every container, irrespective of destination. In terms of technology, the SFI pilots focus mainly on primary inspection using RPMs and gamma-ray or x-ray radiography, whereas the proposed industry-centric regime uses RPMs and medium-energy x-ray scanners for primary inspection, as well as ASPs and medium-energy x-ray scanners again during secondary inspection.[25]

[25] Note that there are some differences across SFI pilots (CBP 2008).

### 7.1. Cost and Coverage

Our results show that the industry-centric regime should be able to provide better inspection coverage than the CSI at a lower unit cost. With (the current) one CSI inspection station per terminal, Terminal A can sustain an inspection rate of 5%, and Terminal B, 1.5%. Under the CSI, equipment and labor costs at both ports total on the order of $100 per inspected container. In contrast, the industry-centric inspection regime has the ability to sustain a 100% inspection load at a lower cost per container. With nine primary inspection stations, both terminals enjoyed few container delays at a cost of roughly $1–$2 per container, plus an additional $11–$13 for each container undergoing secondary inspection.

The above cost figures do not include, for either regime, the cost of real estate required for housing inspection equipment and for staging containers outside the inspection facility. The reason is that we do not have access to reliable estimates of cost of real estate in and around the port facilities at Terminals A and B. In both regimes, stable inspection loads lead to stable requirements for staging area, though at different locations in the terminal. Under the CSI, there is queue buildup at the inspection facility, whereas under the industry-centric regime there is a requirement for staging containers outside the primary inspection facilities, at the terminal gates and quay.

We note that simply scaling up the CSI regime would not significantly affect its unit costs. For example, a 10% inspection rate at Terminal A requires two inspection stations, with per-container costs on the order of $90. Similarly, neither the scanning of containers upon arrival nor the use of clever inspection-scheduling techniques (e.g., earliest-deadline-first instead of first-come, first-served service discipline) significantly change the CSI's capacity limits. Simulation results for these variants show that their performance is similar to that of the original CSI scheme.

The industry-centric regime's coverage and cost advantages follow from two sources. First, the use of higher-capacity drive-through equipment allows for 100% primary inspections. Second, the lower-cost, drive-through inspection station, together with the ability to amortize that cost over a much larger pool of containers, drives down the per-container costs.

In fact, a similar use of this same, lower-cost equipment could also help a CSI-style inspection regime to drive down costs in a similar fashion. Thus, one clear implication of our analysis is that the equipment and inspection protocol used in the industry-centric scheme should be useful for meeting the goal of 100% inspection of U.S.-bound containers at international ports.

### 7.2. Logistics

Even with more cost-effective inspection equipment, the CSI scheme has serious logistical drawbacks that would make it difficult to scale up. As we noted in §§3.2 and 5.1, the disruption arising from the large-scale pulling of containers from stacks, hours before their scheduled departure times, would be likely to degrade the efficiency of terminal operations at a significant cost. To provide a sense of the potential cost associated with delays, we note that the National Association of Manufacturers has recently voiced concern about the cost of the so-called "$10+2$" reporting requirements CBP has proposed to support the intelligence needed for the CSI. The NAM claims that the time required to gather the required information will add roughly three days to containers' stays at terminals, at an annual cost to trade of $5.67 billion per day of delay (Field 2009).[26]

Although these information-driven delays differ from the logistical delays associated with the untimely pulling of containers, the scale of the economic penalties would be similar. Therefore, even though our current data and models do not allow us to simulate the effect of widespread disruptions to terminals' load plans, these process-related concerns are significant enough that such a study should be undertaken before a CSI-style inspection scheme is broadly implemented.[27]

In contrast, the industry-centric regime's scan-upon-entry process eliminates the need to pull containers from terminal stacks. Furthermore, the only information it requires to support its assessment is a description of the cargo to be used in the secondary inspection process. It may therefore be more efficient to modify the CSI scheme to scan U.S.-bound containers upon arrival to the terminal, even if classification of containers by destination is difficult at that stage. Of course, the scanning of U.S.-bound containers upon arrival using drive-through equipment is a scheme that is essentially a restricted form of the industry-centric protocol, and one might ask what is the value of inspecting containers not bound for the United States.

In fact, we believe that there are positive externalities associated with the scanning of all containers, irrespective of destination. When identifying

[26] The "$10 + 2$" rule requires a container's shipper to provide 10 pieces of information concerning each U.S.-bound container's contents and history, and it requires the ocean carrier transporting the container to provide two pieces of information regarding the container's location in the vessel and movements after initial loading.

[27] To numerically evaluate the extent to which the pulling of containers would disrupt terminal cranes' load sequences, one needs to have access to the computer models and algorithms that terminals use to determine them.

potentially dangerous containers, the images provided by such a protocol provide an effective complement to field intelligence and to CBP's current risk-scoring capabilities. Together, systematic NII and intelligence would increase the robustness of the overall maritime security system, creating a safer global standard for maritime commerce, with the collateral benefit of helping to curb the proliferation of nuclear weapons. The latter benefit can be achieved only by scanning all containers. In the related context of commercial baggage screening for airline security, recent research using game-theoretic analysis similarly concludes that all baggage should be screened, regardless of destination (Kunreuther and Heal 2003, 2005).

Furthermore, the systematic use of NII also has economic advantages for the United States. Rather than penalizing only U.S.-bound traffic, inspection costs could naturally be allocated across all containers. An all-encompassing security regime also has the potential to induce terminal operators to bear the costs associated with installing and operating the inspection equipment. (For supportive statements from various terminal operators, see Appendix B in CBP (2008).) From the terminal operators' perspective, this expenditure would be part of the investment required for business-continuity purposes, given the threat from maritime terrorism and the associated systemwide disruption that a major terrorist incident could cause. Given the possibility that these economic benefits do, in fact, induce terminal operators to privately fund such a system, we believe that "industry-centric" is an appropriate name for this inspection regime.

### 7.3. Security Effectiveness

In comparing the industry-centric regime to the CSI regime, it is essential that we address potential differences in security effectiveness across the two protocols. The best way to address this concern would be to report data on false-negative and false-positive rates associated with the different inspection equipment. However, these data are classified and will probably remain so in the foreseeable future.

Our approach to the choice of equipment and the corresponding false-positive rates is to rely on the knowledge and field experience of Dr. Massey. His role as manager responsible for supporting the Second Line of Defense Program and the Megaports Initiative means that he has first-hand knowledge about the performance of the various inspection equipment discussed in this paper.

Even more important are the protocols' impact on false-negative rates. Although we are hard pressed for data on false-negative rates, we nevertheless are in the fortunate position of having to compare the industry-centric protocol with a benchmark, the CSI regime (see CBP 2008, p. 17). This approach of focusing on the relative performance of the two inspection regimes, rather than absolute performance, is consistent with that outlined in Wein et al. (2007, p. 223).

More specifically, the equipment used in CSI inspections include RIIDs and high-energy x-ray scanners, whereas the industry-centric protocol uses superior equipment for passive radiation detection, RPMs and an ASP. Although the industry-centric scheme uses primarily medium-energy x-ray scanners for imaging dense objects, a potential weakness when compared to the CSI's use of high-energy radiography, the industry-centric scheme provides for (and includes the cost of) a backup high-energy x-ray scanner for the exceptional cases that cannot be resolved by a second, controlled medium-energy scan. Thus, the false-negative rate using the industry-centric protocol should be at least as good as that under the CSI.[28]

Two additional points are worth bearing in mind. First, the fact that the CSI does not scale well—and hence can provide inspection coverage to only a small fraction of U.S.-bound containers—heightens the potential security advantage enjoyed by the 100% inspection protocol of the industry-centric scheme. Furthermore, the CSI's limited ability to scale up implies that it would leave maritime trade vulnerable to a major disruption in the aftermath of a potential nuclear incident, when 100% scanning would most likely be enforced, independent of whether or not the recommendations of the 9/11 Commission Act continue to be legally binding.

Second, the scans produced in the industry-centric regime need not be substitutes for the CSI's use of manifest information. Rather, they are useful complements that can help to resolve alarms that may be raised by ATS, as well as to identify potentially dangerous containers that the current system may have missed.

### 7.4. Remaining Challenges

Our simulation results and cost estimates suggest that the industry-centric protocol provides a realistic path to 100% inspection. Nevertheless, a number of significant challenges remain.

If terminal operators are to take responsibility for purchasing, deploying, and operating inspection equipment, then the U.S. government must work with them to establish appropriate technical standards for inspection processes and equipment. Moreover, the government must also ensure the effective operation of inspection equipment by terminal operators. One means of achieving this end would be to set up a

[28] What we are not able to achieve with this analysis is to capture quantitatively the absolute security-related benefit associated with a particular inspection regime. Therefore, we do not try to explore whether less than 100% scanning may be optimal, but treat it as a given constraint.

third-party audit process for the various inspection facilities, along with a centrally controlled audit of the auditors themselves. For details on this approach, we refer to Kunreuther et al. (2002).

There remain information technology (IT) challenges pertaining to the automated alarm generation, as noted in §6.1. There is also a need to deploy IT infrastructure to communicate container-scan and manifest information in real time to CBP officials in the United States (and to other customs officials who want this information for containers being exported to their jurisdictions). This would provide CBP with a window into and control over the automated detection process, as well as the ability to conduct a remote, independent assessment of cargo determined to be at higher risk. It would also provide a valuable audit trail if something untoward were to happen at a U.S. or overseas port. Concerns remain with regard to feasibility, sizing, and costing of such an IT setup. Although we do not address these questions as part of this research, they warrant further investigation.

There are also remain logistical hurdles. Substantial numbers of quayside primary inspections have not yet been undertaken at international ports. Therefore, there are likely to be hiccups in extending the 100% scanning regime to containers that arrive to the port on a vessel/barge. As with the SFI pilots, new pilot projects need to be undertaken to test out the logistics and equipment required for inspecting containers at the quay. From the perspective of terminal operators, Fairnie (2008) strikes a very optimistic note in this regard.

In rolling out any inspection regime at international ports, the United States must also be mindful of the diplomatic and legal challenges involved. In particular, in response to the 100% scanning requirement, there is a possibility that other countries may require reciprocal scanning of U.S. exports (CBP 2008). The United States must have a strategy in place to address this eventuality. Although such reciprocal requirements may increase the cost to the United States, our analysis indicates that this may not be an insurmountable problem.

Finally, we note that, although we have conducted our analysis at two of the world's busiest terminals, one terminal can be quite different from another with respect to layout and cost of operations. It would be hard to imagine a "one-size-fits-all" approach could work for all international terminals.

Despite of the above concerns, our work—which to the best of our knowledge is the first empirically driven, quantitative analysis of this problem—provides insights into the fundamental challenges and potential solutions that the United States will encounter in its efforts to bolster international container security.

### Acknowledgments
The authors gratefully acknowledge the financial support for this research provided by the Wharton Risk Management and Decision Processes Center. The authors also express their gratitude to Dr. Charles Massey for his input and encouragement through the course of this research, and to the management of the terminals from whom they collected data on container movements and who helped refine their understanding of terminal operations.

### Appendix. Cost Calculations
Recall that if $P$ is the principal for which an annuity payment is to be calculated, $n$ is the number of periods over which annuity is to be paid, and $i$ is the interest rate per period, then the annuity payment is $a(i, n, P) = P \div ((1 - 1/((1+i)^n))/i)$.

Here, the number of periods, $n$, represents the number of containers that can be handled over the useful life of the initial investment, $P$. The interest rate, $i$, is the effective rate that accrues between the processing of two containers. By using a constant rate, $i$, we implicitly assume that the time between the instances at which containers are processed is (fairly) constant. We describe how we determine $n$ and $i$.

Let $m$ be the lifetime of the investment $P$, in years, and for concreteness suppose that the investment $P$ is the price of a piece of equipment. If the average time required for the equipment to handle a container is $t$ minutes, then the number of containers that it can handle over its lifetime is $n = 60 \times 24 \times 365 \times m/t$.

For labor costs, we let $P$ be the annual salary and benefits paid to a team of people required to staff a task 24 hours a day, seven days a week, for a year, and we let the duration of the "investment" in these labor costs be $m = 1$ year.

Let $r$ denote the annual discount rate associated with an the investment of $P$, then, the effective interest rate per container is $i = (1+r)^{1/m} - 1$.

### References
Abt, C. C. 2003. The economic impact of nuclear terrorist attacks on freight transport systems in an age of seaport vulnerability. Report, Abt Associates, Cambridge, MA.

Bakshi, N., N. Gans. 2010. Securing the containerized supply chain: Analysis of government incentives for private investment. *Management Sci.* **56**(2) 219–233.

Bennett, A. C., Y. Z. Chin. 2008. 100% container scanning: Security policy implications for global supply chains. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Boske, L. B. 2006. Port and supply-chain security initiatives in the United States and abroad. Report, Lyndon B. Johnson School of Public Affairs, University of Texas, Austin.

Brown, L., L. Zhao. 2002. A test for the Poisson distribution. *Sankhya, Series A* **64**(3) 611–625.

Brown, L., N. Gans, A. Mandelbaum, A. Sakov, H. Shen, S. Zeltyn, L. Zhao. 2005. Statistical analysis of a telephone call center: A queueing-science perspective. *J. Amer. Statist. Assoc.* **100**(469) 36–50.

Fairnie, D. 2008. DP World, director of security, speech at World Customs Organization forum: What future for 100% scanning? Brussels, June 11.

Field, A. 2009. Manufacturers say 10+2 costs \$20 billion. *J. Commerce Online–News Story* (June 19), http://www.joc.com/government-regulation/manufacturers-say-102-costs-20-billion.

Flynn S. E. 2004. *America the Vulnerable: How Our Government Is Failing to Protect Us from Terrorism.* HarperCollins, New York.

Flynn, S. E. 2007. *The Edge of Disaster: Rebuilding a Resilient Nation*. Random House, New York.

Flynn, S. E. 2008. Overcoming the flaws in the U.S. government efforts to improve container, cargo, and supply chain security. Testimony before the Homeland Security Appropriations Subcommittee. U.S. House of Representatives, April 2, 2008.

Green, L., P. Kolesar, W. Whitt. 2007. Coping with time-varying demand when setting staffing requirements for a service system. *Production Oper. Management* **16**(1) 13–39.

Harrald, J. R., H. W. Stephens, J. R. van Dorp 2004. A framework for sustainable port security. *J. Homeland Security Emergency Management* **1**(2) Article 12, http://www.bepress.com/jhsem/vol1/iss2/12/.

Kunreuther, H., G. Heal. 2003. Interdependent security. *J. Risk Uncertainty* **26**(2–3) 231–249.

Kunreuther, H., G. Heal. 2005. IDS models of airline security. *J. Conflict Resolution* **49**(2) 201–217.

Kunreuther, H. C., P. J. McNulty, Y. Kang. 2002. Third-party inspection as an alternative command and control regulation. *Risk Anal.* **22**(2) 309–318.

Law, A. M., W. D. Kelton. 2007. *Simulation Modeling and Analysis*, 4th ed. McGraw-Hill, Boston.

Marine Link. 2004. The 5% myth. Accessed August 16, 2010, http://www.marinelink.com/news/article/the-5-myth/317423.aspx.

Martonosi, S. E., D. S. Ortiz, H. H. Willis. 2006. Evaluating the viability of 100 percent container inspection at America's ports. H. W. Richardson, P. Gordon, J. E. Moore II, eds. *The Economic Impact of Terrorist Attacks*. Edward Edgar Publishing, Northampton, MA, 218–241.

McClure, G. 2007. How safe are our ports? Accessed August 16, 2010, http://www.todaysengineer.org/2007/Sep/port-security.asp.

Policy Research Corporation. 2009. The impact of 100% scanning of U.S.-bound containers on maritime transport. Final report, Commissioned by the European Commission, Directorate-General Energy and Transport. Accessed August 16, 2010, http://ec.europa.eu/transport/maritime/studies/doc/2009_04_scanning_containers.pdf.

Steenken, D., S. Vo$\beta$, R. Stahlbock. 2004. Container terminal operation and operations research—A classification and literature review. *OR Spectrum* **26**(1) 3–49.

U.S. Congress. 2007. H.R.1: Implementing recommendations of the 9/11 commission act of 2007. Pub. L. No. 110-53, 110th Congress. Accessed October 27, 2010, http://intelligence.senate.gov/laws/pl11053.pdf.

U.S Customs and Border Protection (CBP). 2004. FAQ on 24-hour advance vessel manifest rule. Accessed August 16, 2010, http://www.customs.gov/ImageCache/cgov/content/import/carriers/24hour_5frule/24hour_5ffaq_2edoc/v1/24hour_5ffaq.doc. Last accessed on 08/16/10.

U.S Customs and Border Protection (CBP). 2008. Report to congress on integrated scanning system pilots (Security and Accountability for Every Port Act of 2006, Section 231). U.S. Customs and Border Protection, Washington, DC.

U.S. Government Accountability Office. 2005. Container security: A flexible staffing model and minimum equipment requirements would improve overseas targetting and inspection efforts. Report GAO-05-557, U.S. Government Accountability Office, Washington, DC.

U.S. Government Accountability Office. 2008a. Supply chain security: Examination of high-risk Cargo at foreign seaports have increased, but improved data collection and performance measures are needed. Report GAO-08-187, U.S. Government Accountability Office, Washington, DC.

U.S. Government Accountability Office. 2008b. Supply chain security: Challenges to scanning 100 percent of U.S.-bound Cargo containers. Report GAO-08-533T, U.S. Government Accountability Office, Washington DC.

U.S. Government Accountability Office. 2008c. CBP works with international entities to promote global customs security standards and initiatives, but challenges remain. Report GAO-08-538, U.S. Government Accountability Office, Washington, DC.

Wein, L. M., Y. Liu, Z. Cao, S. E. Flynn. 2007. The optimal spatiotemporal deployment of radiation portal monitors can improve nuclear detection at overseas ports. *Sci. Global Security* **15** 211–233.

Wein, L. M., A. H. Wilkins, M. Baveja, S. E. Flynn. 2006. Preventing the importation of illicit nuclear materials in shipping containers. *Risk Anal.* **26**(5) 1377–1393.

Willis, H. H., D. S. Ortiz. 2004. Evaluating the security of the global containerized supply chain. Report TR-214-RC, RAND Corporation, Santa Monica, CA.

**Corrected Version of Record**

In this version of "Estimating the Operational Impact of Container Inspections at International Ports," by Nitin Bakshi, Stephen E. Flynn, and Noah Gans, first published online in Articles in Advance November 19, 2010, Table 2 was corrected to show that steps 5, 6, and 7 together take 60 minutes to complete and steps 9 and 10 together take 40 minutes to complete, and Table 4 was corrected to show that steps 2 and 3 together take 25 seconds to complete.