1-2011

Estimating the Operational Impact of Container Inspections at International Ports

Nitin Bakshi

Stephen E. Flynn

Noah F. Gans
University of Pennsylvania

Recommended Citation

Bakshi, N., Flynn, S. E., & Gans, N. F. (2011). Estimating the Operational Impact of Container Inspections at International Ports. *Management Science, 57* (1), 1-20. http://dx.doi.org/10.1287/mnsc.1100.1252

Estimating the Operational Impact of Container Inspections at International Ports

Abstract
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.

Keywords
homeland security, container inspections, queueing simulation

Disciplines
International Business | Science and Technology Studies

Estimating the Operational Impact of Container Inspections at International Ports

Nitin Bakshi
London Business School
Regent's Park, London
nbakshi@london.edu

Stephen E. Flynn
Council on Foreign Relations
New York
sflynn@cfr.org

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

Working Paper # 2009-05-01

Risk Management and Decision Processes Center
The Wharton School, University of Pennsylvania
3730 Walnut Street, Jon Huntsman Hall, Suite 500
Philadelphia, PA, 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
http://opim.wharton.upenn.edu/risk/

THE WHARTON RISK MANAGEMENT AND DECISION PROCESSES CENTER

Established in 1984, the Wharton Risk Management and Decision Processes Center develops and promotes effective corporate and public policies for low-probability events with potentially catastrophic consequences through the integration of risk assessment, and risk perception with risk management strategies. Natural disasters, technological hazards, and national and international security issues (e.g., terrorism risk insurance markets, protection of critical infrastructure, global security) are among the extreme events that are the focus of the Center's research.

The Risk Center's neutrality allows it to undertake large-scale projects in conjunction with other researchers and organizations in the public and private sectors. Building on the disciplines of economics, decision sciences, finance, insurance, marketing and psychology, the Center supports and undertakes field and experimental studies of risk and uncertainty to better understand how individuals and organizations make choices under conditions of risk and uncertainty. Risk Center research also investigates the effectiveness of strategies such as risk communication, information sharing, incentive systems, insurance, regulation and public-private collaborations at a national and international scale. From these findings, the Wharton Risk Center's research team – over 50 faculty, fellows and doctoral students – is able to design new approaches to enable individuals and organizations to make better decisions regarding risk under various regulatory and market conditions.

The Center is also concerned with training leading decision makers. It actively engages multiple viewpoints, including top-level representatives from industry, government, international organizations, interest groups and academics through its research and policy publications, and through sponsored seminars, roundtables and forums.

More information is available at http://opim.wharton.upenn.edu/risk.

Estimating the Operational Impact of Container Inspections at International Ports

Nitin Bakshi
London Business School
Regent's Park, London, UK
nbakshi@london.edu

Stephen E. Flynn
Council on Foreign Relations
New York
sflynn@cfr.org

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

Abstract

A US law mandating non-intrusive imaging and radiation detection by 2012 for 100% of US-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 various inspection policies being considered. We find that the current inspection regime being advanced by the US Department of Homeland Security and widely supported by the international community can only handle a small percentage of the total load. An alternate inspection protocol which 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.

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). Unlike a long range missile, the millions of shipping containers that are used to transport goods in ocean-going vessels provide terrorists with a way to hide a nuclear device destined for US shores.[1] Further, by using a container, terrorists can potentially achieve mass disruption to global supply chains by creating widespread public anxiety that other containers may have nuclear devices. The resultant stepped-up inspections 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 Security Initiatives in Place

To counter the threat of nuclear terrorism, the US has initiated various security measures, both at domestic and foreign ports. These measures can require the co-operation of foreign nations, trading companies, terminal operators, customs brokers, trucking companies, ocean carriers, and

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

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).

CSI is a security program administered by the US Customs and Border Protection (CBP), an agency which falls within the Department of Homeland Security (DHS). The program, announced in January 2002, uses rules-based software to identify containers bound for the US 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 US, forward manifest information to CSI officials at least 24 hours prior to the container's lading on the vessel that will call on a US port. Once 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 or origin, before they are shipped to US ports. These customs officials use high-energy x-ray radiography (a form of non-intrusive inspection) and hand-held, 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 (RDD).

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

Launched in 2007, SFI is a joint initiative of CBP, the US Department of Energy (DOE), and the US Department of State (CBP 2008-b). It is meant to leverage the learning from other port security initiatives, such as Operation Safe Commerce, and serve as a pilot for satisfying the 100% scanning requirement. Under SFI all (not just US-bound) containers arriving at participating overseas seaports are scanned with both non-intrusive radiographic imaging and passive radiation detection equipment placed at terminal entrance gates.

Optical scanning technology is used to identify containers and classify them by destination. Sensor and image data gathered through this "primary" inspection is then transmitted in near real time to the National Targeting Center in Virginia (CBP 2008-b). 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 as high-risk. High-risk containers then undergo a more sensitive "secondary inspection."

The SFI program has been piloted for one year on a full scale in 3 small international ports, and it has been implemented on a limited capacity basis in 4 additional ports. The 3 full-scale pilots were conducted in Karachi (Pakistan), Puerto Cortes (Honduras) and Southampton (UK). SFI is operating on a limited basis in Hong Kong (GAO 2008-b).

To summarize, the CSI and SFI protocols differ in the pool of containers targeted for inspection, as well as both the timing and tools used for preliminary inspection. The CSI inspection process is geared exclusively to US-bound containers, it begins 24 hours in advance of a container's lading onto an oceangoing vessel, and it uses information contained in the shipping manifests to decide whether or not specific containers require intensive non-intrusive inspection (NII). The SFI protocol uses drive-through portals to scan every container as it enters a port terminal, and the results of these scans in addition to shipping manifest data can trigger the need for more intensive NII. Alternatively, the NII data collected under the SFI protocol may lead CBP to rule out the need to inspect a container that its targeting algorithm identified as warranting an intensive inspection.[2]

1.2 One-Hundred Percent Scanning Requirement

The immediate need for this study arises from a US law enacted in August 2007, "Implementing Recommendations of the 9/11 Commission Act of 2007" (Pub. L.110-53), popularly called the 9/11 Commission Act. The law requires that, before *any* cargo bound for the United States is loaded on a ship at an international port, it must be scanned using 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 2-year increments (US Congress 2007).

This law is a significant deviation from CBP's CSI approach of scanning only cargo it identifies to be high-risk, and immediately generated questions by CBP and European customs officials, by trade associations like the US Chamber of Commerce and the National Association of Manufacturers, and by corporate leaders regarding the operational feasibility of 100% NII scanning. The most common concern that was raised was that the congestion resulting from this security requirement will substantially increase the cost of doing business and hurt commerce (NAM 2008). There are

[2] According to Fairnie (2008), "A key reason DP World has participated in international customs initiatives is to gain trade facilitation and resiliency benefits for ourselves and our customers. In that regard consider this: we understand that the scanning trial led to a plunge in physical inspections of containers from Southampton at US ports - from 1000 containers down to 7 in equivalent time periods."

essentially three broad ways in which the 100-percent 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 port 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 container 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. Again, decreases in terminal efficiency, along with increased lead times, would lead to higher consumer costs.

1.3 Evaluating the Impact of the 100% Inspection Requirement

If not carefully considered, efforts to satisfy the requirement to scan 100% of US-bound containers have the potential to significantly degrade the efficiency of port 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 port 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 datasets 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 US 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 port 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 overlay simulated inspection processes on top of this historical record.

To estimate the effect of container 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 US. If the simulated completion time falls beyond the actual loading time, then we mark the container as being delayed, and if it 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 port 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 various inspection schemes we consider.

1.4 Results and Implications

Our simulation results suggest that a variant of the SFI inspection scheme, that 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 US law.[3] Its use of rapid screening by relatively low-cost drive-though portals can handle 100% of *all* container traffic – that are bound for the US, 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 US-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 US-bound traffic at the smaller port and no more than 1.5% of US-bound traffic

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

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 required for 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. While 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. Our simulation results clearly imply that the equipment and inspection protocol used in the Industry-centric scheme should become the standard by which containers are inspected at 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 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, for loading. This disrupts the highly optimized sequence of containers that terminals use to order yard cranes' movements within the stacks. Under the Industry-centric scheme, in contrast, targeted containers undergo secondary 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 routine early removals of even a small number 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 Section 3 we describe the steps in the container flow through a terminal: when there are no inspections; under the CSI regime; and also under the Industry-centric regime. In Section 4 we describe our research methodology, including a description of the dataset 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 it implementation. Section 6 describes the simulations for the Industry-centric regime. Finally, in Section 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) provides a risk-management framework for securing maritime infrastructure. Boske (2006) reviews the major US-domestic and international initiatives in this regard.

Two books which address this topic in depth are Flynn (2004) and Flynn (2007). The former highlights the US's overall vulnerability to maritime terrorism, while the latter emphasizes the inadequacy of its present set of security initiatives.

In this paper we specifically look at the tradeoff 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 US-domestic ports can be found in Bakshi and Gans (2007). 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) which also aims 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 the paper cannot model the dynamics of terminal operations and is compelled to decouple the analyses of the primary and secondary inspection process. Thus, while 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.

CSI and SFI are two of the key US-maritime security initiatives in place today that address the problem of container security at international ports. CBP (2008-a) is a repository of documents pertaining to CSI, while CBP (2008-b) provides a compilation of documents relevant to SFI. The Government Accountability Office (GAO) periodically reviews these and other maritime security programs of the US government. Examples of such reports, focusing on container inspections at international ports, include GAO (2005) and GAO (2008-a). The latest in this series are GAO (2008-b) and GAO (2008-c), which highlight the difficulties involved with rolling out a 100% inspection regime. In particular, GAO (2008-c) emphasizes the challenges associated with adapting the risk management approach of allocating resources towards the thorough inspection of only high-risk containers, as embodied by the CSI regime, to the new paradigm of inspecting every single container.

We have estimated statistics associated with process times and process outcomes in collaboration with Dr. Charles Massey, as well as the managers of Terminal A and of Terminal B.[4] All statistics regarding the inspection steps (box 11 in Figure 4, and boxes 2 and 5 in Figure 7) were developed in collaboration Dr. Massey. For other, logistical process steps, we developed estimates for process times in consultation with the terminals. The terminals also provided estimates for the costs and lifetimes of the equipment involved in these steps. Personnel costs for the logistical steps were estimated in collaboration with Dr. Massey.[5]

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 5 or 6 high. Empty containers may be stacked even higher.

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

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

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, due to time constraints, exit the terminal as soon they arrive. Figure 1 depicts a high level schematic of the process.

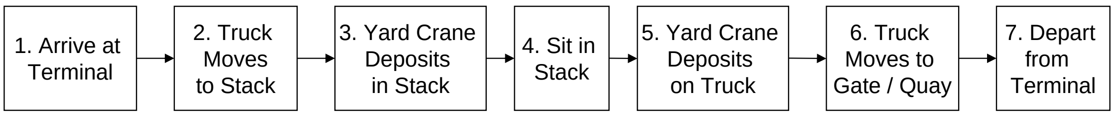

Figure 1: Baseline Container Flow.

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.

To minimize supply-chain costs, all containers should ideally be hot boxes, arriving to a port terminal just before their scheduled departures. This tight turnaround would reduce supply-chain lead times and eliminate the extra space and handling costs associated with moving containers into and out of a yard stack. Nevertheless, most containers arrive well in advance of their scheduled departure times to allow for a full and rapid load-out of scheduled voyage of an ocean carrier capable of carrying upwards of 5,000 containers. Pre-arrival and storage in the terminal yard stacks also provides a buffer that insulates the terminal from uncertainties regarding the time at which ocean-going vessels and, to a lesser extent, trucks and barges may arrive at the terminal.

Departures are similar to arrivals. A truck that carries a container from the terminal to a destination within the hinterland shows up at the terminal gates and is directed to the relevant position among the container stacks. A yard crane picks the container from the stack and places it on the truck's empty chassis, and the truck proceeds to the yard's exit gate, where it checks out. For departures by barge or vessel, an internal truck is directed to drive up to the stack, instead. Again, a yard crane picks the container from its spot in the stack and places it on the truck. Then the internal truck drives to the quayside where it is loaded by a quay crane onto a waiting barge or ocean-going vessel.

3.2 Container Positioning and Terminal Efficiency

At a high level, the typical movement of a container through a port 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. Furthermore, these decisions can be quite complex, and large terminals use sophisticated computer algorithms to help them decide at which spot in the stack each container should be placed (Steenken et al. 2004).

For the purposes of our analysis, two interrelated decisions regarding container movements are worth noting. The first concerns the impact of a container's vertical position in the stack, and the second regards the effect of the positions of groups of containers on the loading of ships, especially ocean-going vessels.

First, suppose a container sitting in the stack is scheduled to depart, a truck (either internal or external) pulling up to the stack waiting for the yard crane to deposit the container on its chassis. If the targeted container sits on top of the stack, then the yard crane simply moves to the container's position, locks onto it, and then moves the box over to the side of the stack, placing the container on the waiting chassis. If, however, other boxes sit on top of the container of interest, the crane must first move each of them away, to other (perhaps temporary) positions in the stack, and then pull the target container as planned. After the container is pulled from the stack, those other container may be moved back to their original spots.

Therefore, a container that is well positioned, at the top of the stack, requires only one crane move to be placed on the truck, while a container lying beneath one or two other containers may require up to three or five crane moves. With each additional move, the terminal uses labor and equipment that could be used to move other containers and incurs additional cost, as well as (perhaps) the capacity of the truck and driver, if they sit and wait for the container.

Second, when a terminal loads a group of containers onto an ocean-going vessel, each container is stowed in a pre-determined position in one of the ship's holds or on its deck. The map of containers' positions within a vessel is called its load plan, and the determinants of a vessel's load plan are varied and complex: containers that are destined for the same port of disembarkation should be in the same hold; heavier containers should be stored beneath lighter ones; weight must balance fore and aft, as well as port to starboard; refrigerated containers must occupy slots supplied with electrical power; and flammable and other hazardous cargo are placed topside, away from the crew's quarters.

A vessel's load plan helps to determine the order in which containers must be pulled from a terminal's stacks. Those that are grouped together on the 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. Since containers that sit beneath others in a ship's hold will be loaded earlier, they will be pulled from the terminal first, and therefore should not sit beneath other containers in the stack. Otherwise, the terminal will require extra crane moves to displace (and replace) the containers that sit atop.

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. This drives up the terminal's costs, and it affects 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.

4 Research Methodology

To assess the operational impact of the various 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 time. 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 dataset as they are relevant in certain scenarios that we analyze. While we store the early arrivals "as is" in the dataset, we assign the records with departure times after the end of the month a proxy departure time which 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 US-bound containers from those that are not. In our database, roughly 13% of the traffic at Terminal A is US-bound, while 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*.[6]

Arrival rates measure the amount of work flowing into the system over time, work that drives the need for terminal, as well as potential inspection, capacity. Figure 2 shows that these rates vary significantly over time at the two terminals. While the long-run average arrival rates at the terminals are 500 to 600 per hour, differences in rates can be up to 5-fold from one hour to the next. This rapid variation in arrival rates suggests that traditional queueing models – such as those used in Bennet and Chin (2008), Wein et al. (2006), and Wein et al. (2007) – which assume a constant arrival rate over a fairly sustained period of time, may not be adequate to capture the terminals' dynamics. Discrete-event simulation provides a more robust means of modeling system performance in these settings.

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. Figure 3 shows the distribution of

[6]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.

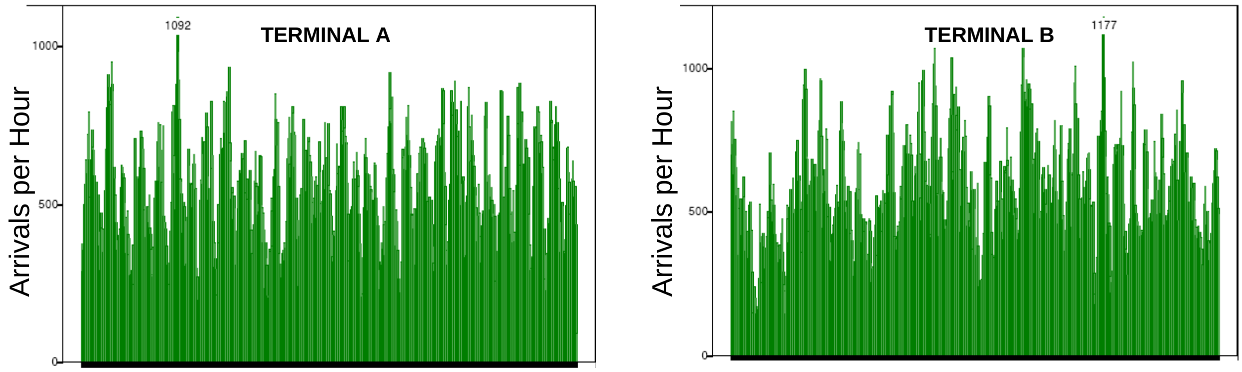

Figure 2: Container Arrival Rate (per hour) over a Month: Terminals A and B.

dwell times (in days) for containers at the two terminals. Each terminal's distribution is calculated over all departures that occurred during that month.[7] The spike close to zero reflects the 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.

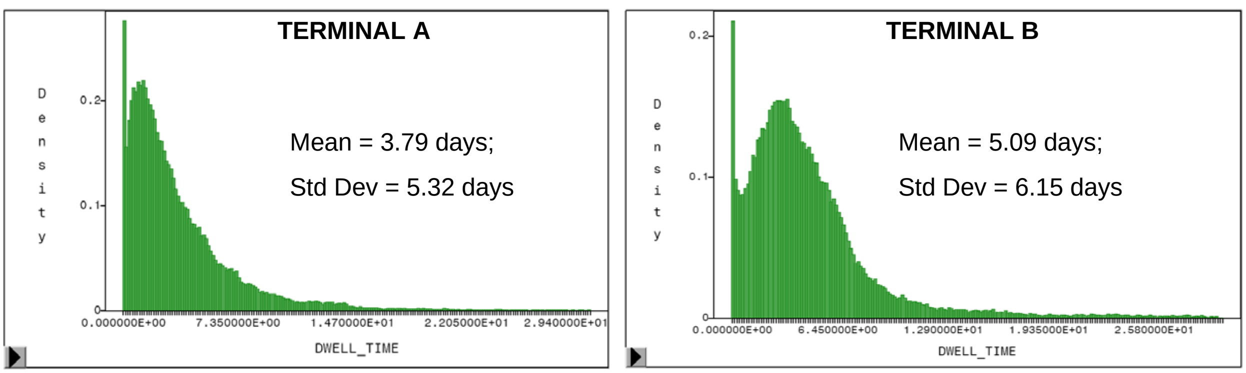

Figure 3: Dwell Time Distributions for Terminals A and B.

4.2 Simulation Models

To capture the potential impact of inspections at a container terminal, we use discrete-event simulation (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 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.

For concreteness, we provide an example of these four elements. An action might be the passive

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

monitoring of a container to determine if it emits radiation. A server that performs the action might be a radiation portal monitor (RPM). The time required to perform the scan may be a random variable (drawn from a predefined probability distribution) that reflects the time required for a truck to drive through the RPM. The container would flow into the RPM just after arriving to the terminal. After the RPM scan is complete, the container might proceed to an NII station that uses gamma or x-ray radiography.

Note that, because the inspection processes have not yet been implemented comprehensively, we do not have detailed historical data from which we can construct probability distributions for process times. Nevertheless, to operationalize our simulation models we must make specific assumptions regarding process-time distributions. For this purpose, we often use lognormal distributions in which the 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.[8] 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 should have no effect on 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.

Given a stream of time stamps for arriving containers, one simulation run will generate a series of inspection actions to be completed for each container, each action associated with a (perhaps) randomly drawn "service" time. As the containers make their way through the simulated inspection process, they will spend time being processed. If, at a given step of the inspection process, the number of arriving containers exceeds the service capacity of the equipment and people dedicated to performing that step, then containers may also spend time waiting in queue.

Indeed, the only way to avoid this queueing is to provide enough 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

[8] For example, the dwell times shown in Figure 3 are roughly lognormally distributed.

less 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 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. 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 Terminal A and B estimate that, for containers that are stacked 2-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 two metrics: one, the fraction of inspected containers whose simulated inspection time exceeds its historical departure time; and two, for each capacity-constrained step of the inspection process, statistics regarding queue length. Because queue lengths are likely to change dramatically over time – in response to highly variable arrival rates – we track both the average queue length and the maximum queue length incurred over the course of the month.

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 our averaging of statistics over multiple simulation "runs" as a method of stabilizing statistical fluctuations that may result from randomness in simulated service times or service outcomes. For background concerning both of these issues, see Law and Kelton (2007).

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 7 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 CSI's 24-hour rule, we exclude containers processed on the last day of each month. Since autumn months have 30 or 31 days, this leaves 22 or 23 days worth of data for the calculation of performance statistics.

Second, within our simulations, elements of the inspection process vary randomly from one container to the next. For example, the time required to transport a container from the stack to the inspection area may be 10 minutes for one container and 15 minutes for another. Similarly, the outcome of one inspection may be an alarm for one container, but that for another may be no alarm. As is common in discrete-event simulations, these types of outcomes are modeled as random variables, and each simulation "run" generates sample outcomes that drive the timing and disposition of the inspection process for all containers over the course of the month.

Due to this underlying uncertainty, the performance metrics associated with any single simulated month will vary randomly from one simulation run to another. To filter out the effects of these low-level random fluctuations we simulate each one-month period several times and calculate the average over these multiple runs. For example, in each simulation run for a given month, we calculate the maximum queue length at an inspection station. In turn, we calculate the average of these maxima, along with a 95% confidence interval for the average. When the relative uncertainty concerning our estimate of the statistic – as represented by the half-width of the confidence interval – falls below 2.5% of the estimated average, we allow ourselves to stop simulating.

More specifically, we bound the number of runs required for every simulation: below by 25 and above by 100. Otherwise, we stop simulating after the half-widths of the confidence intervals of all values of the statistics we track fall at or below 2.5% of their means. Note that, in the body of the paper we simply report our estimated averages and do not report numbers of simulation runs or widths of confidence intervals. We do report all of these data for the CSI simulations in Appendix A, however.[9]

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:

[9]For the Industry-centric simulations, the key metrics pertain to the primary inspection process which involves deterministic process times, and hence confidence intervals are not relevant.

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.[10]

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. 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: $a(i, n, P) = P \div \left(\frac{1 - \frac{1}{(1+i)^n}}{i} \right)$.

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, 7 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: $i = (1 + r)^{\frac{1}{m}} - 1$.

When evaluating costs accrued by the US government, we use a risk-free rate of $r = 3.7\%$. For costs borne by private companies, such as terminal operators, we use an estimated cost of capital for the maritime transportation sector of $r = 7.2\%$ per annum.

5 Model and Results for CSI Regime

The CSI inspection regime applies only to US-bound containers and begins twenty-four hours before a container is to be loaded onto the vessel that will carry it to the US 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. Thus, under CSI, there exists a 24-hour window within which a US-bound container has to be inspected,

[10] A tophandler is a small mobile crane.

if supply-chain lead times are to remain relatively unaffected by security operations. Note that the 24-hour rule implies that US-bound containers cannot be hot boxes.

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

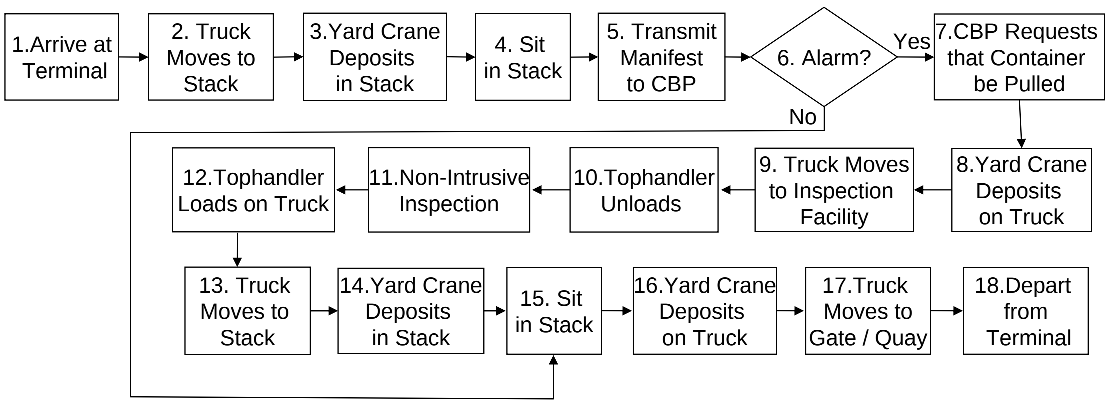

Figure 4: Typical Container Flow for CSI Regime.

Recall that one measure of system performance is the fraction of inspected containers that completes the inspection process before the time-stamp of its actual departure. (See Figure **??**.) In our simulation models, we define the total inspection time to be the elapsed time between the start of the transmission of the container's manifest (box 5) and the end of time required for the container to complete its non-intrusive inspection (box 11).[11]

The other 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 non-intrusive inspection step (box 11). At this step we track both the average and the maximum number of containers waiting to be inspected over the course of a simulated month.

5.1 Model Details

In this section we describe the details of our simulation model. We describe each process step, along with the critical statistics associated with our simulation calculations. We estimate statistics associated with inspection times and outcomes in collaboration with Dr. Charles Massey of Sandia National Laboratories. Statistics associated with other container movements within terminals are

[11] 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.

estimated with the input of the managers at Terminal A and at Terminal B. We summarize these statistics in Table 1.

Transmit Manifest to CBP (5) – Request Container be Pulled (7)

The CSI inspection process begins when an ocean carrier transmits a container's manifest to CBP. In practice, this occurs at least 24 hours before the associated vessel's estimated departure time. For the purposes of the simulation, we assume this occurs precisely 24 hours before the time stamp associated with the container's departure from the terminal, as recorded in our datasets.

Once the manifest is transmitted, CBP's National Targeting Center (located in Arlington, Virginia) compares the manifest to other intelligence it has collected and determines whether or not the container poses a potential threat. If so, the container is added to a list of containers on the vessel that are flagged for inspection by the loading port's customs authorities.

When complete, the National Targeting Center transmits the list of flagged containers to the CBP officials stationed at the port, who then communicate the request to the host nation's local customs authority. In turn, the customs authority transmits the inspection request to the terminal's management so that they locate the targeted containers and bring them to the government's inspection facility.

Probability of Inspection. The fraction of containers that are currently identified as high risk is roughly 5% (McClure 2007). In our interviews with terminal personnel, we found that the number of containers actually flagged for NII was far lower, however. For the purposes of this study – whose aim is to determine the level of inspections that the CSI regime can support – we systematically vary the fraction tagged. If the fraction tagged is f then, in the simulation, each laden US-bound container is tagged with independent and identically distributed (*i.i.d.*) probability f. Empty containers need not undergo NII.

Timing Assumption. While we do not have hard data from CBP concerning the timing of these requests, we have interviewed terminal personnel who have noted that it typically takes several hours, past the 24-hour mark, before a request that a container be pulled reaches terminal management. In the absence of hard data, we have assumed that these requests always take exactly one hour.

Yard Crane Deposits on Truck (8)

If a container is tagged as risky, the terminal operator must move it to an inspection facility where it undergoes inspection by the local customs authority. The first step in this process is the retrieval of the container from the terminal's stacks. An internal truck drives to the position in the stack where the container is resting, and a yard crane places the container on the truck.

The time required to complete this process depends on a number of factors. If both the truck and the yard crane are available and no other container sits atop the container to be removed, then the crane moves into place, latches on, and moves the container to the truck as it waits at the side of the stack, the process taking just a few minutes. Any of these three conditions may be violated, however. If there are containers sitting atop the desired container, then the yard crane must also handle them as described in §3.2, each additional move adding minutes to the operation. If an internal truck is not available, then the yard crane must wait, and this may also add time to the process. Finally, and most importantly, the yard crane itself may not be available for the request. As briefly described in §3.2, it is likely to be busy loading a predetermined sequence of containers into a vessel that departs several hours before the vessel onto which the container in question is to be loaded. The interruption of the loading process to handle this out-of-order container may not happen for some time – tens of minutes or even hours – particularly if there are many such targeted containers and the vessel being loaded is at risk of being delayed.

In the long run (were this inspection regime to be operationalized) terminals may be able to accommodate large numbers of inspections by building in slack (empty slots) into yard cranes' lists of containers to be moved. Nevertheless, we expect that delays would remain significant, and the empty slots required to improve responsiveness would systematically increase terminal handling costs.[12]

Timing Assumption. We assume that the time spent waiting for the yard crane and truck to become available, for the yard crane to pick a container from the stack and load it onto a truck is lognormally distributed with a mean of 15 minutes and a standard deviation of 15 minutes.

Truck Moves to Inspection Facility (9) – Tophandler Unloads (10)

Once loaded on to the truck, the container is transported to the inspection site. The inspection facility may be within the terminal, or it may be at an off-site location, depending on the terminal's layout. The transportation time depends the distance to the site, as well as the amount of traffic congestion between the stack and the site. Terminals A and B are large and busy, and travel times in them are on the order of tens of minutes.

Upon arrival at the inspection site, the container is unloaded by a tophandler. The container may be unloaded directly onto an inspection station, if it is free. Otherwise the container is placed somewhere in a staging area, in a queue of containers to be inspected. We assume that there exists

[12] If x% of slots are made empty, then yard-crane efficiency declines by x% and yard crane costs per container increase by 100/(100-x) percent.

adequate tophandler capacity at the inspection station, so this process step should not require more than a few minutes.

Timing Assumption. We assume that the total process time for both steps is 40 minutes (deterministic). Although there will be some uncertainty associated with this process time, it is likely to be small if we assume that terminals will arrange for adequate truck and tophandler capacity such that these steps are not affected by congestion. However, it is worth noting that constraints on capacity will likely introduce variability into the process time associated with this step.

Non-Intrusive Inspection (11)

At the inspection facility, each container undergoes two forms of inspection. The first is a scan using hand-held radiation isotope identification devices (RIID). RIIDs are used to detect gamma and neutron emissions that would indicate the presence of illicit nuclear materials. The second is radiographic imaging using high-energy (9 MeV) x-ray equipment. This form of NII is used to detect a nuclear device that has lead shielding around it as a means of containing its emission of dangerous isotopes, thereby allowing it to escape detection in RIID scans. Shielding material is dense and can be observed as a large shadow in a radiographic scan. The high-energy x-ray machines provide detailed images, but because of their energy levels they are hazardous to use, requiring the driver of a truck carrying the container to leave the cab and move away from the area.

In the unlikely event that NII does not rule out the possibility that a container houses a nuclear weapon or RDD, the host government, together with US government agencies, would begin a further response. For the purposes of our simulation study, which focuses on the types of capacity required for the higher-volume NII screenings, these further responses need not be considered.

Number of Inspection Stations. To date large CSI ports have operated a single inspection station that uses two RIIDs and one high-energy x-ray radiographic device. Our simulations include results for two setups, one with a single inspection station and another with two. In either case, if more than one or two containers arrive to the inspection area over a short period of time, some will wait in queue before they can be inspected. In the event there is such congestion, we assume containers are taken from the staging queue to an inspection station on a "First Come, First Served" (FCFS) basis.

Timing Assumption. We assume inspection times are lognormally distributed with a mean of 20 minutes and a standard deviation of 20 minutes. This process time includes: time for manual RIID scan; time for a truck driver to position the container within the high-energy radiographic equipment, exit the cab, and move away from the inspection station; time to scan the container;

time for customs officials to interpret the scan; and finally time for the driver to reenter the cab and drive the container away. Note that, if the container had queued for inspection, the simulated inspection time would also include time for a tophandler to load the container onto the internal truck that moves the container to the inspection station.

Tophandler Loads (12) – Yard Crane Deposits in Stack (14)

After an NII is completed, a tophandler places the inspected container back on an internal truck, and the truck drives the container to the terminal's stack, where a yard crane places it in (a possibly new) designated location. The time required for the tophandler and truck to do their work is the same as that required for the trip from the stack to the inspection area: on the order of tens of minutes. The time required for a yard crane to put the container in the stack can be highly variable and is affected by the same factors as other yard-crane operations.

Timing Assumption. We do not include the times for these process steps in our simulation study. More specifically, we assume that any container that completes its NII before its historical departure time – the time at which a quay crane would have locked onto it to place it in a vessel – would not miss its loading. Rather than being returned to the stack, a container that completes its NII shortly before its departure time would be moved by truck directly to the quay. Since this last step can be accomplished in a relatively short amount of time, we do not explicitly include it in the calculations to determine whether or not a container has missed its scheduled departure.

We do account for the equipment and labor costs associated with returning a container to a stack location, post inspection, however. For this purpose we assume that labor and equipment time required to complete the task is the same as that spent in steps 8–10, in which a yard crane retrieves a container from its location in the stack, an internal truck transports it to the inspection facility, and a tophandler unloads it.

5.2 Simulation Results

In this section we present our simulation results for the CSI inspection regime. For Terminal A we vary the fraction of US-bound containers that are inspected from 4% to 20%, and for Terminal B we vary the percentage from 1% to 20%. We use a step size of 1 while varying the inspection rate, with one notable exception. Because there is a steep jump in the fraction of delayed containers at Terminal B as the inspection rate climbs from 1% to 2%, we also simulate a 1.5% inspection rate. For each of these percentages, we run two simulations: one with one NII station and another with two. As noted in §4.3, in Appendix A we provide supporting tables that report the 95% confidence

Box	Process Step	Process Time Distribution	Mean	Standard Deviation
5	Transmit Manifest to CBP			
6	Alarm?	fixed	60 min	0 min
7	CBP Requests Container Be Pulled			
8	Yard Crane Deposits on Truck	lognormal	15 min	15 min
9	Truck Moves to Inspection Facility	fixed	40 min	0 min
10	Tophandler Unloads Truck			
11	NII Inspection	lognormal	20 min	20 min
12	Tophandler Loads Truck			
13	Truck Moves to Stack		omitted	
14	Yard Crane Deposits in Stack			

Table 1: Process-Time Assumptions Used in Simulations of CSI Protocol

intervals for all relevant statistics.

Figure 5 plots the fraction of inspected containers that miss their scheduled departure against the fraction of all (laden) US-bound containers that are inspected. The figure's top plot shows that, for 1 NII station at Terminal A, a 4% inspection rate generates nearly no delayed containers. With a 5% inspection rate – the rate at which CBP targets containers for inspection – 16.3% of tagged containers would have missed their historical vessel loading times. 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. 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.

The bottom plot of Figure 5 shows that, 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 13.7%. 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 US-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 US-bound containers is more than 138% higher at B, and each NII station can handle a proportionately lower fraction of the offered traffic.

In Figure 6 we plot the maximum 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,

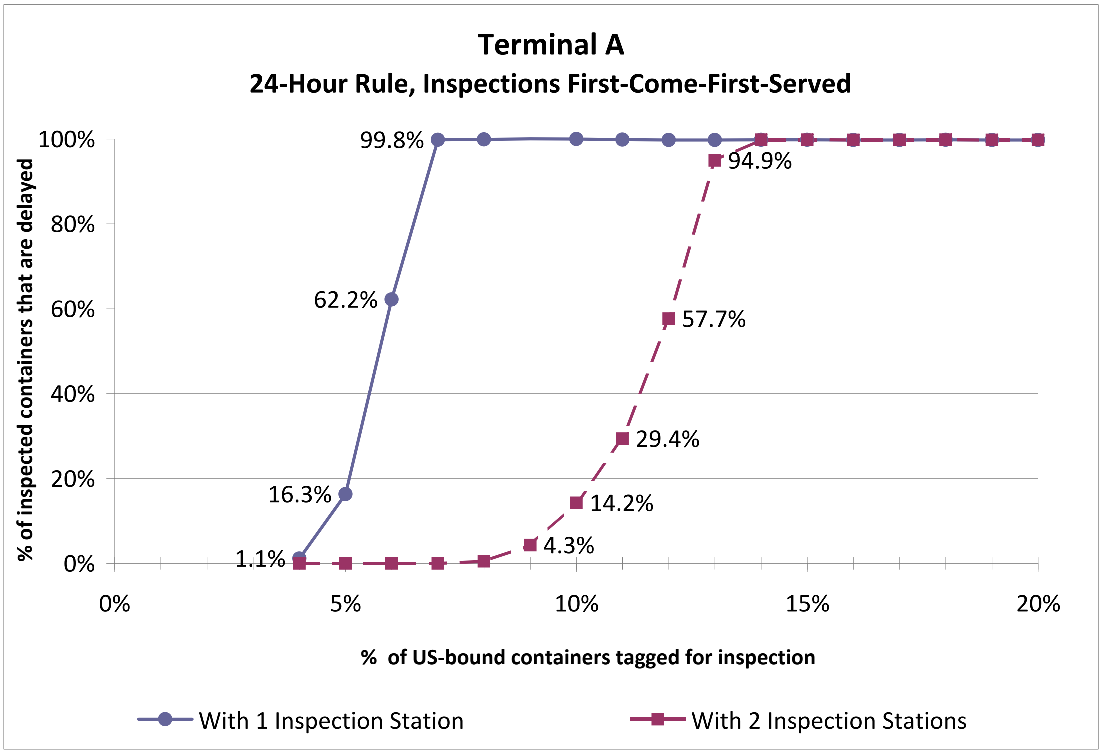

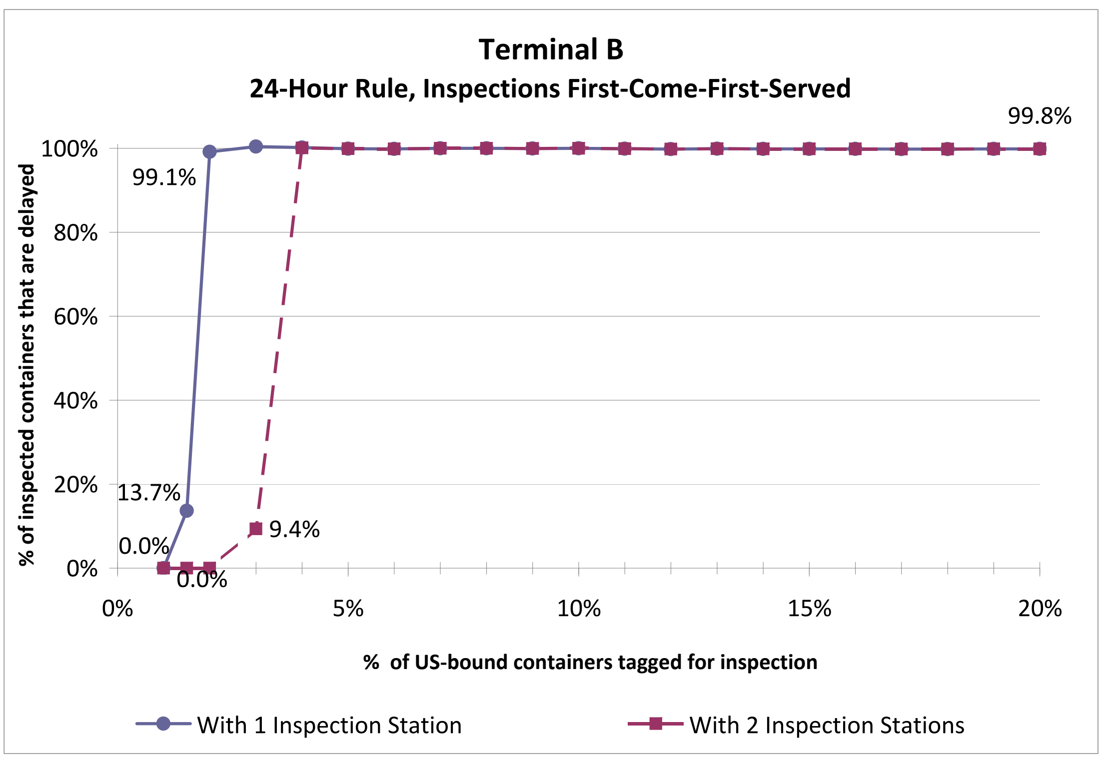

Figure 5: Fraction of Containers Delayed at Terminals A and B under CSI Regime

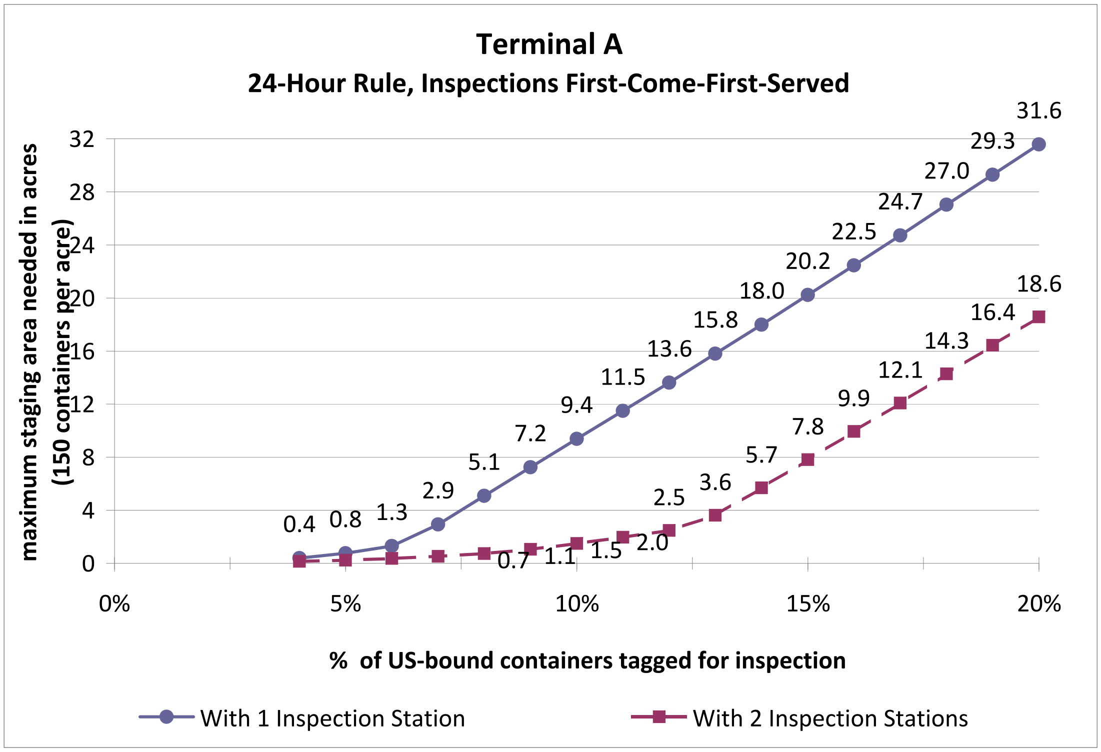

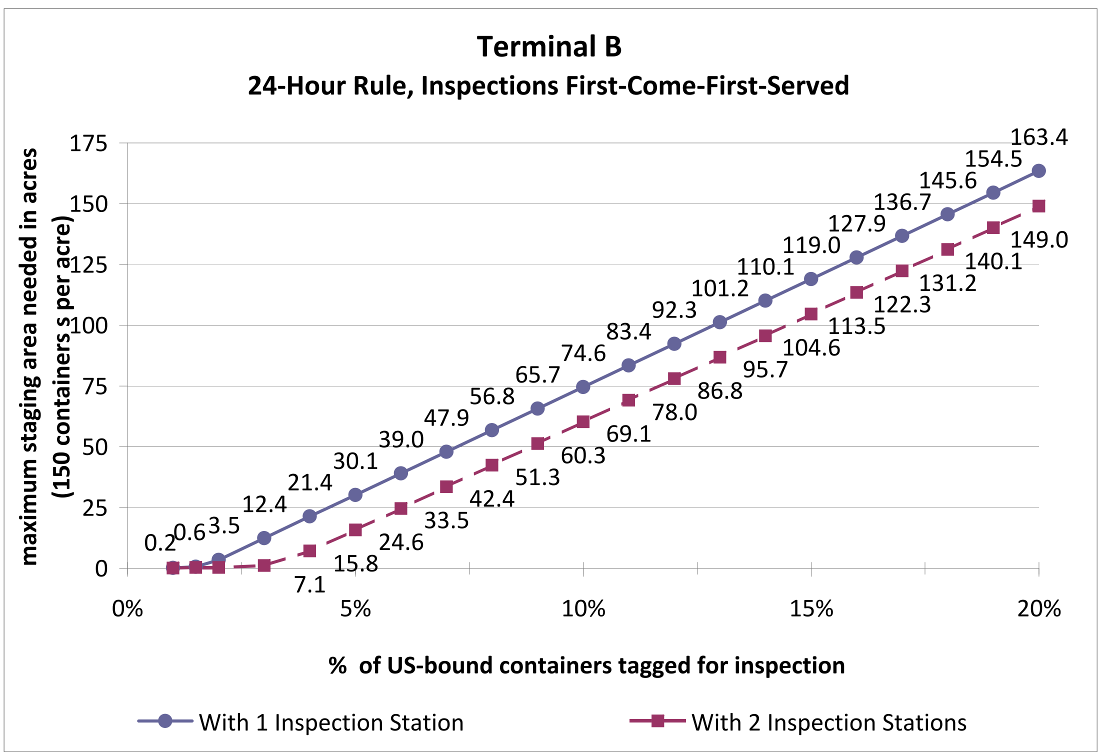

Figure 6: Maximum Staging Area Required at Terminals A and B under CSI Regime

the industry thumb rule for estimated average number of containers that can be accommodated per acre when container stacks are a maximum of 2-high.[13]

The top panel of Figure 6 shows that, for the simulation of Terminal A with one NII station, the required staging area grows linearly with the inspection rate when that rate crosses the 7% mark. Recall that 7% is the threshold rate at which NII utilization hits 100%, and every additional container that is flagged for inspection adds to the queue buildup in the staging area outside the inspection facility. In fact, in these cases the maximum of the queue length occurs at the end of the month, after many tagged but unprocessed containers have arrived to the inspection area.

We note that traditional (steady-state) queueing formulae would predict that, beyond the point of 100% utilization, the long-run average queue length would become infinite. In fact, the only the reason that the size of the staging area remains finite in our simulations is that we track container buildup over only 22 or 23 days. For a given inspection percentage, each additional 22–23 days with analogous container traffic would drive the required area up at roughly the same rate, and over the long run the queue would grow without bound.

Below the 7% threshold the inspection station appears to be stable, and the queue does not grow in the same manner. Rather, it fluctuates with the arrival rate and there is no long-term, systematic buildup. In this range, each added month of simulation would not increase the maximum backlog. We would expect it to remain stable: on the order of 0.4 to 3 acres, depending on the fraction of US-bound containers inspected.

The top plot of Figure 6 also shows that, when the number of NII stations is doubled, 13% is again the point at which the backlog of containers waiting for inspections explodes. Below this point, the queue is stable, and the maximum staging area ranges from 0 to roughly 3.6 acres. The bottom plot of Figure 6 shows that Terminal B's behavior is analogous to that of terminal A. Below the 100% utilization points of 2% and 4%, for one and two NII stations, the queue is stable. Above these thresholds the inspection queue explodes.

5.3 Cost Estimate

Under the CSI protocol, each US-bound container that is inspected passes through steps 5 through 14, and we can estimate the costs associated with each of these process steps. More specifically, we recall from Section 4.4 that, given equipment costs and lifetimes, as well as annual cost per FTE for labor costs, the essential elements required to allocate costs on a per-container basis are process

[13] 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.

times and numbers of pieces of equipment or FTE's. The first six columns of Table 2 provide this information.

		Average	Equipment		Labor		Cost Per Container	
	Process	Time	Unit Cost	Life		Comp.	Equip.	Labor
Box	Step	(min)	($000)	(yrs)	FTEs	($000)	($)	($)
5	Transmit Manifest							
6	Alarm?	n/a	n/a		n/a		n/a	
7	CBP Requests							
8	Yard Crane Deposits	15	1,200	25	1.2	50	1.04	2.18
9	Truck Moves	40	20	10	6.42	15	0.18	3.49
10	Tophandler Unloads	5	450	7	0.8	25	0.69	0.73
11	NII Inspection	20			20	20		32.59
	-2 RIID units		75	5			1.19	
	-1 9MeV x-ray		7,500	5			58.31	
12	Tophandler Loads	5	450	7	0.8	25	0.69	0.73
13	Truck Moves	40	20	10	6.42	15	0.18	3.49
14	Yard Crane Deposits	15	1,200	25	1.2	50	1.04	2.18
	Total						63.34	45.38

Table 2: Cost per Inspected Container at Terminal A, 5% Inspection Rate, 1 NII Station

The process times listed in Table 2 match those used in the simulation analysis, described in Section 5.1. In most cases labor costs are straightforward: they represent the number of FTEs and 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 5 people per shift is required to man the inspection station. With 3 shifts a day, and a backup team of 5 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 US-bound containers are inspected at a 5% rate at Terminal A and at a 1.5% rate at Terminal B. We also assume that US government money is used to finance the inspection scheme and apply an annual risk-free rate of 3.7%.

The far right two columns Table 2 show the discounted 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 US-bound containers, rather than just the 5% that are inspected, then per-container cost drops to $5.44. With 2 servers and a 10% inspection rate, the analogous cost numbers are $92.42 and $9.24

per container, respectively.[14]

For Terminal B, the discounted inspection costs are \$131.18 per inspected container, when 1.5% of the containers are inspected with 1 server. When cost is shared by all US-bound containers the expense drops to \$1.97 per US-bound container. With 2 servers and 3% inspection rate, the cost figures are \$110.91 and \$3.33 respectively.

6 Model and Results for Modified-SFI/Industry-centric Regime

Under this inspection scheme every container that arrives at the terminal immediately undergoes a primary inspection. Those containers that trigger an alarm during primary inspection are then tagged for more careful secondary inspections. The entire dwell time of a container is available to complete the inspection, as opposed to just 24 hours in the case of the CSI regime.

A high-level schematic of the container flow in the Industry-centric inspection regime is provided in Figure 7. The Industry-centric inspection protocol is outlined in boxes 2-5. Boxes 6-11 correspond to boxes 13-18 in the CSI protocol, shown in Figure 4. The system performance measures used in this context are the same as the ones used for the CSI regime. In the Industry-centric simulation, queuing occurs at the primary inspection step (box 2) and the secondary inspection step (box 5). We track the average and maximum queue length attained at both these steps.

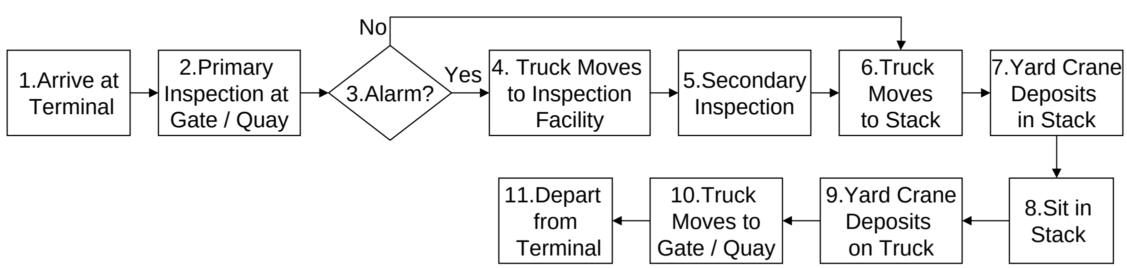

Figure 7: Typical Container Flow for Industry-centric Regime.

6.1 Model Details

In this section we describe the details of the Industry-centric simulation model. Again, we describe each process step, along with the critical statistics associated with our simulation calculations. As with the CSI simulation, we have estimated statistics related to inspection times and outcomes

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

in collaboration with Dr. Charles Massey of Sandia National Laboratories. We summarize these statistics in Table 3.

Primary Inspection at Gate/Quay (2) - Alarm? (3)

As noted above, every container undergoes primary inspection upon its arrival to the terminal. This inspection is comprised of an RPM scan followed by non-intrusive imaging using medium energy x-ray radiography. Both scans are conducted using drive-through portals.

We note that, for simplicity, our simulation assumes that all primary inspection stations are pooled together and that all containers – whether arriving via truck or vessel – are sent to this central pool. In reality, however, these stations would be split into two pools, one for arrival trucks at the terminal gates and the other for arrival vessels at quayside. Our simplifying assumption will tend to underestimate the number of inspection stations required to provide a given level of waiting-time performance, hence provide a lower bound on the cost incurred for primary inspections. Nevertheless, the simulation results reported in §6.2 and cost estimates developed in §6.3 suggest that the overall level of primary inspection cost is low enough that the incremental cost of extra stations would not be substantial.

During a primary inspection, there exist three potential ways in which a container can trigger an alarm: radiation from the container can trigger the RPM alarm; or dense material may be detected during x-ray radiography; or both tests could indicate something suspicious. If any of these events occurs, the container is directed to a secondary inspection facility, irrespective of whether or not it is US-bound. If primary inspection does not trigger an alarm, however, the container is deemed benign and sent to a designated yard location where it awaits its scheduled departure, as in the baseline case.

Timing Assumption. A truck carrying a container rolls past an RPM at slow speed (about 5 mph), while the RPM captures a sample of the emitted radiation. This scan time is essentially the time the truck takes to drive through an RPM (Wein et al. 2007). Similarly, the time spent in capturing the x-ray image of the container is equal to the time it takes for a truck moving at 5 mph to roll past the x-ray machine. We assume that the time to complete the primary inspection – roll through the RPM scanner, travel the distance between the RPM and x-ray portals, and complete the x-ray scan – is a fixed 25 seconds.

Alarm Rate. We assume that 5% of the containers trigger the RPM alarm, while x-ray radiography indicates dense material in 1% of the containers. We also assume that the results of the RPM and radiographic scans are statistically independent of each other. The assumption that the two tests are independent is an approximation we make in the absence of concrete data that suggests otherwise. However, we believe these percentages to be weighted towards the higher end of likely alarms based on the alarms experienced by the Department of Energy in its overseas operations.

We also assume that the alarms are triggered through an automated process. This is straightforward to implement for an RPM, by setting a threshold level of each kind of radiation that is considered dangerous. For x-ray radiography, automated detection requires the use of rules-based software that can pick up any unexpected dense material inside the container. Similar technology is used in other contexts such as food safety, and could be adapted for use in the context of container security (Kwon et al. 2008).

Number of Inspection Stations. The primary inspection capacity is a critical determinant of the congestion and queue buildup at the terminal gates. In our simulations we conduct a sensitivity analysis with respect to the number of primary inspection stations (sets of RPM and x-ray scanners) and vary them from 5 up to 9.

Truck Moves to Inspection Facility (4) - Secondary Inspection (5)

A container is sent to the inspection facility, if required, on the same truck that carries it to the primary inspection station. In contrast to the CSI regime, a container undergoing secondary inspection does not need to be unloaded and reloaded by tophandlers at the inspection facility, since secondary inspection is conducted in a drive-through fashion that is similar to that used for primary inspections.

During secondary inspection the container is again subject to two forms of NII: a scan by an Advanced Spectroscopic Portal (ASP), followed by another scan using medium-energy x-ray radiography. Both tests are conducted in a drive-through fashion.

While RPMs are unable to distinguish between naturally occurring radioactive material present in common substances, such as fertilizers and ceramics, and more dangerous materials, the ASP can better discriminate what the source of radioactivity may be. The second medium-energy x-ray scan is conducted under more controlled circumstances and at slower drive-through speeds, when compared to the primary stage, and results in better image resolution.

Timing Assumptions. We assume that the ASP inspection time is lognormally distributed with a mean of 45 seconds and a standard deviation of 45 seconds. The distribution of the inspection time for the medium energy x-ray machine is identical to (but independent of) that for the ASP.

Number of Inspection Stations. We assume that one set of inspection stations is available for secondary inspection - that is, one each of the ASP and medium energy x-ray scanner. This capacity level proves to be more than sufficient, since the secondary-inspection load in the Industry-centric regime turns out to be quite small.

Box	Process Step	Process Time Distribution	Mean	Standard Deviation
2	Primary Inspection at Gate / Quay			
3	Alarm?	fixed	25 sec	0 sec
4	Truck Moves to Inspection Facility	omitted / negligible		
5	Secondary Inspection			
	- ASP scan	lognormal	45 sec	45 sec
	- Medium energy x-ray scan	lognormal	45 sec	45 sec

Table 3: Process-Time Assumptions Used in Simulations of Industry-centric Protocol

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 5 to 9. As before, the performance measures of interest are the fraction of inspected containers that miss their historical departure times and the sizes of the queues of containers waiting to be inspected.

We find that the number of containers that miss their scheduled departure is steady at approximately 2500. These containers have very short dwell times, under 6 hours. Moreover, none of these containers is US-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 dropoff 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 inspection regime.

Under the Industry-centric regime, there are points at which containers may wait for inspection – primary and secondary – hence there are two sets of queues. Even with only one inspection station, the secondary inspection queue is not significant, however. The average queue length varies between 3 and 4 containers, while the maximum queue length attained varies between 52 and 62, as the number of primary servers is increased from 5 up to 9. 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. Since 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, in our simulations, the arrival rate varies considerably from hour to hour. Nevertheless, the capacity of single station is adequate to prevent significant delays.

Conversely, the queue of containers waiting for primary inspection can become significant, and this can have serious economic consequences. For example, tolerance for congestion at terminal gates is particularly low at many 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 trucks, the containers cannot be stacked, and the associated real estate requirements can be large.

In Figure 8 we plot the maximum and mean queue lengths attained during primary inspection, as a function of the number of inspection stations deployed. The figure's top panel shows the results for Terminal A, and the bottom panel those for Terminal B. For Terminal A we find that, using 8 primary servers, mean queue lengths are quite acceptable: 0.3 containers at both Terminal A and Terminal B. Because of extreme arrival-rate peaks, the maximum queue lengths at the terminals can be large: 136 at Terminal A and 53 at Terminal B. At roughly 100 containers per acre this can become large - this translates into 1.36 and 0.53 acres at the terminals.[15]

At this point it is worth recalling from §6.1 that our simulation models make the simplifying assumption that all inspection stations are pooled together in a single location. If separate pools of stations were placed at the terminal's in-gate and at the quay, however, each smaller pool might require one extra station in order to maintain the same waiting-time performance as that achieved with a single queue. This would increase the cost of primary inspection equipment by 10/8, or 25%. At the same time, together the numbers of containers waiting in the two queues would total something on the order of 132 or 53, at most; each of the two queues would have a length that is a fraction of the total.

6.3 Cost Estimate

As in the case for CSI regime, the Industry-centric regime gives rise to equipment and labor costs which we allocate over the set of inspected containers. In Table 4 we provide the details of the cost breakdown mapped onto the process steps outlined in Figure 7.

Besides differences in the testing equipment used, the logistics of the process for Industry-centric regime differs from that for 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

[15] 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.

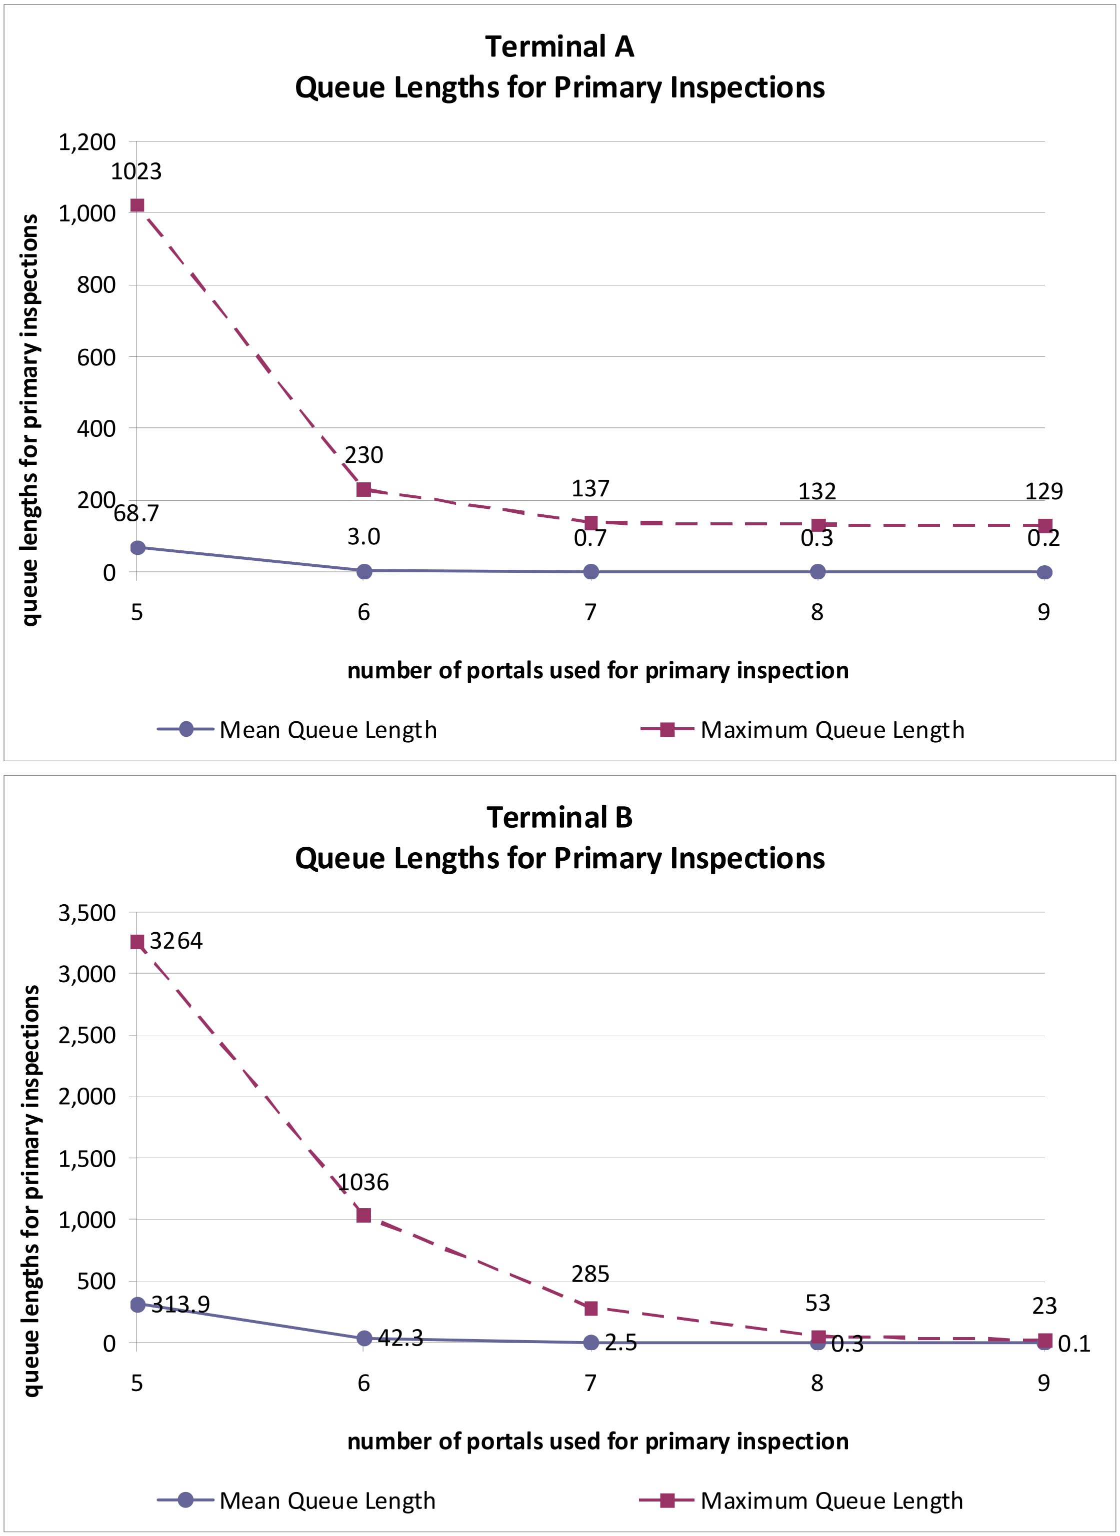

Figure 8: Maximum and Mean Queue Length during Primary Inspection: Industry-centric Regime.

operational overhead on account of the inspection regime. 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 4 and those calculated for CSI concerns the discount rate. For CSI, we assumed that the US 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 4's results show that, even with this higher discount rate, the per-container costs under the Industry-centric regime are much lower than those under CSI. At Terminal A, we find that primary inspection costs roughly $1.63 per container, while the cost of inspection is about $14.14 for each container that undergoes both primary and secondary inspection. The per-container cost reduces to just $2.47 if secondary inspection costs are allocated across all container traffic, as opposed to just the inspected containers. At Terminal B, primary inspection costs $1.6 per container. For containers undergoing both primary and secondary inspection, the per container cost is $12.32. If the secondary inspection cost is allocated across all container traffic then the per container cost turns out to be $2.33.

Box	Process Step	Average Time (sec)	Equipment Unit Cost ($000)	Equipment Life (yrs)	Labor FTEs	Labor Comp. ($000)	Cost Per Container Equip. ($)	Cost Per Container Labor ($)
2	Primary Inspection	25			4	10	1.59	0.01
	- 8 RPMs		250	5				
	- 8 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		
	- 1 ASP	45	600	5			9.51	1.20
	- 1 medium-energy x-ray	45	4,000	5				
	- 1 9MeV x-ray (backup)		7,500	5	5	45		
	Total						11.11	1.21

Table 4: Cost per Inspected Container at Terminal A: 8 Primary Stations, 1 Secondary Station

Again, we recall from §6.1 and §6.2 that, to adjust for the fact that our simulation model assumes a single pool of primary inspection stations may lead us to underestimate the cost of primary inspection equipment by an amount that is on the order of 25%. From Table 4, we see that 25% of $1.59 per container is only about 40 cents per container.

7 Discussion and Conclusions

Our simulation results show that the modified-SFI/Industry-centric regime should be able to provide better inspection coverage than CSI at a lower unit cost. With (the current) one CSI inspection station per terminal, Terminal A could sustain an inspection rate of 5% and Terminal B 1.5%. Under CSI, equipment and labor costs at both ports totalled on the order of $100 per inspected container. In contrast, the Industry-centric inspection regime had the ability to sustain a 100% inspection load at a lower cost per container. With 8 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.

In both regimes, stable inspection loads led to stable requirements for staging area, though at different locations in the terminal. Under CSI, the above inspection rates required 0.8 acres of secondary inspection staging area at Terminal A and 0.3 acres at Terminal B. Under the Industry-centric regime, the maximum queue lengths at the primary inspection station were a bit higher, requiring on an average 1.36 acres at Terminal A and 0.53 acres at Terminal B.

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 2 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 would significantly change CSI's capacity limits. Simulation results for these variants, reported in a separate Appendix that can be requested from the authors, show that their performance is similar to that of the original CSI scheme.

The Industry-centric regime's cost and coverage advantages follow from two sources. First, the use of higher-capacity drive-through equipment allows for 100% primary inspections. Second, the lower cost per 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 simulation analysis is that, no matter which inspection scheme is used, the equipment and inspection protocol used in the Industry-centric scheme should become the standard by which containers are inspected at international ports.

Still, 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 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 potentially significant cost. In addition, CBP is actively pursuing additional data from US importers 24-hours in advance of loading so that CBP can improve its ability to target high risk containers. US importers must bear the additional cost associated with providing this data. To provide a sense of the potential cost, we note that the National Association of Manufacturers (NAM) has recently voiced concern about the cost of the so-called "10+2" reporting requirements CBP has proposed to support the intelligence needed for CSI, claiming that the time required to gather the required information will add 2 to 5 days to containers' stays at terminals, at an annual cost to trade of $8.5 billion per day of delay (NAM 2008).[16] Given this concern, the NAM has requested that CBP undertake a pilot program so that the impact of "10+2" can be more properly evaluated.

While these information-driven delays differ from the logistical delays associated with the untimely pulling of containers, the scale of the economic penalties is 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.[17]

In contrast, the Industry-centric regime's scan-upon-entry process all together avoids the need to pull containers from terminal stacks. Further, the only data 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 US-bound containers upon arrival to the terminal, even if classification of containers by destination is difficult at that stage. Of course, the scanning of US-bound containers upon arrival using drive-through equipment is a scheme that is essentially a restricted form of SFI, and one might ask what is the value of inspecting containers not bound for the US.

In fact, we believe that there are positive externalities associated with the scanning of all containers, irrespective of destination. The images provided by such a protocol provide an effective complement to field intelligence and CBP's risk-scoring capabilities when identifying potentially dangerous containers. 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..

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

[17] 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.

Furthermore, the systematic use of NII also has economic advantages for the US. Rather than penalizing only US-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. From their perspective, this expenditure would be part of the investment required for business continuity purposes, given the threat from maritime terrorism and the associated system-wide disruption that a major terrorist incident could cause. Given the possibility that these economic benefits do, in fact, induce the terminal operators to privately fund such a system, we believe that "industry-centric" is an appropriate name for this inspection regime.

Still, challenges also remain with the successful implementation of the Industry-centric regime. If terminal operators are to take responsibility for purchasing, deploying and operating inspection equipment, then the US government must work with them to establish appropriate technical standards for inspection processes and equipment. Moreover, government must also ensure the effective operation of inspection equipment by the terminal operators. One means of achieving this end would be to set up a 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 also remain IT challenges. Under the Industry-centric regime, a large part of the primary inspection's anomaly-detection process must be automated. There is also a need to deploy IT infrastructure to communicate container-scan and manifest information in real time to CBP officials in the US and 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 and the ability to conduct their own remote independent assessment of cargo they determine to be at higher risk.. It would also provide for a valuable audit trail if something untoward were to happen at a US or overseas port. Concerns remain with regard to feasibility, sizing and costing of such an IT setup. While we do not address these questions as part of this research, they warrant further investigation.

Finally, we note that we have conducted our analysis at two of the world's busiest terminals. Terminals can be quite different from each other with respect to their layout and cost of operations. It would be hard to imagine a "one-size-fits-all" approach could work for all international terminals. Nevertheless, 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 US will encounter in its efforts to bolster international container security.

References

Abt, C. C. 2003. The Economic Impact of Nuclear Terrorist Attacks on Freight Transport Systems in an Age of Seaport Vulnerability. Report by Abt Associates Inc.

Bakshi, N. and N. Gans. 2007. Securing the Containerized Supply Chain: An Economic Analysis of C-TPAT. Working paper, The Wharton School, University of Pennsylvania.

Bennett, A. C. and Y. Z. Chin. 2008. 100% Container Scanning: Security Policy Implications for Global Supply Chains. Master's Thesis, MIT.

Boske, L. B. 2006. Port and Supply-Chain Security Initiatives in the United States and Abroad. Report by Lyndon B. Johnson School of Public Affairs, UT Austin.

CBP. 2004. FAQ on 24-Hour Advance Vessel Manifest Rule. http://www.customs.gov/ImageCache/cgov/content/import/carriers/24hour_5frule/24hour_5ffaq_2edoc/v1/24hour_5ffaq.doc. Accessed online on 05/24/08.

CBP. 2008-a. CBP documents on CSI. http://www.cbp.gov/xp/cgov/trade/cargo_security/csi/. Accessed online on 05/24/08.

CBP. 2008-b. CBP documents on SFI. http://www.cbp.gov/xp/cgov/newsroom/fact_sheets/trade_security/sfi/. Accessed online on 05/24/08.

Eagle Marine Services. 2008. http://www.eaglemarineservices.com/images/lgGGS1.gif. Accessed online on 07/10/08.

Fairnie, D. 2008. DP World Director of Security, Speech at World Customs Organisation Forum: What Future for 100% Scanning? Brussels, June 11, 2008.

Flynn S. E. 2004. America the Vulnerable: How Our Government is Failing to Protect Us from Terrorism. Harper Collins.

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

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.

GAO. 2005. Container Security: A Flexible Staffing Model and Minimum Equipment Requirements Would Improve Overseas Targetting and Inspection Efforts. Report GAO-05-557.

GAO. 2008-a. 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.

GAO. 2008-b. Supply Chain Security: Challenges to Scanning 100 Percent of U.S.-Bound Cargo Containers. Report GAO-08-533T.

GAO. 2008-c. CBP Works with International Entities to Promote Global Customs Security Standards and Initiatives, but Challenges Remain. Report GAO-08-538.

Harrald, J. R., H. W. Stephens and J. R. vanDorp 2004. A Framework for Sustainable Port Security. *Journal of Homeland Security and Emergency Management*, Vol. 1 : Iss. 2, Article 12.

Kunreuther, H. C., P. J. McNulty, and Y, Kang. 2002. Third-Party Inspection as an Alternative Command and Control Regulation. *Risk Analysis*, 22:2, 309-318.

Kwon, J. S., J. M. Lee, and W. Y. Kim. 2008. Real-time detection of foreign objects using X-ray imaging for dry food manufacturing line. *IEEE International Symposium on Consumer Electronics, 2008*, 14-16 April, 1-4.

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

Marine Link. 2004. The 5% Myth. http://www.marinelink.com/Story/The+5%25+Myth-15423.html. Accessed online on 08/29/08.

Martonosi, S. E., D. S. Ortiz and H. H. Willis. 2006. Evaluating the Viability of 100 Percent Container Inspection at America's Ports. Report by RAND Corporation.

McClure, G. 2007. How Safe Are Our Ports? http://www.todaysengineer.org/2007/Sep/port-security.asp. Accessed online on 08/14/08.

National Association of Manufacturers 2008. Comments of the National Association of Manufacturers on Notice of Proposed Rulemaking by U.S. Customers and Border Protection, Department of Homeland "Importer Security Filing and Additional Carrier Requirements" Docket ID: USCBP-2007-0077. http://www.nam.org/∼/media/Files/PDF/2008%2008/comments.ashx. Accessed online on 09/21/08.

Steenken, D., S. Voβ, and R. Stahlbock. 2004. Container terminal operation and operations research - a classification and literature review. *OR Spectrum*, 26:1, 3-49.

US Congress. 2007. H.R. 1: Implementing Recommendations of the 9/11 Commission Act of 2007. http://www.govtrack.us/congress/bill.xpd?bill=h110-1. Accessed online on 05/24/08.

Wein, L. M., A. H. Wilkins, M. Baveja, and S. E. Flynn. 2006. Preventing the Importation of Illicit Nuclear Materials in Shipping Containers. *Risk Analysis*, 26:5.

Wein, L. M., Y. Liu, Z. Cao, and S. E. Flynn. 2007. The Optimal Spatiotemporal Deployment of Radiation Portal Monitors Can Improve Nuclear Detection at Overseas Ports. To appear in *Science & Global Security.*

A Appendix: Confidence Intervals for CSI Simulation

24-Hour Rule, Inspections First-Come-First-Served: Terminal A								
	1 Inspection Station				2 Inspection Stations			
Inspection Rate	Fraction Delayed	Half Width 95% C.I.	Max. Staging Area	Half Width 95% C.I.	Fraction Delayed	Half Width 95% C.I.	Max. Staging Area	Half Width 95% C.I.
4%	1.1%	0.6%	0.4	0.02	0.0%	0.0%	0.1	0.01
5%	16.3%	1.5%	0.8	0.03	0.0%	0.0%	0.2	0.01
6%	62.2%	3.6%	1.3	0.03	0.0%	0.0%	0.4	0.02
7%	99.8%	0.5%	2.9	0.09	0.0%	0.0%	0.5	0.02
8%	99.9%	0.8%	5.1	0.17	0.5%	0.3%	0.7	0.02
9%	100.0%	0.9%	7.2	0.18	4.3%	0.8%	1.1	0.04
10%	100.0%	0.8%	9.4	0.16	14.2%	1.2%	1.5	0.04
11%	99.8%	0.7%	11.5	0.16	29.4%	1.3%	2.0	0.05
12%	99.8%	0.7%	13.6	0.18	57.7%	2.9%	2.5	0.05
13%	99.8%	0.7%	15.8	0.19	94.9%	1.4%	3.6	0.09
14%	99.8%	0.6%	18.0	0.17	99.8%	0.6%	5.7	0.20
15%	99.8%	0.6%	20.2	0.18	99.8%	0.6%	7.8	0.19
16%	99.8%	0.5%	22.5	0.19	99.8%	0.5%	9.9	0.19
17%	99.8%	0.5%	24.7	0.21	99.8%	0.5%	12.1	0.18
18%	99.8%	0.5%	27.0	0.21	99.8%	0.5%	14.3	0.22
19%	99.8%	0.5%	29.3	0.22	99.8%	0.5%	16.4	0.21
20%	99.8%	0.5%	31.6	0.23	99.8%	0.5%	18.6	0.23

24-Hour Rule, Inspections First-Come-First-Served: Terminal B								
	1 Inspection Station				2 Inspection Stations			
Inspection Rate	Fraction Delayed	Half Width 95% C.I.	Max. Staging Area	Half Width 95% C.I.	Fraction Delayed	Half Width 95% C.I.	Max. Staging Area	Half Width 95% C.I.
1%	0.0%	0.0%	0.2	0.01	0.0%	0.0%	0.1	0.00
1.5%	13.7%	2.0%	0.6	0.02	0.0%	0.0%	0.3	0.02
2%	99.1%	0.8%	3.5	0.11	0.0%	0.0%	0.3	0.01
3%	100.4%	0.9%	12.4	0.25	9.4%	1.4%	1.1	0.03
4%	100.2%	0.8%	21.4	0.21	100.1%	0.7%	7.1	0.22
5%	99.9%	0.6%	30.1	0.22	99.9%	0.6%	15.8	0.25
6%	99.9%	0.6%	39.0	0.24	99.9%	0.6%	24.6	0.27
7%	100.0%	0.6%	47.9	0.29	100.0%	0.6%	33.5	0.30
8%	100.0%	0.5%	56.8	0.28	100.0%	0.5%	42.4	0.29
9%	99.9%	0.5%	65.7	0.32	99.9%	0.5%	51.3	0.35
10%	100.0%	0.5%	74.6	0.38	100.0%	0.5%	60.3	0.36
11%	99.9%	0.4%	83.4	0.36	99.9%	0.4%	69.1	0.34
12%	99.8%	0.4%	92.3	0.39	99.8%	0.4%	78.0	0.36
13%	99.9%	0.4%	101.2	0.40	99.9%	0.4%	86.8	0.37
14%	99.8%	0.4%	110.1	0.44	99.8%	0.4%	95.7	0.40
15%	99.8%	0.4%	119.0	0.45	99.8%	0.4%	104.6	0.42
16%	99.8%	0.4%	127.9	0.48	99.8%	0.4%	113.5	0.44
17%	99.8%	0.4%	136.7	0.51	99.8%	0.4%	122.3	0.50
18%	99.8%	0.4%	145.6	0.53	99.8%	0.4%	131.2	0.50
19%	99.8%	0.4%	154.5	0.50	99.8%	0.4%	140.1	0.47
20%	99.8%	0.4%	163.4	0.53	99.8%	0.4%	149.0	0.50

Figure 9: 95% Confidence Intervals for Simulations of CSI Regime.