

The Value of Advance Disruption Information on Transportation Gateways: The Inventory a Warning Replaces

James Rappold

jrappold@rollins.edu

Rollins College

Research Article

Keywords: Attack warning, Value of information, Cyber-kinetic attack, False alarms, Inventory hedging, Transportation gateway

Posted Date: October 26th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11252684/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

THE VALUE OF ADVANCE DISRUPTION INFORMATION ON TRANSPORTATION GATEWAYS: THE INVENTORY A WARNING REPLACES

James A. Rappold

*Crummer Graduate School of Business
Rollins College, Winter Park, FL 32789, USA*

Email: jrappold@rollins.edu

ORCID: [0009-0003-7161-4384](https://orcid.org/0009-0003-7161-4384)

Corresponding author: James A. Rappold, jrappold@rollins.edu

THE VALUE OF ADVANCE DISRUPTION INFORMATION ON TRANSPORTATION GATEWAYS: THE INVENTORY A WARNING REPLACES

Author names and affiliations withheld for anonymous review.

October 2026

Abstract

An attack on a transportation gateway can be detected before it takes effect. The warning is imperfect and costly to act on. We examine what a warning is worth to the defender.

A facility is replenished through a gateway that an attacker closes at random for a random duration. A detector warns of an attack L periods ahead, with false alarm probability α and miss probability β . The defender hedges by acquiring inventory through the gateway's limited surge capacity. We define the value of advance disruption information through the lens of inventory as the reduction in the defender's minimum average cost of holding inventory, carrying a backlog, and surging.

The optimal policy is a band with two levels for each belief about the attack. We prove it under a discounted criterion and stated conditions, and computed average-cost policies have the same form. When alarms are sparse, the bands sort inventory into three tiers: units held whenever a closure is possible, units acquired on an alarm, and units never held. As a first-order approximation, the value is concave in the lead time and stops rising once the lead time suffices to acquire the middle tier. A gateway near capacity acquires little inventory in any lead time, so a warning can be worth least where needed most. Beyond a false alarm probability near the ratio of the holding cost to the surge premium, most of the value disappears. Iso-value curves price a detector by its operating point.

Keywords: Attack warning; Value of information; Cyber-kinetic attack; False alarms; Inventory hedging; Transportation gateway.

1 Introduction

A deliberate attack on a transportation gateway differs from most natural disasters in one respect that matters to its defender. The attack can be detected before it takes effect. A cyberattacker who intends to halt a container terminal, the ground operations of an airport, or the dispatching of a railroad must first gain access, move through the network, and position the payload. Each step can leave a trace, whether an intrusion alert, a sensor anomaly, or a pattern of interference with satellite navigation. The interval between the first trace and the physical effect is a *warning lead time*. An earthquake offers seconds; a hurricane offers days, which a forecast supplies, and a protest that gathers to blockade a port or a road offers the notice its organizers give. What we say of an attacker below applies to the hurricane and the protest as well.

The warning is imperfect, and acting on it is expensive. A detection system raises false alarms, and because attacks are rare, the false alarms can outnumber true alarms many times over, even when the detector is accurate (Axelsson, 2000). Some attacks arrive with no alarm at all. The defender who acts on an alarm can acquire inventory on the far side of the gateway, reserve capacity on an alternate mode, or reroute freight. It incurs the cost of the precaution whether or not the attack follows. When ransomware reached the business systems of Colonial Pipeline in May 2021, the company itself shut the pipeline down for five days as a precaution (Cybersecurity and Infrastructure Security Agency, 2023). The defender who ignores the alarm incurs no cost of precaution, and suffers a loss only if the gateway closes.

When it comes, the disruption can be long, and its length is hard to predict. When the NotPetya malware reached A.P. Moller-Maersk in June 2017 it halted booking and terminal operations across the company's network. The core systems took about ten days to rebuild (Greenberg, 2018), and the company put the cost at USD 250 to 300 million (Maersk, 2018). A ransomware attack on the Port of Seattle in August 2024 left its airport open but disrupted its check-in systems and its baggage and flight-information displays. The Port reported the last of them restored on 11 September, about three weeks later, and its public website restored about three months after the attack (Port of Seattle, 2024; Coogan and Bosch, 2024). An attack on the ground systems of LOT Polish Airlines in June

2015, by contrast, was over in about five hours (Reuters, 2015). Hurricane Sandy, forecast days ahead, closed the Port of New York and New Jersey before its landfall on 29 October 2012; the port reopened to vessel traffic within five days, and its container terminals began reopening within a week (Port Authority of New York and New Jersey, 2012; U.S. Department of Energy, Office of Electricity Delivery and Energy Reliability, 2012; Universal Cargo, 2012). These accounts show how widely closures vary in length, but they are not a sample from which to estimate a distribution, so Section 6 tests the model against a heavy tail instead.

These observations raise the question that a port authority, an airport operator, or a logistics command must answer before buying any detection capability. What is a warning worth? A detection system presents itself to its buyer through three numbers. Its lead time L is the interval between the alarm and the effect. Its *false alarm probability* α is the probability that it raises an alarm in a period in which no campaign is in view. Its *miss probability* β is the probability that an attack arrives without a true alarm. In the language of hypothesis testing, with no campaign in view as the null hypothesis, α and β are the detector's Type I and Type II error probabilities and $1 - \beta$ is its power; the test is repeated every period, so false alarms accumulate over an open spell. We define the *value of warning*, $V(L, \alpha, \beta)$, as the reduction in the defender's minimum long-run average cost when the signal is available, so that a detection system priced at w per period is worth buying if and only if

$$V(L, \alpha, \beta) > w. \quad (1.1)$$

We compute V , show how it depends on each of the three numbers, and identify the lead time beyond which, as a first-order approximation, V stops rising (which depends on the surge capacity), and the false alarm probability beyond which most of the value disappears.

In reality, an attack costs more than inventory. It can cost lives, destroy property and infrastructure, and stop the flow of relief, and none of that is in V . We adopt a narrow focus on purpose. We quantify what advance disruption information is worth through the lens of inventory, the pre-positioning of supply before a closure and the backlog after it, which is one portion of the total cost of an attack and the portion a logistics command controls. Our decision framework does not

depend on that choice. The cost function of Section 3 can be replaced by another, such as the cost of a delayed evacuation or of relief aircraft held idle. The value of warning and the pricing of a detector are then defined as they are here, but the thresholds must be derived afresh.

We study a single network link. Suppose a facility (or a region) consumes a commodity at a steady rate and is replenished through one gateway such as an airport or seaport. An attack is initiated with probability λ in each period in which none is under way, takes M periods to mount, and then closes the gateway for a random duration T . A detector raises an alarm $L \leq M$ periods before the closure. While the gateway is open, the defender can move more than current consumption through it, up to a surge capacity C , and can accumulate anticipation inventory against the closure. While it is closed, the defender serves demand from that inventory, and demand that cannot be served is backordered at a cost for every period it waits, until the gateway reopens and its surge capacity clears it. The defender cannot see the attacker's countdown and acts on a belief about it, so its problem is a partially observed Markov decision process. A *belief* is the defender's probability distribution over the number of periods that remain before the gateway closes, including the possibility that no campaign is under way at all. Specifically, it is the posterior distribution of the countdown, given what the defender has observed in every period so far, namely whether the gateway was open and whether the detector raised an alarm. We solve that problem exactly on the belief, and we approximate it in closed form by following each unit of inventory separately, which is what yields the policy thresholds. We assume that the attacker is not strategic, which is to say that λ does not respond to what the defender does, the inventory it holds, the hedge it acquires on an alarm, or the detector it buys. We believe this assumption is appropriate because it isolates what the warning is worth before the attacker adapts to it. Section 7 discusses what a strategic attacker changes.

We call the framework *advance disruption information*. It is the supply-side counterpart of advance demand information, in which a customer's early order tells a firm about demand L periods ahead (Hariharan and Zipkin, 1995; Gallego and Özer, 2001). Here, an alarm tells the defender about supply conditions L periods ahead. The two differ in two respects. First, an advance order

is a commitment, or close to one, and an alarm is not. Many alarms are false. Second, the event being signaled is rare and severe, so the base rate of attack governs what an alarm means, and an accurate detector can still be a worthless one.

Two bodies of literature are close to our problem. Inventory theory shows how to hedge a link that fails at random (Parlar and Berkin, 1991; Tomlin, 2006; Snyder et al., 2016). It also shows how to use information about the state of supply, whether exact as in Song and Zipkin (1996) and Jakšič et al. (2011) or a noisy reading of a hidden environment in Arifoğlu and Özekici (2010, 2011). It has little to say about a warning: a signal that arrives a lead time ahead of a disruption, is often wrong, and is described by its errors. Detection theory and the economics of intrusion detection price false alarms and misses with care (Gaffney and Ulvila, 2001; Cavusoglu et al., 2005), as does the older cost-loss model of forecast value in meteorology (Murphy, 1977; Richardson, 2000). In the cost-loss model, the protective action is a single choice with a fixed cost and a fixed loss, and the detection models keep that structure while adding the detector’s configuration and the attacker’s incentives. In none of them is there inventory to acquire, a lead time in which to acquire it, or a capacity limit on how fast it can be acquired.

We make four contributions in this paper. First, without a warning, the defender’s optimal quiet inventory is, as a first-order approximation, a fractile of the equilibrium distribution of the closure duration T , so the variability of T governs it. The critical ratio weighs the holding cost of the wait for the attack against the backlog cost of the closure, and it sits well below the newsvendor’s. This is the *uninformed defender*’s problem, and its cost is the benchmark from which we measure the value of warning. When T has a geometric tail, the optimal inventory grows with the logarithm of $\lambda b/h$, the attack probability times the ratio of the backlog cost to the holding cost, and when the tail is heavy, it grows faster. The recorded durations of closures after attacks on gateways range from hours to months.

Second, for the *informed defender* with a single warning raised L periods before the closure, the *tripwire*, the optimal order under each belief is set by a band with two levels, which we prove under a discounted criterion and find in every average-cost instance we solve. Without a surge premium,

the band reduces to a modified base-stock policy limited by the surge capacity. When alarms are sparse, the policy has a simple shape. The defender establishes and restores a quiet inventory S_0 , surges to an alert inventory S_A on an alarm, and lets consumption draw the inventory back down over time. That shape categorizes each unit of inventory into one of three tiers. The lowest S_0 units are held whenever a closure is possible and protect against the attacks the detector misses. The next $S_A - S_0$ units, the hedge, are acquired only on an alarm, in time for the closure. Units above S_A are never held, and the defender accepts that exposure. Following each unit separately gives a closed-form approximation to the value of warning, as a sum of cost-loss decisions, one per unit. Thus, for a warning system with attributes (L, α, β) , the approximation gives a rule for when to hold a unit of *part-time inventory*, just in case.

Third, we show that the value V is non-decreasing in L . As a first-order approximation, it is also concave in L and constant beyond a critical lead time L^* , the number of periods the surge capacity needs to acquire the middle tier. The increase in V from one more period of lead time, which we call the *value of lead time*, depends on the surge capacity.

Fourth, as a first-order approximation, we find in closed form a critical false alarm probability α^* , the *cry-wolf threshold*, beyond which acting on an alarm no longer pays. The exact policy keeps acquiring a fraction of a unit on an alarm well beyond it, but most of the warning's value is gone by then. When the uninformed defender holds inventory, the threshold is approximately $(1 - \beta)h/p$, the ratio of the holding cost to the surge premium, scaled by the detection probability, and the attack probability λ drops out. With a premium below about the holding cost, the threshold exceeds $1 - \beta$, and the approximation acquires on every alarm. When attacks are too rare for the uninformed defender to hold anything, the threshold is nearly proportional to λ . The rarer the attack, the fewer false alarms the detector may raise. The value V is also the most a detector is worth paying for, by (1.1), and it prices a detector by its operating point. The curves of constant first-order value in the plane of false alarm probability against detection probability have a closed-form slope, the rate at which the defender trades false alarms for detections.

The remainder of the paper is organized as follows. In Section 2, we review the related literature. In Section 3, we present the advance disruption information framework and model the gateway, the threat, the warning, and the optimal policy. In Section 4, we solve for the uninformed defender. In Section 5, we solve for an informed defender with the single warning and derive the value of warning and its thresholds. In Section 6, we present a numerical study of a single-runway airfield, solved exactly and by simulation. We verify the closed-form approximations against the exact solution, show that lead time keeps adding value longer at high utilization, price the false alarms of a detector, and test the model against a heavy-tailed closure, random demand, and lost demand. In Section 7, we discuss strategic attackers and the limits of the model. We present conclusions in Section 8.

2 Literature Review

Four areas of the literature are relevant to this paper: inventory theory under supply disruption, models for supply state information, the theory of imperfect warning signals, and transportation security.

2.1 Inventory Under Supply Disruption

The study of a supply source that alternates between available and unavailable begins with Parlar and Berkin (1991), who add random supply interruptions to the economic order quantity model. Parlar et al. (1995) treat periodic review with a Markovian supply state and show that an order-up-to policy is optimal. Gürler and Parlar (1997) consider two such suppliers, and Arreola-Risa and DeCroix (1998) allow a mix of backorders and lost sales during the interruption. Schmitt et al. (2010) combine disruptions with stochastic demand and give closed-form approximations to the base-stock level. Snyder et al. (2016) review the field.

Two results from this area shape the uninformed defender’s policy. First, Tomlin (2006) shows that the right defense depends on the disruption profile, inventory suiting frequent and short

disruptions, and a contingent response suiting rare and long ones. An attack on a gateway is rare and closes it for a long time, which is the profile for which quiet inventory is least attractive and information most valuable. Second, Saghafian and Van Oyen (2016) show that average reliability alone can mislead. Two suppliers equally available on average differ in how long and how often they fail, and the one with longer and rarer disruptions is the one worth backing up. Section 4 carries the point a step further, since both the mean and the tail of the closure duration set the uninformed defender's quiet inventory.

2.2 Supply State Information

Advance information has been studied mainly on the demand side. Hariharan and Zipkin (1995) show that a demand lead time offsets a supply lead time one for one, and Gallego and Özer (2001) show that with advance demand information a base-stock policy remains optimal in a modified inventory position that nets out the known orders. Özer and Wei (2004) add a capacity limit and show that information substitutes for capacity and for inventory. When the advance information is imperfect, because announced orders may change or fail to materialize, the policy keeps its form, but the information loses value (Tan et al., 2007; Gayon et al., 2009; Benjaafar et al., 2011).

On the supply side, Song and Zipkin (1996) let an observed Markov chain modulate the supply lead time and show that the optimal policy is a base-stock policy whose level moves with the observed state. Jakšič et al. (2011) and Altuğ and Muharremoglu (2011) give the firm a forecast of its supplier's capacity over an information horizon and again find state-dependent base-stock policies. In Gao et al. (2017) the forecast comes from a signal that evolves as a Markov chain, and in Arifoğlu and Özekici (2010) and Arifoğlu and Özekici (2011) the environment that modulates supply, demand, and costs is a hidden Markov chain read through a noisy signal; the manager carries its posterior, and base-stock policies survive with levels that move with the posterior. Tomlin (2009) studies a firm that learns a supplier's reliability from experience.

Two papers come closest to ours. Saghafian and Van Oyen (2012) study a newsvendor that reserves flexible backup capacity on a perceived reliability of its supplier, which may be in error,

and then observes the supplier's state and exercises the backup. Their information is imperfect, as ours is, but their problem is a single season with one recourse, and so it has no lead time in our sense and no recurring false alarm. In a working paper, Tomlin and Snyder (2007) study a firm that observes a threat level which modulates the probability of disruption and adjusts its base stock to it. Their threat level is the nearest antecedent of our alarm state. It is a state that the firm observes, and it is not described by the operating characteristics of a detector, so it does not price one by its lead time and its two error probabilities; imperfect threat information is the direction they name for future work.

2.3 The Value of Imperfect Warning Signals

The value of information is the difference between the best the decision maker can do with it and without it (Raiffa and Schlaifer, 1961; Howard, 1966), and Blackwell (1953) orders signals by that value in every decision problem. For a binary warning of a binary event, the standard treatment is the cost-loss model of meteorology. The decision maker incurs a cost to protect against a loss, and protects when the probability of the event exceeds the ratio of cost to loss (Murphy, 1977). Richardson (2000) writes the value of a forecast in that model through its hit rate and its false alarm rate, and Provost and Fawcett (2001) show that lines of equal expected cost in the plane of those two rates are straight, with a slope fixed by the base rate and the cost ratio. Both facts reappear in Section 5, one unit of inventory at a time.

The same calculus has been applied to intrusion detection. Axelsson (2000) shows that the base rate of intrusions is so low that the false alarm rate, not the detection rate, limits the usefulness of a detector. Gaffney and Ulvila (2001) choose a detector's operating point by expected cost, and Cavusoglu et al. (2005) show that a detector configured at the wrong point can have negative value to the firm. The cry-wolf effect of Breznitz (1984), in which credibility is lost with each false alarm, and the strategic-warning literature (Betts, 1982; Grabo, 2002) make Axelsson's point behaviorally. Our threshold α^* is its rational counterpart, the false alarm probability beyond which a defender without resentment, in our first-order approximation, acquires nothing on an alarm.

Warnings with a lead time have been modeled for hurricanes, where the forecast sharpens as landfall nears. Regnier and Harr (2006) and Regnier (2008) model the forecast as a Markov chain and value the option to wait for a better one. Lodree and Taskin (2009), Taskin and Lodree (2011), and Galindo Pacheco and Batta (2016) use forecast updates to stock and position relief supplies before the storm. Two differences separate that setting from ours. First, a storm that has formed is a real event whose forecast sharpens as it nears, while most alarms signal nothing. Second, the threat to a gateway has no season, so the false alarms recur without end.

When the threat is hidden and the signal is noisy, the defender’s problem is a partially observed Markov decision process on the posterior of the hidden state (Smallwood and Sondik, 1973; Krishnamurthy, 2016). The hidden demand state of Treharne and Sox (2002) is the inventory example. The unit-by-unit analysis of Section 5 follows the single-unit decomposition of Muharremoglu and Tsitsiklis (2008).

2.4 Transportation Security

Models of defended infrastructure have concentrated on where to defend. Brown et al. (2006) formulate the defense of critical infrastructure as a sequential game between a defender and an attacker, and Tambe (2011) collects the security games that schedule patrols and checkpoints at airports and ports. These models fix what the defender knows at the moment of attack. Gkonis and Psaraftis (2010) model container security as a problem of interdependent security, and Yolmeh and Baykal-Gürsoy (2018) schedule rail patrols against a strategic adversary. The cyber threat to gateways has been treated mainly as a matter of guidance and governance, by Adams et al. (2021) for ports under physical, cyber, and hybrid threats, by Karahalios (2025) for autonomous ships approaching port, by Oka et al. (2026) for rail, and by O’Kelly (2025) for hubs across the modes, following the agenda for maritime security set out by Véronneau and Yoho (2020). False alarms enter this literature through screening, where Dzhongova et al. (2017) measure the false alarm rates of detectors of liquid explosives. The models in this literature ask where to defend against a

strategic attacker; we take the attacker as given and ask what the defender knows, when it knows it, and what that knowledge is worth.

2.5 The Position of This Paper

We borrow the disrupted-supply model and its order-up-to policies from the first body of literature, the idea that advance information can be netted into the inventory position from the second, the description of a detector by its false alarm and miss probabilities and the belief state from the third, and from the fourth we take the application setting. What is new is their combination and the results it yields. No paper we have found describes an advance warning of disruption by a lead time, a false alarm probability, and a miss probability together with a capacity limit. So we focus on what a detector with a given operating point and a given lead time is worth to the defender of a gateway. The thresholds of Section 1 are consequences of that combination. The sufficient lead time L^* needs a lead time and a capacity limit, the threshold α^* needs false alarms that recur, and the pricing of a detector in the plane of its two error probabilities needs both.

3 Advance Disruption Information

We build the model of a warning's value to the defender in five parts: the gateway and its costs, the attack threat, the warning, the reading of the model at a real gateway, and the defender's optimal policy, from which the value of warning follows. We present the notation in Table 1 and show the sequence of events within a period in Figure 1.

3.1 The Gateway and Costs

Let time be divided into periods $n = 1, 2, \dots$, and we have in mind a period of one day. A facility consumes a commodity at a constant and known demand of d units per period, and is replenished through a single gateway which, in each period, is either *open* or *closed*. We measure inventory in periods of demand, so let $d = 1$ throughout. Let x_n be the *net inventory* at the start of period n , the

Table 1: Notation.

<i>The gateway (§3.1)</i>		<i>The warning and the belief (§3.3, §3.5)</i>	
x_n	net inventory at the start of period n	W_n	signal; $W_n = 1$ is an alarm
q_n	order in period n , $0 \leq q_n \leq C$	L	lead time of the detector, $1 \leq L \leq M$
C, ρ	surge capacity, $C > 1$, and utilization $\rho = 1/C$	α	false alarm probability
h	holding cost per unit per period	β	miss probability
b	backlog cost per unit per period	w	price per period of the detector
p	surge premium per unit beyond demand	$\mathcal{H}_n$	the defender's information through period n
$c(x, q)$	cost of a period, (3.3)	π_n	belief in period n , the posterior on K_n
B	backlog at a reopening, $[T - S_0]^+$ (§4)	$\pi_0, \bar{\pi}$	belief at a reopening, and the quiet belief
		ω	probability that an alarm is true, (5.1)
<i>The threat (§3.2)</i>		<i>Costs (§3.5, §4, §5)</i>	
K_n	countdown: 0 when quiet, then $M, \dots, 1$	$G(x)$	expected cost of a closure begun at x , (3.9)
M	campaign length, a known constant	$\tilde{G}(S)$	first-order closure and recovery cost, (5.2)
λ	probability a campaign is initiated	J, J_1	cost of an attack cycle, uninformed, informed
Ω	open spell, reopening to the next closure, (3.5)	$A(Q)$	expected cost change from a hedge Q , (5.4)
T	duration of a closure	$v(x, \pi)$	relative value function, (3.12)
Λ	attack cycle, reopening to reopening, $\Omega + T$	ϵ	critical ratio of the quiet inventory, (4.4)
<i>Inventory levels (§3.5, §4, §5)</i>		<i>Values (§3.5, §5)</i>	
S_0, S_A	quiet and alert inventory levels	g_i^*	minimum long-run average cost of defender i :
Q	hedge, $S_A - S_0$, in periods of demand		$i = 0$ uninformed, 1 informed, 2 clairvoyant
S_{lo}, S_{hi}	levels of the band, functions of belief π	V	value of warning, $g_0^* - g_1^*$
S^*	uninformed defender's optimal quiet level	$\bar{V}$	value of clairvoyance, $g_0^* - g_2^*$
S_0^*, S_A^*	informed defender's optimal levels	V_1	first-order value of warning, (5.7)
<i>Key values, marked with a star (§5)</i>			
L^*	the lead time that suffices; V_1 is constant beyond it (§5)		
α^*	first-order cry-wolf threshold, beyond which nothing is acquired on an alarm (§5)		

inventory on hand when $x_n \geq 0$ and a backlog of $-x_n$ when $x_n < 0$. The defender observes whether the gateway is open. If it is open, the defender orders $q_n \in [0, C]$, which arrives immediately, where $C > 1$ is the *surge capacity*, the most the gateway can move in one period. If it is closed, $q_n = 0$. Demand that cannot be served is backordered and served when the gateway can next deliver it. We therefore have

$$x_{n+1} = x_n + q_n - 1. \quad (3.1)$$

We can equivalently read a positive x_n as the number of periods of a closure the defender can cover. The surge capacity need not be an integer, and we are interested in a gateway that is busy in normal times. Its *capacity utilization* is

$$\rho = \frac{1}{C}, \quad (3.2)$$

the fraction of the gateway's capacity that everyday demand uses while the gateway is open, and a gateway with $C = 1.1$ runs at $\rho = 0.91$.

The defender incurs three costs. Inventory carried into the next period costs $h > 0$ per unit. A backlog costs $b > 0$ per unit for every period that it waits. Each unit ordered beyond current demand costs a surge premium $p \geq 0$, representing the overtime, the charter, or the priority handling that moves freight through the gateway faster than its normal pace. We define the cost of period n as

$$c(x_n, q_n) = p \cdot [q_n - 1]^+ + h \cdot [x_n + q_n - 1]^+ + b \cdot [-(x_n + q_n - 1)]^+, \quad (3.3)$$

with $[a]^+ = \max\{0, a\}$ for a real. Under the backlog assumption, every unit of demand is served sooner or later, so the acquisition cost of a unit is the same under every policy, and we set it to zero.

We make three assumptions about the gateway. First, demand is constant, as in the disruption models that begin with Parlar and Berkin (1991); Schmitt et al. (2010) is the extension to random demand. The uncertainty that governs the hedge is on the supply side, and constant demand lets each unit of inventory be followed separately; Section 6 discusses random demand. Second, unserved demand is backordered. Freight that cannot pass a closed gateway waits for it, and a need for fuel, water, or medical supplies that goes unmet today is still there tomorrow; the alternative, in which unserved demand is lost, is also discussed in Section 6. Third, an order arrives in the period it is placed. A period is long relative to the transit through the gateway, and the constraint that matters is the surge capacity because net inventory is capable of rising by at most $(C - 1)$ per period. So a hedge of Q periods of demand takes $\lceil Q/(C - 1) \rceil$ periods to acquire, and a backlog of the same size takes as long to clear. At $\rho = 0.91$, that is ten periods of surging for each period of cover or of backlog. The surge capacity thus plays the part of a lead time twice, before the closure and after it.

3.2 The Attack Threat

An attack is a campaign that takes a constant $M \geq 1$ periods to mount, from its initiation to the closure it causes. We speak of an attacker, but the model asks only for a threat that announces itself before it closes the gateway. A hurricane that forms and tracks toward a port mounts a campaign in this sense, and its lead time is the forecast's; a protest that gathers to blockade a road or a terminal mounts one too, and its lead time is the notice its organizers give. Section 2 cites the models that treat the hurricane. Let K_n be the *countdown* in period n , the number of periods that remain before the gateway closes. When no campaign is under way we write $K_n = 0$ and say that the gateway is *quiet*. In a quiet period the gateway stays quiet with probability $1 - \lambda$, and with probability $\lambda \in (0, 1)$ an attacker initiates a campaign. The countdown then takes the values $M, M - 1, \dots, 1$ in successive periods, and the gateway closes in the period after it reaches 1. If $K_n = k \geq 1$, the gateway is open in periods $n, \dots, n + k - 1$ and closes in period $n + k$. The value 0 is thus a label for the quiet state, and a campaign's countdown ends at 1. The countdown moves as follows:

$$\begin{aligned}
 P\{K_{n+1} = 0 \mid K_n = 0\} &= 1 - \lambda && \text{(the gateway is quiet),} \\
 P\{K_{n+1} = M \mid K_n = 0\} &= \lambda && \text{(attacker initiates a campaign),} \\
 P\{K_{n+1} = k - 1 \mid K_n = k\} &= 1, \quad k = 2, \dots, M && \text{(the campaign runs down),} \\
 P\{\text{closed in } n + 1 \mid K_n = 1\} &= 1 && \text{(the closure begins).}
 \end{aligned} \tag{3.4}$$

Once closed, the gateway remains closed for T periods, where T is a positive integer random variable, independent of all else, with $E[T^2] < \infty$. The probability that a gateway closure lasts at least j periods is $P\{T \geq j\}$. The gateway then reopens in the quiet state. Restricted to open periods, the countdown is a Markov chain on $\{0, 1, \dots, M\}$ with transition probabilities (3.4). Since the defender sees only the closure and, from Section 3.3, a signal emitted from the countdown, the countdown is the hidden state of a hidden Markov model, and the defender's problem of Section 3.5 is a partially observed Markov decision process. Let an *open spell* be the run of open periods from a reopening to the next closure, and let Ω be its length, a random variable. The open spell has two

parts: the quiet periods, whose number is geometric with mean $1/\lambda$, and then the M periods of the campaign that ends it. Let $\Lambda = \Omega + T$ be the *attack cycle*, the random time from one reopening to the next. The mean lengths of an open spell and of an attack cycle are

$$E[\Omega] = \frac{1}{\lambda} + M, \quad E[\Lambda] = E[\Omega] + E[T] = \frac{1}{\lambda} + M + E[T]. \quad (3.5)$$

We assume there is one campaign at a time, which is reasonable when attacks are rare. We also assume an attacker who does not respond to the defender's inventory, which is the subject of Section 7. We take the campaign length M to be a known constant. The defender never observes a campaign. It knows M in the way it knows λ and the distribution of T , as a parameter estimated from the record. M is the interval from an intruder's initial access to the attack's effect. Intrusion-lifecycle reports measure the *dwelt time*, the time from compromise to detection (Mandiant, 2025), which is $M - L$ when the report's compromise and detection mark the campaign's start and the detector's alarm; when a ransom note is the detection, the dwelt time runs to the effect and is M itself. One representative value is enough, because a campaign is invisible to every detector until its countdown reaches L . Until then, it looks to the defender exactly like a quiet gateway, so M enters the defender's problem only through the safe window after a reopening (Section 4) and the mean attack cycle (3.5).

Were M random and bounded, with $M \geq L$ for every campaign, the safe window would run to the shortest campaign rather than to M , and $E[M]$ would replace M in (3.5). From the longest campaign onward, the belief on the countdowns that a detector can see, and the probability that the gateway closes next period, are exactly what they are under a constant M , whatever its distribution. The thresholds of the sections that follow rest on that belief, and the first-order cost accounting of Section 4 would carry the shortest and the mean campaign length in place of M . The belief between the shortest and the longest campaign, however, is a transient that depends on the whole distribution. We do not analyze it, and we keep M constant.

The campaign is what makes a warning possible. A detector with lead time L presupposes an attack that is decided at least L periods before it lands. The campaign length M is how far ahead the attacker decides, and the lead time L is how far ahead the defender can see. Note that (3.5) does not contain L . We keep M and L separate so that detectors with different lead times face the same threat. A closure that arrived without any such run-up could not be foreseen, and there would be nothing for Bayes' rule to update. For the same reason the lead time cannot exceed the campaign length. A detector set to fire $L > M$ periods ahead is set at a countdown the campaign never reaches, so it never raises a true alarm, and a detector that watches the last $L > M$ periods of the campaign watches all M of them. A lead time longer than M is therefore at best a lead time of M , so we assume $L \leq M$.

3.3 The Warning

In each open period, the defender sees a signal $W_n \in \{0, 1\}$ before ordering, where $W_n = 1$ is an *alarm*. The sequence of signals is the advance disruption information, and Section 5 values it. We describe a detector by its *lead time* $L \in \{1, \dots, M\}$, its *false alarm probability* α , and its *miss probability* β . We study the *single warning*. The detector is a tripwire that the campaign crosses L periods before the closure, so the campaign is visible at countdown L only. A detector that reports in each of the campaign's last L periods, the *persistent warning*, accumulates evidence over a stream of signals and calls for a belief-state analysis that we leave to a separate study. We assume that the warning signals W_n are independent given the countdowns, with

$$P\{W_n = 1 \mid K_n = k\} = \begin{cases} 1 - \beta, & k = L, \\ \alpha, & k \in \{0, 1, \dots, M\} \setminus \{L\}. \end{cases} \quad (3.6)$$

The first line is the detector with the campaign in view. It raises an alarm unless it misses, which it does with probability β . The second line is the detector with nothing in view, and that is so in three situations: the gateway is quiet ($k = 0$); a campaign is under way but has not yet come

within the detector's reach ($k > L$); and the campaign has already passed it ($0 < k < L$). In each of these the alarm is noise, raised with the same probability α whether or not an attack is coming, so a signal can only tell a visible campaign from everything else. The false alarm probability is therefore the probability of an alarm in any period in which the campaign is not visible, of which the quiet periods are the most numerous. The earliest true alarm comes when the countdown is L , so the defender has L periods before the gateway closes, and at the surge rate of Section 3.1 it can raise its net inventory by at most $(C - 1)L$ in them. That is why a busy gateway needs more lead time to acquire a given hedge.

The three probabilities do not refer to the same exposure. The attack probability λ and the false alarm probability α are probabilities per period, so their size depends on the length of the period. A detector that raises one false alarm a month has $\alpha \approx 1/30$ when the period is a day. The miss probability β is a probability per look at a visible campaign, and the tripwire has one look, so β is the probability that the attack arrives without a true alarm. In the terms of a hypothesis test whose null hypothesis is that no campaign is in view, α is the Type I error of a test repeated every open period and β the Type II error of its one decisive look.

We assume $\alpha + \beta \leq 1$, which says that an alarm is at least as likely when the campaign is visible as when it is not; a detector that violates it satisfies it once its silence is read as its alarm. The assumption therefore costs nothing, and it fixes what an alarm means: evidence of a campaign, with likelihood ratio $(1 - \beta)/\alpha \geq 1$, which every threshold of Section 5 uses. It marks the region on which the value of warning is monotone in α and β (Proposition 4). On the line $\alpha + \beta = 1$, the value vanishes, and beyond it an alarm is evidence that no campaign is in view and the value can rise again.

3.4 Modeling a Gateway

The model is abstract, and Table 2 illustrates possible inputs for a port or an airfield. Three things about the table deserve notice. First, the three costs are the defender's own, and none of the quantities requires the defender to know anything about the attacker's intentions. Second, the three

numbers that describe the detector are the ones its vendor can state: the false alarm probability is from system tuning, the miss probability is from its sensitivity, and the lead time is the interval between the alert and the effect. The false alarm probability counts periods, not alerts. It is the chance of at least one alarm in a period, after whatever rule turns raw alerts into one signal per period, so it is not a per-item false-positive rate or one minus a precision. The sensitivity and the lead time may move together on one detection threshold, so the two need not be separately available; our comparisons hold one fixed while the other moves. Third, the closure duration is the least well known of the table's quantities and the most consequential, and its variability sets the hedge. We show this in Section 4.

3.5 The Defender's Optimal Policy and the Value of Warning

Because the countdown is hidden, the defender acts on a belief about it. Let $\mathcal{H}_n$ be the σ -field generated by the gateway's status and the signals through period n (Krishnamurthy, 2016), the defender's information when it places order q_n . Each period yields one binary status and one binary signal, so $\mathcal{H}_n$ is generated by finitely many events. Ordering adds nothing to it, because an order is a function of what the defender had already seen. The information records only what the gateway and the detector have shown, and the defender's experience and judgment enter the model through its parameters. The defender's *belief* in period n is the posterior distribution of the countdown given that information, π_n with $\pi_n(k) = P\{K_n = k \mid \mathcal{H}_n\}$ for $k = 0, \dots, M$. In each open period, the defender first learns that the gateway is open, then sees the signal, and then orders inventory, as Figure 1 shows. The belief that carries into the order is π_n , which Bayes' theorem forms from the belief before the signal and the signal just seen. When the period brings an alarm,

$$P\{K_n = k \mid \mathcal{H}_n\} = \frac{P\{W_n = 1 \mid K_n = k\} P\{K_n = k \mid \mathcal{H}_{n-1}, \text{ open in } n\}}{\sum_j P\{W_n = 1 \mid K_n = j\} P\{K_n = j \mid \mathcal{H}_{n-1}, \text{ open in } n\}}, \quad (3.7)$$

with $P\{W_n = 1 \mid K_n = k\}$ from (3.6). When the period is silent, each of those probabilities is replaced by its complement. The belief before the signal, on the right of (3.7), is not last period's

Table 2: The objects of the model at a seaport or an airfield.

Object	In the model	At a gateway, with examples
Period and unit	one period; demand $d = 1$ per period	a day, and a day's requirement of what the gateway carries: a day's lift for a relief operation, a day's fuel for a base, a day's containers
Surge capacity C , utilization $\rho = 1/C$	the most the gateway moves in a period; the fraction of it that everyday traffic uses	the maximum daily throughput of a runway, a berth, or a rail yard; a gateway that runs near capacity in normal times is busy in our sense, and its surplus $C - 1$ is what a hedge is acquired from and a backlog cleared with
Campaign length M	periods from the initiation of an attack to the closure	the interval from an intruder's initial access to its kinetic effect, roughly the reported dwell time plus the lead time; it belongs to the attacker
Closure duration T	periods the gateway is closed; its mean sets the recovery, and its tail sets the hedge (Section 4)	the time to restore operations or the systems behind them, including forensics and rebuilding: about ten days at Maersk, about three weeks at Seattle for the airport's systems, five hours at LOT; records of comparable closures would estimate its mean and tail
Detector (L, α, β)	alarm probability $1 - \beta$ when the campaign is visible, and probability α in any other open period; the campaign comes into view L periods before closure	an intrusion detection system on the terminal's operational technology network (Adams et al., 2021); a monitor for interference with satellite navigation, of the kind that made a Baltic airport's satellite-based approaches unusable and suspended an airline's service to it for a month (C4ADS (Center for Advanced Defense Studies), 2019; ERR News, 2024); a monitor on the navigation systems of ships approaching a port, against the cyber threats that Karahalios (2025) assesses
Holding cost h	per unit carried into the next period	the cost of carrying a day's supply at the facility: storage, spoilage, capital
Backlog cost b	per unit of demand, per period it waits	a day's requirement unmet for a day, called an unmet unit-period; at a terminal, demurrage and stranded freight
Surge premium p	per unit ordered beyond current demand	overtime, extra sorties, or chartered lift through the gateway




Figure 1: The sequence of events in a period. The shaded lane is hidden from the defender, and the unshaded lane is what it sees and does. A closed period has no signal and no order.

belief, because the countdown has moved one step and the gateway is open. The gateway closes in period $n + 1$ with probability $P\{K_n = 1 \mid \mathcal{H}_n\}$, and if it stays open, the transition probabilities of (3.4) carry the belief forward,

$$P\{K_{n+1} = k \mid \mathcal{H}_n, \text{ open in } n + 1\} = \frac{\sum_{j \neq 1} P\{K_n = j \mid \mathcal{H}_n\} P\{K_{n+1} = k \mid K_n = j\}}{1 - P\{K_n = 1 \mid \mathcal{H}_n\}}. \quad (3.8)$$

At a reopening, the belief is π_0 , which puts all of its mass on the quiet state. A belief could, in general, lie anywhere in the set of probability distributions on $\{0, 1, \dots, M\}$, a continuum. Here it does not. The countdown is deterministic once a campaign starts, and a signal older than L periods carries no information.

Lemma 1 (The information state is finite). *The belief in period t of an open spell depends on the signal history only through t and the last L signals, and for $t \geq M + 1$ through the last L signals alone. Only finitely many beliefs are therefore reachable.*

Section 6 therefore solves the defender's problem, on the finite inventory range assumed below, with no approximation of the belief.

Suppose the gateway closes with net inventory x . After period j of the closure, the net inventory is $x - j$, and at the reopening it is $x - T$. Let $G(x)$ be the expected cost incurred over the closure.

Lemma 2 (The cost of a closure). *For every x ,*

$$G(x) = \sum_{j \geq 1} \mathbb{P}\{T \geq j\} \left(h [x - j]^+ + b [j - x]^+ \right), \quad (3.9)$$

G is convex and linear between consecutive integers, and for integer $x \geq 0$,

$$G(x + 1) - G(x) = h \mathbb{E}[\min\{T, x\}] - b \mathbb{E}[T - x]^+ = (h + b) \mathbb{E}[\min\{T, x\}] - b \mathbb{E}[T]. \quad (3.10)$$

Therefore G is minimized at the smallest integer x for which $\mathbb{E}[\min\{T, x\}] \geq \frac{b}{b+h} \mathbb{E}[T]$.

This is a newsvendor result with one change. The critical ratio $b/(b + h)$ is the familiar fractile taken in the *equilibrium distribution* of the duration, $\mathbb{E}[\min\{T, x\}]/\mathbb{E}[T]$, known from renewal theory as the distribution of the age, counting the period in progress, of a closure observed at a random moment (Ross, 1996). The change arises because both costs are incurred per period, not per unit. One more unit is held in every period of the closure until it is consumed, and then spares one unit of backlog in each of the $[T - x]^+$ periods that remain. Long closures therefore weigh on the hedge through the second moment of T , which is why we assume it is finite. The level that minimizes G is the level worth holding for the closure alone; the quiet inventory of Section 4 pays to hold it through the wait for the attack as well, and sits lower. In the backlog case, the closure is only part of what an attack costs. At reopening, a backlog of B takes $\lceil B/(C - 1) \rceil$ periods to clear.

Proposition 3 (Stability). *If $(C - 1) \mathbb{E}[\Omega] < \mathbb{E}[T]$, then under every policy, the backlog grows without bound and the long-run average cost is infinite. If*

$$(C - 1) \mathbb{E}[\Omega] > \mathbb{E}[T], \quad (3.11)$$

then there is a policy with finite long-run average cost.

We assume condition (3.11), which states that the surplus capacity, $C - 1$ per period over the mean length $\mathbb{E}[\Omega]$ of an open spell, must be greater than the demand that accumulates over a

closure, and we leave the boundary case of equality aside. The backlog at successive closures obeys the queueing recursion of Lindley (1952), and (3.11) is its drift condition. This just affirms that no warning can restore system stability to a gateway that lacks sufficient capacity.

As a partially observed Markov decision process, the defender's problem has the countdown K_n for its hidden state, moving by (3.4); its observations are the gateway's status and the signal W_n , emitted by (3.6); its action is the order q_n ; and its cost is (3.3) in each open period and G over each closure. We follow the standard treatment and replace the hidden state by the belief (Smallwood and Sondik, 1973; Krishnamurthy, 2016). Given $\mathcal{H}_n$, the distribution of everything the defender will see and incur depends on the past only through x_n and π_n . The countdown is Markov and π_n is the conditional distribution of its present value. The signals are independent given the countdown, and the inventory moves by (3.1). The pair (x_n, π_n) is therefore a state in the Markov sense, and a policy that conditions on it alone gives up nothing. Decisions are made in open periods only. A closure is a transition of random length with no decision in it and expected cost G by Lemma 2.

We write the state as (x, π) , with the subscript n suppressed, and call it the *information state*, since it carries everything the defender knows. A policy chooses the order from it, we evaluate the policy by its long-run average cost, and we assume that π is a belief that some history produces. Two restrictions make the state space finite, and we assume both. First, we confine x to a finite range, whose edge the process reaches with negligible probability (Section 6 reports how little). Second, we restrict orders to a lattice of small step. On the states it can reach, the model is communicating. A closure arrives with probability one under any order rule and returns the belief to π_0 , and a quiet stretch of any length has positive probability, during which the defender can steer the inventory to any level in its range. The minimum average cost g^* therefore does not depend on the initial state, and it and the optimal policy are those of a finite-state Markov decision process with the average-cost criterion (Puterman, 1994, Sections 8.3.3 and 9.5). With a relative value function v ,

they satisfy the optimality equation for the long-run average cost

$$\begin{aligned}
v(x, \pi) = \min_{0 \leq q \leq C} \bigg\{ & c(x, q) - g^* \\
& + \pi(1) \cdot \left(G(x + q - 1) - g^* \cdot \mathbb{E}[T] + \mathbb{E}[v((x + q - 1) - T, \pi_0)] \right) \\
& + (1 - \pi(1)) \cdot \mathbb{E}[v(x + q - 1, \pi^+) \mid \pi] \bigg\}, \quad (3.12)
\end{aligned}$$

where π^+ is the belief of the next period, which (3.7) and (3.8) produce from π and the next signal, and the expectation on the last line is over that signal. The first line is the period's cost net of the average g^* . The second is the closure that may follow, its T periods each net of g^* , which is the semi-Markov form of the equation (Puterman, 1994, Section 11.4), with the recovery folded into $v(x + q - 1 - T, \pi_0)$. The third is the open period that may follow instead. Equation (3.12) is written before truncation, with each closure a single transition. Section 6 solves its daily form, in which each closed day is a period of its own with no order, by relative value iteration on the finite range described there.

We compare three defenders, whose costs give a worst case (highest cost), a middle case, and a best case (lowest cost). They are: (0) the *uninformed* defender that has no detector and sees only whether the gateway is open; (1) the *informed* defender that sees the signals of a detector (L, α, β) ; and (2) the *clairvoyant* defender that can see the attacker's countdown. Let g_0^* , g_1^* , and g_2^* be their minimum average costs per period, respectively. We define the *value of warning* and its ceiling as

$$V(L, \alpha, \beta) = g_0^* - g_1^*, \quad \bar{V} = g_0^* - g_2^*. \quad (3.13)$$

V is what the advance disruption information is worth to the defender. We compare V directly against the price w per period of the detector, as in (1.1). Four properties hold for any duration distribution and any costs.

Proposition 4 (Properties of V). *If (3.11) holds, $\alpha + \beta \leq 1$, and $1 \leq L \leq M$, then:*

$$(a) \ 0 \leq V(L, \alpha, \beta) \leq V(L, 0, 0) \leq \bar{V};$$

(b) $V(L, \alpha, \beta) = 0$ when $\alpha + \beta = 1$;

(c) V is non-increasing in α and non-increasing in β ;

(d) V is non-decreasing in L .

Property (a) bounds the value of any detector between nothing and the value of clairvoyance. The informed defender can ignore the signals, and the clairvoyant defender, who sees the countdown, can draw the signals of any detector. Property (b) says that a detector whose alarm is as likely with the campaign in view as without carries no information or value. The signals are then independent of the countdown, Bayes' theorem leaves every belief unchanged, and the informed defender's problem is the uninformed one. Property (c) is Blackwell's principle of informativeness. A better detector can imitate a worse one by corrupting (garbling) its own signals at random (Blackwell, 1953). Property (d) says that a detector with lead time $L + 1$ can imitate one with lead time L by delaying each signal by a period.

The form of the optimal policy follows from a kink in the ordering cost, which the surge premium puts at $q = 1$. The unit that replaces current demand moves at no premium, and every unit beyond it costs p , as Figure 2(a) shows.

Theorem 5 (The optimal policy is a band). *Under the discounted criterion with discount factor $0 \leq \gamma < 1$ per period, if the inventory is unbounded and $p(1 - \gamma) < b$, then for every reachable belief π there are levels $S_{lo}(\pi) \leq S_{hi}(\pi)$, which depend on γ , such that the optimal order in state (x, π) is*

$$q^*(x, \pi) = \begin{cases} \min\{C, S_{lo}(\pi) - x + 1\}, & x < S_{lo}(\pi), \\ 1, & S_{lo}(\pi) \leq x \leq S_{hi}(\pi), \\ \max\{0, S_{hi}(\pi) - x + 1\}, & x > S_{hi}(\pi). \end{cases}$$

If $p = 0$ the two levels coincide, and the policy is a base-stock policy, limited by the surge capacity.

Theorem 5 is easiest to read one belief at a time. At a belief π , where the net inventory settles depends on the direction from which it arrives. From below, the defender surges only up to the lower




Figure 2: The ordering cost kink (a), and a schematic of the net inventory under the optimal policy of Theorem 5 through a false alarm (b). The shading is the band $[S_{l_0}(\pi_n), S_{hi}(\pi_n)]$ at the belief of each phase, and q is the order.

level $S_{l_0}(\pi)$, because a unit above it is not worth the premium. From above, it lets consumption draw the inventory down only to the upper level $S_{hi}(\pi)$, because a unit already on hand would cost the premium to acquire again. Between the levels, in the *hold region*, it orders $q = 1$, replacing demand and nothing more, and the inventory stays where it is. The premium opens this gap. The belief moves the band, and that is how the policy of Section 5 arises. From period $M + 1$ of an open spell on, the belief stays fixed for as long as the last L periods bring no alarm (Lemma 1), and we call it the *quiet belief*, $\bar{\pi}$. Its lower level is the quiet inventory, $S_0 = S_{l_0}(\bar{\pi})$, which the defender establishes after a reopening. For the uninformed defender, which sees no alarm, every period from $M + 1$ on is quiet, and the upper level never binds. An alarm raises the belief that a closure is near, and the band with it. The lower level rises, and Section 5 approximates it by a single alert inventory S_A , toward which the defender surges. Thus S_0 and S_A are both lower levels, at different beliefs.

Figure 2(b) follows the net inventory through a false alarm. When the alarm proves false, the belief returns exactly to $\bar{\pi}$. A false alarm is revealed when the gateway is still open L periods after it, which rules out the one countdown a true alarm implied. At that moment the alarm leaves the window of the last L signals in Lemma 1, and if no new alarm has come, those signals are silent again. The band falls back with the belief, and the inventory, now above it, comes down only to $S_{hi}(\bar{\pi})$ and stays there. The part of the hedge $S_A - S_0$ below that level waits, at no further premium, for the next alarm or the next closure. If another alarm comes first, the surge starts from $S_{hi}(\bar{\pi})$ rather than S_0 , a head start. If a closure comes first and consumes the inventory, the defender restores it after the reopening from below to S_0 and no further, as Figure 3 shows after a true alarm.

The inventory thus returns to S_0 only from below. The first-order analysis of Section 5 sets the hold region aside and lets consumption draw the hedge down to S_0 after every false alarm, as Figure 3 also shows, and Section 6 measures what the hold region adds.

We base the proof on the convexity of the cost-to-go in the net inventory, which the linear dynamics and the convex costs give by induction, and on the observation that the lower level minimizes the cost-to-go plus the premium while the upper level minimizes the cost-to-go alone. Without the premium, the band is the modified base-stock policy of Federgruen and Zipkin (1986a,b) for a source of limited capacity, and the proof follows their convexity argument. We state the theorem under the discounted criterion with the inventory unbounded because there the convexity argument is exact. For discount factor γ , the condition $p(1 - \gamma) < b$ says that the premium is less than the discounted cost of carrying a unit of backlog forever, which is what guarantees a finite lower level. When the premium exceeds $b/(1 - \gamma)$, no optimal order exceeds one unit. The defender never surges to clear a backlog, the lower level is $-\infty$, and the policy has one level. As γ approaches 1, the condition holds for every finite premium, because a backlog that is never cleared costs b in every period while the premium is incurred once. We do not prove the band under the average cost, which is the criterion of this paper because the value of warning is a cost per period, to be compared with a price per period. Carrying a structural result from the discounted to the average cost needs conditions (Puterman, 1994, Section 8.11) that we do not verify for an unbounded inventory, and on the finite range on which we compute, the edges can break the convexity argument. The discounted theorem is our reason to expect a band under the average cost.

We make claims of three kinds. First, the finite information state, the cost of a closure, stability, and the information comparisons (Lemmas 1 and 2, Propositions 3 and 4) hold for the model as stated. Second, the band is proved under the discounted criterion, and Section 6 finds it under the average cost in every reachable state of every instance solved, on a finite range. Third, the tiers, the lead time that suffices, and the cry-wolf threshold of Sections 4 and 5 are first-order approximations, which Section 6 measures against the exact solution.

4 The Uninformed Defender and Its Cost

In this section, we solve the defender's problem without a detector. The result is the benchmark, the cost g_0^* , from which we measure every value of warning. We first show that the uninformed defender's problem is simpler than the general one. This is because, after a reopening, a closure is impossible for a stretch of $M + 1$ periods, which we call the *safe window*, and after that the wait for one is memoryless. We then compute the cost and the quiet inventory S^* that minimizes it. We close with how the tail of T shapes S^* .

4.1 The Safe Window

Let the integer $t \geq 1$ count the periods since the last reopening, so that $t = 1$ is the first open period, in which the gateway is quiet. By (3.4), a campaign can be initiated no sooner than period 2, and its countdown reaches 1 no sooner than period $M + 1$. A closure can begin in period $t + 1$ only if $K_t = 1$. The first period in which a closure can begin is therefore $t = M + 2$. We describe what happens after that in the following lemma.

Lemma 6 (The safe window). *The gateway is open with certainty in periods $1, \dots, M + 1$ of an open spell. For every $t \geq M + 1$, conditional on the gateway being open in period t , the uninformed defender's belief about the countdown is*

$$\begin{aligned} P\{K_t = 0 \mid \text{open in } t\} &= (1 - \lambda)^M, \\ P\{K_t = k \mid \text{open in } t\} &= \lambda(1 - \lambda)^{k-1}, \quad k = 1, \dots, M, \end{aligned} \tag{4.1}$$

and the probability that the gateway closes in period $t + 1$ is exactly λ .

Lemma 6 divides an open spell into a safe window and the risky periods that follow, by what the defender knows. Section 3.2 divided the same spell into quiet periods and a campaign, by what the attacker does. In the safe window, the defender has $M + 1$ periods in which to clear any backlog from the last closure and to restore the quiet inventory in preparation for the next one. From period $M + 1$ on, the closure time is memoryless with hazard λ , so the number of risky periods

is geometric with mean $1/\lambda$. The two parts share period $M + 1$, the last that is open and the first after which the gateway may close. So, less that overlap, their lengths add to Ω and their means to $E[\Omega] = 1/\lambda + M$, as in (3.5). The distribution (4.1) is the belief before the signal, the prior that we update on an alarm in Section 5.

4.2 The Quiet Inventory

From here on, in this section and the next, we trade exactness for closed forms. We have none for the exact problem, and the approximations below give the quiet inventory, the hedge, and the value of warning by minimizing explicit costs over one or two levels. Section 6 measures them against the exact solution.

Because the belief is constant from period $M + 1$ on, the optimal order in those periods depends on the net inventory alone, and the defender orders up to a level, the form that Theorem 5 gives under the discounted criterion and that Section 6 finds under the average cost. Let S_0 be this level and call it the *quiet inventory*. In the cycle we describe, the theorem's upper level never binds, since the net inventory is restored to the lower level in the safe window and only a closure changes it after that. In the safe window, the defender clears the backlog as fast as it can, since every period of delay costs b per unit waiting. It restores the quiet inventory as late as it can, since holding during the safe window protects against nothing. Both moves run at the surge rate. The defender orders C , demand takes 1, and the net inventory rises by $C - 1$ per period, the most it can. We call this the *quiet inventory policy* with level S_0 , and we compute its cost per attack cycle in what we call the *first-order approximation*, which makes three simplifications. First, we treat the ramps and the drawdowns as continuous and the levels as real numbers. Second, we neglect the closures that begin while the ramp of a long recovery is still unfinished. Third, we neglect the units that survive a short closure and are held through the safe window.

Proposition 7 (Cost of a quiet inventory). *Let $B = [T - S_0]^+$ be the backlog at a reopening and $U = \min\{T, S_0\}$. For a level $S_0 \geq 0$, the total cost of one attack cycle under the quiet inventory*

policy is, in the first-order approximation,

$$J(S_0) = \frac{hS_0}{\lambda} + h E\left[S_0 U - \frac{U^2}{2}\right] + \frac{hS_0^2}{2(C-1)} + \frac{bC}{2(C-1)} E[B^2] + p E[T], \quad (4.2)$$

and the average cost is $J(S_0)$ divided by the mean length of the attack cycle (3.5).

The five terms of (4.2) are the holding cost of the quiet inventory while the defender waits for the attack, its holding cost while the closure consumes it, the holding cost of acquiring it, the cost of the backlog, and the surge premium. The random variables U and B compare a time with an inventory, which they can because inventory is measured in periods of demand. An inventory of S_0 units covers S_0 periods of a closure. So $U = \min\{T, S_0\}$ is the part of the closure the quiet inventory covers, during which the inventory falls by one unit a period from S_0 , and $B = [T - S_0]^+$ is the part it does not cover, in which the backlog grows by one unit a period. Two of the terms deserve comment. First, the premium is $p E[T]$ whatever the level, because every unit consumed during a closure must be replaced by a unit surged beyond current demand, and no unit is surged for any other reason. The premium therefore does not enter the choice of S_0 , but it will enter the choice of a hedge on an alarm, since a hedge acquired on a false alarm is drawn down without being consumed by a closure. Second, the backlog costs b per unit-period both while it accumulates during the closure and while it is cleared during the recovery. The closure costs $b E[B^2]/2$ and the recovery $b E[B^2]/(2(C-1))$, together $b E[B^2]/2$ multiplied by $C/(C-1)$, which is $1/(1-\rho)$. A gateway at $\rho = 0.5$ incurs its backlog cost twice, once as it forms and once as it clears, and a gateway at $\rho = 0.91$ incurs it eleven times. The recovery and the restoration are where utilization enters the model.

Proposition 8 (The quiet inventory is a fractile). *The cost (4.2) is convex in S_0 , and the optimal quiet inventory, which we write S^* , is, in the first-order approximation, the smallest $S_0 \geq 0$ for which*

$$\frac{bC}{C-1} E[T - S_0]^+ \leq h \left(\frac{1}{\lambda} + E[\min\{T, S_0\}] + \frac{S_0}{C-1} \right). \quad (4.3)$$

In particular, $S^* > 0$ if and only if $\lambda b C E[T] > h(C-1)$, and (4.3) then holds with equality.

The left side of (4.3) is what a marginal unit of inventory saves. It removes one unit from the backlog in each period of the closure beyond S_0 , of which there are $E[T - S_0]^+$ on average, and it saves the recovery of that unit as well. The right side is what the marginal unit costs to hold: for the mean number of periods $1/\lambda$ until the attack, for the periods of the closure it survives, and for the periods of the ramp that restores it. Since $E[T - S_0]^+ = E[T] - E[\min\{T, S_0\}]$, condition (4.3) is a fractile condition in the equilibrium distribution of Lemma 2, extended linearly between integers, with critical ratio ϵ :

$$\epsilon = \frac{\frac{bC}{C-1} - \frac{h}{\lambda E[T]} - \frac{h S_0}{(C-1) E[T]}}{\frac{bC}{C-1} + h} \leq \frac{E[\min\{T, S_0\}]}{E[T]}. \quad (4.4)$$

The ratio depends on the level itself, so the condition is implicit. It is the newsvendor's $b/(b+h)$ of Lemma 2 with two changes. The recovery multiplies the backlog cost by $C/(C-1)$ and raises the ratio. The holding cost of the wait and of the ramp counts against it and pulls it well down. At the instance of Section 6, where $\rho = 0.67$, the recovery raises the ratio from 0.91 to 0.97, and the wait and the ramp pull it down to 0.52, so the quiet inventory is 2.6 days where the closure alone would hold 9. When attacks are rare, the holding cost h/λ dominates the right side, and (4.3) becomes

$$E[T - S_0]^+ \approx \frac{h}{\lambda b} \cdot \frac{C-1}{C} = \frac{h}{\lambda b} (1 - \rho). \quad (4.5)$$

The quiet inventory is therefore the level at which the expected uncovered length of a closure, $E[T - S_0]^+$, falls to the ratio of the holding cost to the attack probability times the backlog cost, scaled by $1 - \rho$, the unused share of capacity.

4.3 Light and Heavy Tails

From (4.5), the quiet inventory is a function of the tail of the duration, and the tail's shape decides how fast the inventory grows as attacks become more frequent or more costly.

Corollary 9 (Distribution tails). *For a geometrically distributed duration, in which a closure ends in each period with probability r , so that $P\{T \geq j\} = (1-r)^{j-1}$ for $j \geq 1$ and $E[T] = 1/r$, we have $E[T - S]^+ \approx (1-r)^S/r$, with equality at integer levels, and*

$$S^* \approx \frac{\log(\lambda \cdot b \cdot C / (r \cdot h \cdot (C-1)))}{\log(1/(1-r))} = \frac{\log(\lambda \cdot b / (r \cdot h \cdot (1-\rho)))}{\log(1/(1-r))},$$

or zero when the right side is negative. The level grows with the logarithm of $\lambda \cdot b \cdot C / (h \cdot (C-1)) = \lambda \cdot b / (h \cdot (1-\rho))$.

Under a geometric duration, a tenfold rise in the cost of an attack adds a fixed number of periods to the quiet inventory. Saghafian and Van Oyen (2016) find the same form, a base-stock level logarithmic in the ratio of backlog cost to holding cost, under geometric disruptions and exponential demand. Under a heavier tail, the inventory grows faster, as a power of that ratio when the duration has a Pareto distribution. In Section 6, we examine a Pareto-distributed duration, because closure durations can vary widely. For example, recovery from a simple mechanical failure ends when the part is replaced. Recovery from a complex attack includes forensics, the rebuilding of systems and infrastructure, and the restoration of trust in them.

Two conditions on utilization deserve to be stated here. First, the closed-form approximations of this section and the next assume that the recovery and the restoration after a typical closure, $E[T]/(C-1)$ periods in all, fit inside the safe window of $M+1$ periods; beyond that, closures overlap recoveries and the accounting no longer holds. We discuss this in Section 6. Second, a warning relieves the ramp only within its lead time, since the hedge it allows cannot exceed $(C-1)L$, so in the instances of Section 6 the value of a warning rises with utilization only until its lead time binds, and falls beyond that; what keeps rising is the lead time that still adds value, within the limits that the stability condition and the campaign length set on C and L .




Figure 3: The sequence of events under a single warning, in the first-order policy of this section, with the ramps drawn as fluid paths, for an example $L = 3$, $M = 4$, $C = 3$, and $Q = 4$.

5 The Value of a Single Warning, the Tripwire

In this section, we examine the case where the detector has one look at the campaign, when its countdown reaches L . As throughout, the defender minimizes its long-run average cost. Under that criterion, Section 6 finds the optimal policy to be a band, the form that Theorem 5 proves under discounting, and Figure 3 illustrates the first-order policy that approximates it over a false and a true alarm. We follow each marginal unit of inventory through a single alarm, which yields the value of the warning, the lead time that suffices, the cry-wolf threshold, and the price of a detector. All of these are closed-form approximations, first order in the sense of Section 4.2, with each alarm treated as isolated. We measure each of these quantities against the exact solution in Section 6.

In Figure 3, the defender holds the quiet inventory S_0 until an alarm, and then acquires the *hedge* Q , the amount above the quiet inventory, surging at $C - 1$ units per period to the *alert inventory* $S_A = S_0 + Q$. An alarm in period n informs the order q_n of that period, which is reflected in x_{n+1} . With $L = 3$ and a hedge that takes two periods to acquire, the defender waits one period, which saves a period of holding cost, and surges in the next two. A false alarm is revealed when the lead time expires with the gateway open, and the hedge is drawn down by consumption, the defender ordering nothing until the inventory falls back to S_0 .

On a true alarm, the defender surges to S_A , and at the end of the lead time the gateway closes. The inventory S_A covers the first periods of the closure, a backlog accumulates over the rest, and after the reopening the defender surges again to clear it. The reopening marks the start of the safe window, and the recovery and the restoration of S_0 fit inside it. The restoration comes as late as the window allows. No closure can begin inside it, so holding there protects against nothing. Both drawdowns run at the demand rate of one unit per period, since nothing arrives in either, and the closure's runs on into backlog. At a surge capacity near 1, both ramps are far shallower than the drawdowns.

5.1 The Hedge on an Alarm

Consider the defender at the belief (4.1) of a mature open spell, holding the quiet inventory S_0 when an alarm arrives. Let $\pi(L) = \lambda(1 - \lambda)^{L-1}$ be the prior probability that the countdown stands at L . From Bayes' theorem, the alarm is true with probability

$$\omega = \frac{\pi(L) \cdot (1 - \beta)}{\pi(L) \cdot (1 - \beta) + (1 - \pi(L)) \cdot \alpha}. \quad (5.1)$$

When attacks are rare, ω is small unless false alarms are nearly as rare. The defender decides how many extra units Q to acquire before the lead time expires. It acquires them just in time, at $C - 1$ per period over the last $Q/(C - 1)$ periods, because in this accounting, in which a false alarm is noise and no closure is nearer than the tripwire, holding them earlier protects against nothing. If the alarm is true, the closure begins with inventory S_A . If it is false, the defender draws the extra units down by ordering nothing for Q periods, and those units were surged for nothing. They replace no demand lost to a closure, so their premium pQ is incurred over and above the $p E[T]$ per attack cycle of Section 4.

Let $\tilde{G}(S)$ approximate the cost of a closure and its recovery starting with inventory S , the closure cost of Lemma 2 plus the recovery cost that follows. In the first-order approximation of

Section 4.2,

$$\tilde{G}(S) = h \mathbb{E} \left[SU - \frac{U^2}{2} \right] + \frac{bC}{2(C-1)} \mathbb{E}[B^2], \quad (5.2)$$

with $U = \min\{T, S\}$ and $B = [T - S]^+$, and we write the optimality conditions with derivatives.

The derivative

$$\tilde{G}'(S) = h \mathbb{E}[\min\{T, S\}] - \frac{bC}{C-1} \mathbb{E}[T - S]^+ \quad (5.3)$$

is the right side of (3.10) evaluated at a real level, with the recovery's share of the backlog cost, $b/(C-1)$ per unit-period, added to b . $\tilde{G}$ is the closure and recovery part of (4.2), which is therefore $hS_0/\lambda + hS_0^2/(2(C-1)) + \tilde{G}(S_0) + p \mathbb{E}[T]$. Acquiring a hedge $Q \in [0, (C-1)L]$ on the alarm then changes the expected cost by

$$A(Q) = \frac{hQ^2}{2(C-1)} + (1-\omega) \left(pQ + \frac{hQ^2}{2} \right) + \omega \left(\tilde{G}(S_0 + Q) - \tilde{G}(S_0) \right). \quad (5.4)$$

Proposition 10 (The hedge is a sum of cost-loss decisions). *A is convex, and the optimal hedge Q^* is the largest Q in $[0, (C-1)L]$ for which*

$$\omega \cdot \frac{b \cdot C}{C-1} \mathbb{E}[T - S_0 - Q]^+ \geq (1-\omega)(p + hQ) + \frac{hQ}{C-1} + \omega \cdot h \cdot \mathbb{E}[\min\{T, S_0 + Q\}], \quad (5.5)$$

and $Q^* = 0$ when (5.5) holds for no $Q > 0$.

Observe that (5.5) is the cost-loss rule of Murphy (1977), applied at the margin. A unit is acquired when the probability of the event exceeds the ratio of its cost to its loss. The loss is the backlog it averts, amplified by the factor $1/(1-\rho)$, and the cost is its premium and its holding, incurred whether or not the alarm is true. The cost of the marginal unit rises with Q while its loss falls, so the hedge grows until that ratio reaches ω .

5.2 Three Tiers and the Value of Warning

Over an attack cycle, the campaign is seen with probability $1 - \beta$ and missed with probability β , and false alarms arrive in the open periods in which the countdown is not L . We let an informed defender jointly decide the quiet inventory S_0 and the hedge Q . As a first-order approximation, the cost of an attack cycle under quiet inventory S_0 and alert inventory $S_A = S_0 + Q$ is

$$\begin{aligned}
 J_1(S_0, S_A) = & \frac{hS_0}{\lambda} + \frac{hS_0^2}{2(C-1)} + pE[T] \\
 & + \beta \tilde{G}(S_0) + (1 - \beta) \left(\frac{hQ^2}{2(C-1)} + \tilde{G}(S_A) \right) \\
 & + \alpha (E[\Omega] - 1) \left(pQ + \frac{hQ^2C}{2(C-1)} \right).
 \end{aligned} \tag{5.6}$$

The first line is what the alarm does not touch, carried over from (4.2): the holding cost of the quiet inventory while the defender waits for the attack, the holding cost of acquiring it, and the premium on the units a closure consumes, which must be replaced whatever the levels. The second line is the closure and its recovery. With probability β the campaign is missed and the closure begins at S_0 ; with probability $1 - \beta$ it is seen, the hedge is acquired just in time at a holding cost of $hQ^2/(2(C-1))$ over the ramp, and the closure begins at S_A . The third line is the false alarms. They arrive with probability α in each open period of the cycle in which the countdown is not L , $E[\Omega] - 1$ of them on average, and each costs the premium pQ on units surged for nothing, together with the holding of the hedge while it is acquired and drawn down, $hQ^2C/(2(C-1))$, the two ramps of (5.4) combined. The count is $E[\Omega] - 1$ because the open spell of (3.5) has $E[\Omega]$ open periods on average, the detector is silent through the $E[T]$ closed periods that complete the cycle, and the one open period in which the countdown stands at L belongs to the true alarm. The count also includes the first $M - L + 1$ periods of the spell, in which the countdown cannot yet be L and an alarm is known to be false. The defender ignores those alarms, so (5.6) counts false alarms that are never acted on.

We take the first-order value of the warning, V_1 , to be the saving over the same accounting with no hedge, $J_1(S_0, S_0)$, which is the uninformed cost (4.2), divided by the mean attack cycle $E[\Lambda]$

so that, like V , it is a value per period and can be compared with the price w of the detector, as in (1.1):

$$V_1(L, \alpha, \beta) = \frac{\min_{S_0} J_1(S_0, S_0) - \min_{S_0 \leq S_A \leq S_0 + (C-1)L} J_1(S_0, S_A)}{E[\Lambda]}. \quad (5.7)$$

With one accounting in both terms, and within this approximation, a detector on whose alarms the defender acquires nothing is worth nothing, as it must be. The ignored false alarms raise only the cost with a hedge, so V_1 is conservative on that account.

Proposition 11 (Three tiers). *Let the informed defender's levels (S_0^*, S_A^*) minimize (5.6), and let S^* be the uninformed defender's quiet inventory of Proposition 8, which minimizes $J_1(S, S)$. Then $S_0^* \leq S^* \leq S_A^*$.*

The proposition sorts the units of cover into three tiers: units $s \leq S_0^*$ are held in every risky period, units $S_0^* < s \leq S_A^*$ are acquired on an alarm, and units $s > S_A^*$ are never held. A unit is held in every risky period if and only if the backlog it averts in the attacks the detector misses, together with the premium and holding it saves on every alarm by not having to be acquired, exceeds the cost of carrying it between attacks.

The band of Theorem 5 describes the policy within one information state, and the tiers describe the inventory across states. A unit of the middle tier lies below the band's lower level after an alarm and above it in the quiet state. The quiet tier therefore grows with β and with the cost of a false alarm, since a costlier false alarm favors holding a unit always over acquiring it on alarms. The alarm-triggered units are the ones the uninformed defender either held through every risky period or did without, and in this sense we can read a warning system as a way to hold inventory part time. Which uninformed units move down a tier and which move up depends on where the uninformed level sat relative to the cost-loss rule (5.5), and V_1 collects the gain of every unit that moves.

5.3 The Value of Lead Time

The lead time enters (5.7) only through the constraint $Q \leq (C - 1)L$, the most the surge acquires in L periods. The factor $C - 1$ is the gateway's capability to surge and the rate at which it turns

lead time into cover. A gateway near its capacity turns even a long lead time into little cover, so a warning of a given lead time can be worth less to it. As shown in Section 6, at $C = 1.25$, eight periods of lead time acquire only two periods of cover.

Proposition 12 (The lead time that suffices). *Let (S_0^*, S_A^*) minimize (5.6) with the constraint $Q \leq (C - 1)L$ removed, let $Q^* = S_A^* - S_0^*$, and let $L^* = \lceil Q^*/(C - 1) \rceil$. Then $V_1(L, \alpha, \beta)$ is non-decreasing and concave in L , and is constant for $L \geq L^*$.*

A period of lead time is worth buying when its increment in value exceeds its increment in price, and in the first-order approximation the increments vanish beyond L^* , the time the surge capacity needs to acquire the alarm-triggered tier. When L^* exceeds the campaign length M , no admissible detector reaches the plateau. At $\rho = 0.91$ a hedge of two periods of cover needs a lead time of twenty periods, and at $\rho = 0.5$ it needs two. A busy gateway therefore keeps gaining from lead time far longer, while a gateway with a large surge capacity needs almost none. We read concavity as the observation that the units acquired last are the units worth least. The constraint binds on the last units acquired, and each period of lead time releases units of falling value. The exact value, too, rises by ever smaller increments at every utilization of Figure 4.

5.4 The Cry-Wolf Threshold

As α rises, ω falls, the cost-loss test (5.5) becomes harder to pass at every level, and the hedge shrinks toward zero. The cost (5.6) is minimized with no hedge for every α at or above a value α^* , which we call the *cry-wolf threshold*. It is not the point at which V vanishes, which Proposition 4 puts at or below $1 - \beta$ and Section 6 finds numerically a hundredth below it at the instance; there a sliver of value survives in the hold region of Theorem 5 almost until the signal is noise. It is the point beyond which the first-order hedge is zero and, in the exact solution, most of the warning's value is gone; Section 6 shows the exact policy still acquiring a fraction of a unit on an alarm well beyond it. Proposition 13 gives it in closed form.

Proposition 13 (The cry-wolf threshold). *Let S^* be the uninformed quiet inventory of Proposition 8, and let $p > 0$. The cost (5.6) is minimized with no hedge if and only if $\alpha \geq \alpha^*$, where*

$$\alpha^* = \frac{(1 - \beta) (-\tilde{G}'(S^*))}{p (E[\Omega] - 1)}, \quad (5.8)$$

with $-\tilde{G}'(S^*) = h/\lambda + hS^*/(C - 1)$ when $S^* > 0$ and $bCE[T]/(C - 1)$ when $S^* = 0$. If $p = 0$ and $\beta < 1$, the minimum has a positive hedge for every α . When the uninformed defender holds inventory,

$$\alpha^* \approx (1 - \beta) \frac{h}{p} \quad \text{when } \lambda S^*/(C - 1) \text{ and } \lambda(M - 1) \text{ are small,} \quad (5.9)$$

and when it holds nothing, α^* is nearly proportional to λ .

Two things about (5.9) deserve our notice. First, the attack probability has dropped out. A false alarm can arrive in about $1/\lambda$ periods of an attack cycle, and the loss averted by the marginal unit is about h/λ . We can see why. At the uninformed optimum, that unit is worth exactly its holding cost until the next attack, so the two cancel. A warning that costs nothing to heed but the surge premium is worth heeding as long as false alarms are rarer than $(1 - \beta)h/p$, however rare the attacks, provided the uninformed defender still holds inventory. Second, the threshold is a ratio of two costs the defender knows, the cost of carrying a unit over the cost of moving it faster than the gateway's normal pace, scaled by the detection probability. Without a premium, a false alarm costs only the holding of what it acquires, that holding vanishes at the margin, and the approximation hedges on every alarm, even on the noise line $\alpha = 1 - \beta$, where an alarm carries no information. *It is the premium that makes a false alarm expensive, and it pulls the threshold below the noise line.* It does so once p exceeds $-\tilde{G}'(S^*)/(E[\Omega] - 1)$, which is 0.97 at the instance of Section 6. Below that premium the approximation fails on the noise line, because it values alarms that carry no information. At $p = 0$ and $\alpha = 1 - \beta = 0.90$ it gives $V_1 = 0.13$ per day, where $V = 0$ by Proposition 4. Off the line it can still be close, and at the instance's detector with $p = 0$ it gives $V_1 = 1.167$ against $V = 1.169$. The threshold belongs to the first-order approximation. The exact

value can be positive above it, in the hold region of the band, and we report, in Section 6, how much of it survives.

5.5 Pricing a Detector

A detector presents itself to its buyer as an operating point $(\alpha, 1 - \beta)$ on its receiver operating characteristic, together with a lead time and a price. By (1.1) the detector is worth buying if its value exceeds its price, and we obtain from (5.6) how its value depends on its operating point.

Proposition 14 (Iso-value curves). *For fixed tiers (S_0, S_A) , V_1 is affine in α and in $1 - \beta$, a constant plus a multiple of each. At each point of the plane of α against $1 - \beta$ at which the optimal tiers (S_0^*, S_A^*) are unique and the hedge $Q = S_A^* - S_0^*$ is positive, the curve of constant V_1 has slope*

$$\frac{d(1 - \beta)}{d\alpha} = \frac{(E[\Omega] - 1) (pQ + hQ^2C/(2(C - 1)))}{\widetilde{G}(S_0^*) - \widetilde{G}(S_A^*) - hQ^2/(2(C - 1))}. \quad (5.10)$$

The slope is the detection probability the defender will give up per point of false alarm probability. Its numerator is the cost of a false alarm times the number of periods in which one can occur, and its denominator is the gain from a true alarm. When attacks are rare, the numerator carries the factor $1/\lambda$ and the curve is steep. A defender facing a rare attack should give up many points of detection to remove one point of false alarm probability. For fixed tiers, these curves are the iso-performance lines of Provost and Fawcett (2001). The detector worth buying is the one whose value exceeds its price by the most, and among detectors of one price, it is the one whose operating point lies on the highest curve of constant value. With one unit of cover and the ramp set aside, we recover from (5.10) the slope of the cost-loss model of Richardson (2000).

Remark 15 (What the warning does to the policy). Under advance demand information with firm orders, the optimal policy is a base-stock policy in an inventory position that nets out the known orders (Gallego and Özer, 2001). The counterpart here is that, on a true alarm, the defender's target moves from S_0 to S_A . The anticipated closure is netted into the position as if it were demand already announced. What has no counterpart is the drawdown. An announced order is a commitment,

and the units staged for it are consumed. An alarm is not, and the units staged for a false one are surged for nothing. That asymmetry is the whole of the cry-wolf threshold, and it is why advance disruption information is not advance demand information with the sign changed.

6 Numerical Study: A Single-Runway Gateway

We illustrate the results on one instance, verify the closed-form approximations against the exact solution of the defender's problem, and report what the exact solution shows that the first-order analysis does not. We produce every exact number below with a program that solves the optimality equation (3.12) by relative value iteration on the finite state space of Lemma 1, and every first-order number from (4.2) and (5.6), minimized over real levels.

6.1 The Instance

We take a period to be a day and the unit to be a day's consumption at a facility supplied through a single-runway airfield. A campaign is initiated with probability $\lambda = 0.02$ per day, a mean wait of fifty quiet days. It takes $M = 8$ days to mount, and it closes the airfield for a geometric duration with mean $E[T] = 4$ days. The airfield can move $C = 1.5$ days of consumption in a day, a utilization of $\rho = 0.67$. We compute on a lattice of a quarter of a day, confine the net inventory to a backlog of at most 65 days and an inventory of at most 25, where an order that would carry the inventory above the ceiling is excluded and a closure that would carry the backlog below the floor stops there, with its cost counted at the floor. The variants of Observation 6 use the same range except the two that say otherwise: the Pareto closure and the sweep over the backlog cost, which enlarge it. We use one numerical tolerance throughout, 10^{-7} . Relative value iteration pins the relative value of the closed state with no net inventory at zero and stops when the span of its increments falls below 10^{-7} , and we treat a value of warning below 10^{-7} as zero. The stationary distribution of the solved chain, from which Table 4 and the boundary probabilities come, is computed by iterating its transition matrix. The range is wide enough. In each of the 1,156 exact solves of the model

of Section 3 that this section reports, the stationary probability at the backlog boundary is below 10^{-7} . We set holding costs at $h = 1$ per unit-day, a backlog cost at $b = 10$ per unit-day, and the surge premium at $p = 5$ per unit. We equip the airfield with a tripwire detector of lead time $L = 4$, false alarm probability $\alpha = 0.05$, and miss probability $\beta = 0.10$. The stability condition (3.11) holds with room to spare, $(C - 1) E[\Omega] = 29 > 4 = E[T]$, and the recovery and restoration after a typical closure, $E[T]/(C - 1) = 8$ days in all, fit inside the safe window of $M + 1 = 9$ orders. Not every closure fits. A closure of T days consumes T units, which the surge replaces at $C - 1 = 0.5$ units a day, so its recovery and restoration last $2T$ days, more than the nine of the window when $T \geq 5$. Since a closure ends each day with probability $1/4$, the ramp runs past the window after $(3/4)^4 = 32$ percent of closures. Beyond the window, a new closure begins with probability λ each day (Lemma 6), so one begins before the ramp is done with probability $1 - (1 - \lambda)^{\lfloor 2T-9 \rfloor^+}$, whose mean over T is 0.04. A new closure thus begins during an unfinished ramp in 4 percent of cycles, which is the error that the accounting of Section 4 neglects.

6.2 Verification

In Table 3 we compare the closed-form approximations of Sections 4 and 5 against the exact solution. The exact uninformed cost is 5.907 per day, and the tripwire reduces it to 5.038. The warning is therefore worth $V = 0.869$ per day, or 15 percent of the uninformed cost. A perfect detector with the same lead time would be worth $V(4, 0, 0) = 1.464$, and the value ceiling is $\bar{V} = 2.750$. We compute the clairvoyant defender's cost as that of a perfect detector with $L = M$, which reports every campaign in its first period and is never wrong, so it sees the countdown. The tripwire thus captures 59 percent of what a perfect detector with its lead time would, and 32 percent of the ceiling. We ran a simulation of two million days under the optimal policy, on the same lattice and range, after a burn-in of fifty thousand days, in two hundred batches of ten thousand days from seed 12345 of NumPy's default generator; it returns 4.953 ± 0.065 , one standard error of the batch means, against the program's 5.038, within 1.3 standard errors. The fractile condition (4.3) gives a quiet inventory of 2.6 days, against the program's 3. The difference reflects the three simplifications

Table 3: Closed-form approximations against the exact solution. $L = 4$, $\alpha = 0.05$, $\beta = 0.10$, other parameters as in the instance. Costs per day; the first row is the instance. The first-order cost is (4.2) divided by $E[\Lambda]$, and V_1 is defined in (5.7).

C	p	Uninformed cost			Value of warning		
		g_0^*	first order	error	V	V_1	error
1.5	5	5.907	5.783	-2.1%	0.869	0.707	-18.6%
1.5	2	5.682	5.590	-1.6%	1.042	0.983	-5.7%
2	5	4.531	4.501	-0.6%	1.122	0.803	-28.4%
2	2	4.325	4.308	-0.4%	1.442	1.205	-16.4%
3	5	3.687	3.661	-0.7%	0.888	0.662	-25.5%
3	2	3.494	3.467	-0.8%	1.189	1.039	-12.6%

of the first-order approximation (Section 4.2). We find the first-order uninformed cost within 2.1 percent of the exact one. The first-order value V_1 is 0.707 against 0.869, 19 percent low. We find it low in every one of the six cases of the table, by 6 to 28 percent, with the larger errors at the larger premium. It is not a bound. At a backlog cost of 40 it is 0.945 against an exact 0.867. A policy that acquires a fixed hedge just in time and draws it down after a false alarm is one the defender can follow, but (5.6) accounts for its cost only as a first-order approximation, neglecting the alarms that overlap and pricing a false alarm's drawdown by a ramp the exact policy need not follow, so the error runs either way.

Table 4 splits the two costs into holding, backlog, and premium, from the stationary distribution of the solved chain. The warning's value is saved holding. The informed defender carries 1.07 days of inventory on average against 2.33, and incurs 0.11 more backlog cost and 0.28 more premium per day to do so. The uninformed premium, 0.374 per day, exceeds the $p E[T]/E[\Lambda] = 0.323$ of the first-order accounting because the exact policy draws surviving inventory down through the safe window and surges it back. The waiting is what makes inventory worth holding at all; the warning lets the defender hold inventory part time.

Table 4: The exact average cost per day at the instance, by cost component.

	holding	backlog	premium	total
uninformed defender	2.333	3.200	0.374	5.907
informed defender	1.074	3.315	0.649	5.038
value of warning	1.259	−0.115	−0.275	0.869
change in cost	−53.9%	+3.6%	+73.6%	−14.7%

6.3 Findings

Observation 1 (The warning moves most of the quiet inventory into the alert tier). Once it has recovered from a closure, the uninformed defender holds 3 days of cover in every risky period. With the tripwire, the quiet band is $[0.5, 1.0]$, and the informed defender holds half a day to a day. On an alarm, the lower level of the band jumps to 3 and then rises to 4, 4.5, and 5 as the closure nears, but those are targets, and at half a day per period, the inventory entering the closure is about 3 days. The three tiers of Proposition 11 are therefore about 1 day held through the risky periods and 2 days acquired on the alarm, which is what the first-order tiers $S_0^* = 1.02$ and $S_A^* = 3.02$ say. The defender would acquire more if the lead time or surge capacity allowed. After a closure whose backlog it clears inside the safe window, the defender holds nothing through that window and moves to the quiet band in period 9, the first risky period. When a false alarm is revealed, it returns to the quiet band and draws the hedge down, as the first-order analysis assumes. The surge capacity binds throughout the lead time, as we show next.

Observation 2 (Lead time keeps adding value longer at high utilization). In Figure 4, we show V against L at four utilizations, from $C = 3$ ($\rho = 0.33$) to $C = 1.25$ ($\rho = 0.8$). At $C = 3$ the value at $L = 2$ is within 6 percent of its value at $L = 8$. At $C = 2$ that takes $L = 3$. At $C = 1.5$ it is still rising at $L = 8$, from 0.869 at $L = 4$ to 1.271. And at $C = 1.25$ it has reached only 0.728 at $L = 8$ and shows no sign of leveling off. The first-order sufficient lead times, $L^* = 2, 3, 6, 8$ at $C = 3, 2, 1.5, 1.25$, mark where the gain from a further period of lead time becomes small. They do not mark where the exact value stops rising. At the instance's $C = 1.5$ it rises another 8 percent beyond $L^* = 6$. The curves rise at different rates because a period of lead time acquires $C - 1$ units




Figure 4: The value of the single warning against its lead time, for surge capacities $C = 1.25$ ($\rho = 0.80$), 1.5 ($\rho = 0.67$), 2 ($\rho = 0.50$), and 3 ($\rho = 0.33$); other parameters as in the instance.

of cover. The first period is worth 0.11, 0.24, and 0.54 at $C = 1.25$, 1.5 , and 2 , nearly in proportion to $C - 1$, while at $C = 3$ one period acquires almost the whole hedge and is worth 0.78 of the 0.90 that eight give. The curves therefore cross, and which of the four gateways values a warning most depends on its lead time: $C = 3$ at $L = 1$, $C = 2$ from $L = 2$ to 5 , and $C = 1.5$ from $L = 6$.

Observation 3 (The value of the lead time depends on the surge capacity available). A gateway seldom changes its utilization quickly, and for its utilization, there is a desirable warning lead time. Short of it, each further day of lead time acquires $C - 1$ more units of cover and adds value. Beyond it, the surge already acquires the whole hedge within the lead time, so a further day adds little and is worth buying only if it costs little. In the first-order approximation the sweet spot is the lead time that suffices of Proposition 12, beyond which more lead time adds nothing. In Figure 5 we show V against utilization at five lead times and eight surge capacities, from $\rho = 0.17$ ($C = 6$) to $\rho = 0.8$ ($C = 1.25$), all else as in the instance. At a fixed utilization, the curves of the lead times at or beyond the sweet spot nearly coincide, and those short of it fall away. Taking the sweet spot to be the lead time beyond which a further day adds less than 5 percent to the value already reached, it is at most two days up to $\rho = 0.4$, three at $\rho = 0.5$, four at $\rho = 0.57$, and seven at the instance's $\rho = 0.67$. The instance's detector, with four days of lead time, falls short of its sweet spot. Seven




Figure 5: The value of the single warning against capacity utilization $\rho = 1/C$, for lead times $L = 2, 3, 4, 6$, and 8 ; other parameters as in the instance.

days would be worth 1.248 against its 0.869, and an eighth would add 2 percent more. At $\rho = 0.8$ the sweet spot lies beyond the campaign, and the eighth day still adds 11 percent.

Observation 4 (Most of the value is gone long before the signal is worthless). In Figure 6 we show V against α for premiums of 0, 5, and 15. Using bisection, we locate the false alarm probabilities at which the computed value falls below 10^{-7} . Rounded to three decimals, these are $\alpha = 0.885$, 0.889, and 0.891 for $p = 0, 5$, and 15, respectively, about a hundredth below the noise line $1 - \beta = 0.90$. These mark where the value becomes numerically negligible on the quarter-unit lattice at the stated tolerance, not exact zeros of V . At a positive premium, the value survives that far because of the band's hold region. An informative signal lets the defender time, at no premium, where within its quiet band it sits. The value of that timing is small but lasts until the signal carries almost no information. The cry-wolf threshold (5.8) is the one that matters to a defender, beyond which most of the value is gone. At $p = 5$ it is $\alpha^* = 0.17$, whose leading term $(1 - \beta)h/p$ in (5.9) is 0.18. By $\alpha = 0.15$ the warning has lost 75 percent of its value, from 1.25 to 0.31, and by $\alpha = 0.20$, just beyond the threshold, 84 percent. By $\alpha = 0.30$ it has lost 93 percent, and the exact policy at the quiet belief, starting from no stock on hand, still orders half a day more on an alarm than on silence. At $p = 15$ the threshold is 0.06, and by $\alpha = 0.05$ the value is 0.38 against 1.24. With no

premium, the first-order analysis has no threshold, and the value declines steadily until it becomes numerically negligible just short of $1 - \beta$, because a false alarm then costs only the holding of what it acquires.

Observation 5 (The price of a detector). Figure 7 draws the curves of constant V in the plane of the false alarm probability against the detection probability. The heavy line is the diagonal $\alpha = 1 - \beta$, on which $V = 0$ by Proposition 4, and Observation 4 finds the value numerically negligible a hundredth before it. The curves are steep. From exact solves, by central differences of 0.001 in each probability, the defender gives up 3.4 points of detection probability for each point of false alarm probability at the instance, 3.0 at $(\alpha, 1 - \beta) = (0.10, 0.90)$, and 1.8 at $(0.20, 0.90)$. The closed-form approximation (5.10) gives 7.0 at the instance and 6.2 at $\alpha = 0.10$. The sign and the order of magnitude are correct, but the approximation doubles the slope, for a different reason at each point. At the instance it doubles the marginal cost of false alarms and values detection about right. At $\alpha = 0.10$ the errors trade places, and it values false alarms about right and halves the marginal value of detection. At $\alpha = 0.20$ the first-order hedge is already zero, because 0.20 lies beyond the cry-wolf threshold, $\alpha^* = 0.17$. The closed-form approximation (5.10) needs a positive hedge, so it no longer applies. A detector priced at w per day is worth buying if its operating point lies above the curve $V = w$. At this instance, cutting the false alarm probability by one point is worth several points of detection probability, as (5.10) predicts when attacks are rare. A buyer comparing detectors of the same price and lead time should therefore weigh their false alarms heavily.

Observation 6 (Robustness). We test the model against three departures from it: a heavy-tailed closure in place of the geometric one, random demand in place of constant demand, and lost demand in place of backorders. First, suppose the closure duration has a heavy tail. We replace the geometric duration by a Pareto duration of index 2.5, censored at 80 days. For the same mean, we use $P\{T \geq s\} = \min\{1, (2.7208/s)^{2.5}\}$ for $s \leq 80$, with the tail beyond 80 folded into the last value. The quiet inventory falls from 3 to 2.5 days, the uninformed cost rises from 5.91 to 5.97,




Figure 6: The value of the single warning against its false alarm probability, for surge premiums $p = 0, 5$, and 15 ; $L = 4$, $\beta = 0.10$, other parameters as in the instance.




Figure 7: Curves of constant value of warning in the $(\alpha, 1 - \beta)$ plane of false alarm probability against detection probability; $L = 4$, other parameters as in the instance. The heavy line is the diagonal $\alpha = 1 - \beta$, on which $V = 0$. Below it, where $\alpha > 1 - \beta$, lie detectors outside the model.

and the value of warning falls from 0.869 to 0.857. At equal mean, the shape of the tail matters little to the value at this instance, and Corollary 9 matters only as attacks become costlier or more frequent. The heavy tail does put mass at deep backlogs. The Pareto case is solved on a backlog range of 160 days, since the 65-day range of the other solves understates its cost by 0.10 per day, so it does matter. We also double b repeatedly, from 40 to 640. For that sweep we censor the Pareto closure at 200 days instead, with scale 2.7163 to keep its mean, and solve it on a backlog range of 320 days with the inventory ceiling raised to 57. The upper level of the uninformed defender's quiet band rises through 8, 10.5, 13, 15.5, and 18 under the geometric closure and through 7, 11, 18, 28, and 43 under the Pareto closure. Under the geometric closure the level grows with $\log b$, as Corollary 9 predicts. Its least-squares slope on these five points is 3.61, against the corollary's $1/\log(4/3) = 3.48$, and each doubling of b adds two and a half days of cover. Under the Pareto closure the level grows over this range as a power of b instead. The slope of its logarithm in $\log b$ is 0.66, against the exponent $1/(2.5 - 1) = 2/3$ that (4.5) gives for the uncensored tail, and each doubling multiplies the cover by about 1.6. At $b = 640$ the Pareto closure calls for 43 days of cover against 18. The tail, not the mean, sets how fast the cover must grow as attacks become costlier.

Second, we relax constant demand. Section 3.1 assumes it because the uncertainty that governs the hedge is on the supply side, and because constant demand lets the analysis of Sections 4 and 5 follow each unit of inventory separately. The exact solution needs no such assumption, so it can test whether the value of warning survives uncertain demand. We let demand be 0.5, 1, or 1.5 days with probabilities one quarter, one half, and one quarter, drawn independently each period after the order is placed, with the premium incurred on orders beyond the mean. The uninformed cost rises to 6.28, and the value of warning falls only to 0.855. Doubling the spread, to 0, 1, or 2, gives 7.19 and 0.674. The informed defender's quiet band widens, to $[0.5, 2.0]$ and $[1.5, 3.5]$, because it now carries inventory against demand uncertainty as well, and the warning's share of the cost falls from 15 to 9 percent.

Third, under the case of lost demand, in which a unit unmet is lost at b once and nothing is recovered, and in which we keep the acquisition cost at zero so that a lost unit saves nothing, the

warning is worth nothing at $b = 10$. The uninformed defender holds no inventory. In cost-loss terms, an alarm with posterior $\omega = 0.26$ makes a unit of cover worth 2.6 against a premium of 5, so nothing is worth acquiring. At $b = 40$ the warning recovers a value of 0.295 per day, 11 percent of the uninformed cost of 2.58. The cost of waiting is what makes a warning worth acting on; Table 4 shows where the saving is then realized.

7 Discussion

The strategic attacker. We assumed that λ does not respond to the defender's posture, and two of our results point to what an attacker who does respond could do. The first is the cry-wolf threshold. An attacker who can induce false alarms, by probing the network without attacking it, need only raise α past α^* , about $(1 - \beta)h/p$ by (5.9), to strip the defender's detector of most of its value at our instance, without touching its sensitivity. That is cheap for a detector that is cheap to trigger, and it is one reason we treat a detector's false alarm probability as an adversarial quantity rather than an engineering constant. The second is the drawdown. An attacker who observes the defender's inventory sees that the hedge acquired for a false alarm is drawn down over the Q periods that follow, and that the quiet inventory is rebuilt last after a reopening. Both are windows in which the defender is exposed, and a patient attacker times the closure to them. Against such an attacker, we expect the quiet tier to be worth more. Examining the two needs a model of its own. What we can say is that the value we compute is the value against an attacker who does not adapt; the first tactic lowers it by Proposition 4, and we expect the second to as well.

An uncertain lead time. The campaign length M and the lead time L are constants, and a true alarm therefore dates the closure exactly. That is what lets the defender acquire Q just in time, and we treat it as an idealization. A detector whose lead time is uncertain dates the closure with an error. The defender then weighs earlier acquisition and the holding it costs, against exposure to an early closure, and the hedge is carried for longer whether the alarm is true or false. In practice we expect that departure from just in time to be the rule. We expect a defender that cannot date the

closure to begin surging toward S_A on the alarm, and to carry the hedge until the closure comes or the alarm proves false. It orders just in case rather than just in time. The extra holding, h on each unit of the hedge for every period it is carried early, is the cost of hedging the timing uncertainty. It is set against the backlog the defender would incur if the closure arrived before the ramp was complete, at b per unit-period and amplified by the recovery. That is the cost-loss rule once more, with the timing error in place of the false alarm, and since the backlog cost is many times the holding cost, we expect it to favor surging at once, unless the timing error is small against the ramp. We expect the single warning to lose value as the error grows, since the date is most of what it carries; a detector that reports repeatedly can filter the countdown as it runs, and we expect it to lose less. We have computed neither, and both are directions for future study.

8 Conclusion

We asked what advance disruption information is worth to the defender of a capacity-limited transportation gateway, and we answered it for a detection system described by its lead time, its false alarm probability, and its miss probability. Its value is the reduction in the defender's long-run average cost. It is a cost per period and it is the most the defender should pay for a detector.

The analysis rests on two features of the model. The first is that the hedge is a sum of cost-loss decisions, one per unit of cover, whose cost rises and whose loss falls with the order in which the units are acquired. That is what gives the value of warning a closed-form approximation, and what ties the inventory problem to the detection literature that prices a single protective action. The second is that unserved demand waits. We find that the recovery from a closure is what makes utilization matter, through the factor $C/(C - 1) = 1/(1 - \rho)$ on every unit of backlog and through the time a hedge takes to acquire.

What the defender should take from this. Three numbers describe a detector, and we say what each is worth. A period of lead time is worth what it adds to the value, and, as a first-order approximation, it adds nothing beyond L^* , the time the surge capacity needs to acquire the alarm-

triggered tier. A busy gateway needs more of it than an idle one, and a gateway near capacity may need more than any detector can give. A false alarm probability beyond α^* , about $(1 - \beta)h/p$ when the uninformed defender holds inventory, means that acquiring on an alarm has stopped paying in the first-order approximation and, at our instance, that most of the detector's value is gone. That number can be checked on the back of an envelope from two costs the defender already knows. A detection probability is worth buying at the rate (5.10), which says how many points of detection probability a point of false alarm probability is worth, and the detector to buy is the one whose value exceeds its price by the most, which at equal price is the one whose operating point lies on the highest iso-value curve. None of these numbers requires the defender to know the attacker's intentions. They require the defender to know its own costs, its own surge capacity, and the distribution of how long a closure lasts, which is the quantity the historical record of Section 1 suggests is the least well known and the most consequential.

Future research. There are three interesting directions for future study. The first is the optimal policy under the persistent warning of Section 3.3, a detector that sees the campaign in each of its last L periods, so that the defender acts on a stream of signals rather than on a single look. The second is the setting against a strategic attacker of Section 7, which our thresholds locate but do not solve. The third is a network of gateways, where the question becomes at which of them to invest in detection. There, the two-stage structure of the problem after an alarm suggests that stochastic programming will carry what the single link's dynamic program cannot.

Statements and Declarations

Funding. No funding was received to assist with the preparation of this manuscript.

Competing interests. The author has no relevant financial or non-financial interests to disclose.

Ethics approval. Not applicable.

Data and code availability. This article uses no external data. The numerical results of Section 6 were computed with solver and verification code developed for this study, from the parameters stated in the text. The code and its outputs are available from the author on reasonable request.

Use of generative AI. During the preparation of this manuscript, the author used large language models to support specific phases of the research. Computational code for the numerical study was reviewed and debugged using Claude Opus 5.5 (Anthropic), and ChatGPT 6 Pro (OpenAI) was used to expand the scope of the literature search and act as a conversational adversary to evaluate the paper’s arguments against opposing perspectives. Following the use of these tools, the author manually reviewed, verified, and controlled all final code, source integrations, and core arguments, maintaining full accountability for the integrity and accuracy of the findings.

Appendix A: Proofs

Proof of Lemma 1

Fix a period t of an open spell and a hypothesis $K_t = k$ consistent with the gateway being open, so $k \in \{0, \dots, M\}$. Under it, the countdown in period $t - i$ was $k + i$ if $k \geq 1$ and $k + i \leq M$, and the gateway was quiet in period $t - i$ otherwise. The signal of period $t - i$ was therefore emitted at countdown L , where (3.6) gives it alarm probability $1 - \beta$, if and only if $k \geq 1$ and $k + i = L$, which requires $i \leq L - 1$. For $i \geq L$ every hypothesis gives that signal the alarm probability α , so its likelihood is the same factor under every hypothesis and cancels from Bayes’ rule (3.7). The posterior is thus a function of the prior in period t , which depends on t alone and is constant for $t \geq M + 1$ by Lemma 6, and of the signals of periods $t - L + 1, \dots, t$, of which there are finitely many combinations. □

Proof of Lemma 2

Number the periods of the closure $j = 1, \dots, T$. Nothing is ordered, so the net inventory after period j is $x - j$, and that period costs $h [x - j]^+ + b [j - x]^+$. Period j occurs if and only if $T \geq j$, which has probability $P\{T \geq j\}$, and T is independent of x . Summing over j gives (3.9). Only the periods $j < x$ incur holding, and the backlog $[j - x]^+$ grows only linearly in j , so the sum is finite when $\sum_j j P\{T \geq j\}$ is. Exchanging the order of summation gives $\sum_{j \geq 1} j P\{T \geq j\} = \sum_{t \geq 1} P\{T = t\} t(t+1)/2 = E[T(T+1)]/2$, which is finite because we assume $E[T^2] < \infty$ (Section 3.2). Each term of (3.9) is convex in x and linear between consecutive integers, so G is too. For integer $x \geq 0$, raising x by one raises $[x - j]^+$ by one when $j \leq x$ and lowers $[j - x]^+$ by one when $j \geq x + 1$, so

$$G(x+1) - G(x) = h \sum_{j=1}^x P\{T \geq j\} - b \sum_{j>x} P\{T \geq j\} = h E[\min\{T, x\}] - b E[T - x]^+.$$

Since $E[\min\{T, x\}] + E[T - x]^+ = E[T]$, the right side is $(h + b) E[\min\{T, x\}] - b E[T]$. A convex function that is linear between integers is minimized at the first integer at which its forward difference is non-negative, which is the smallest x with $E[\min\{T, x\}] \geq b E[T]/(b + h)$. $\square$

Proof of Proposition 3

The attacker is not strategic, so the sequence of open spells and closures does not depend on the policy. Let Ω_i and T_i be the lengths of the i th open spell and of the closure that ends it. The pairs are independent and identically distributed, with $E[\Omega_i] = E[\Omega]$.

Insufficient capacity. At most C is ordered in an open period and nothing in a closed one, so by (3.1) the net inventory after n periods is at most $x_1 + C N_n - n$, where x_1 is the net inventory at the start of the first period and N_n is the number of open periods among the first n . The attack cycles $\Lambda_i = \Omega_i + T_i$ are independent and identically distributed, so they form a renewal process, and the open periods of a cycle are its reward. By the renewal-reward theorem (Ross, 1996), N_n/n converges almost surely to $E[\Omega]/E[\Lambda]$, the long-run fraction of open periods. The condition

$(C - 1) E[\Omega] < E[T]$ is $C E[\Omega] < E[\Omega] + E[T] = E[\Lambda]$. Let $\eta = 1 - C E[\Omega]/E[\Lambda] > 0$. Then $x_n \leq x_1 - \eta n + o(n)$ almost surely, under every policy. The backlog therefore grows at least linearly in n , the backlog cost of period n is at least $b(\eta n - o(n))$, and the average cost over the first n periods is at least about $b \eta n/2$, which grows without bound.

Sufficient capacity. Consider the policy that holds no inventory and clears any backlog as fast as it can: $q_n = \max\{0, \min\{C, 1 - x_n\}\}$. We start from $x_1 \leq 0$. A positive initial inventory is consumed in finitely many periods and adds only a transient. Let Z_i be the backlog at the start of the i th closure. During the closure it rises by T_i , and over the open spell that follows it falls by $(C - 1)\Omega_{i+1}$ or to zero, so

$$Z_{i+1} = [Z_i + U_i]^+, \quad U_i = T_i - (C - 1)\Omega_{i+1},$$

which is the recursion of Lindley (1952) with independent and identically distributed increments. Its drift is $E[U_i] = E[T] - (C - 1) E[\Omega]$, which is negative if and only if (3.11) holds. From Lindley (1952), the recursion then has a stationary distribution Z , and from Kiefer and Wolfowitz (1956) we have that $E[Z] < \infty$. Over the i th closure and the open spell after it the backlog is at most $Z_i + T_i$ for $T_i + \Omega_{i+1}$ periods, so those periods cost at most $b(Z_i + T_i)(T_i + \Omega_{i+1}) + p(C - 1)\Omega_{i+1}$, whose expectation is finite because Z_i is independent of (T_i, Ω_{i+1}) and $E[T^2] < \infty$. The closures recur at intervals of finite mean, so the long-run average cost of this policy is finite. $\square$

Proof of Proposition 4

We use one principle throughout. It is the garbling argument of Blackwell (1953). Suppose a defender with detector D can construct, from the signals of D seen so far and from private randomization, a sequence $\tilde{W}_n$ whose joint distribution with the countdown process is that of the signals of a detector D' . Then every policy available under D' is available under D at the same cost, and $g_D^* \leq g_{D'}^*$. By (3.6) the signals of a detector are independent given the countdowns, and the distribution of W_n depends on the countdowns only through whether the campaign is visible

in period n . It therefore suffices to match, period by period, the probability of an alarm when the campaign is visible and when it is not. We prove (c) first, because (a) rests on it, then (b), (a), and (d).

Part (c). If $\alpha + \beta = 1$, the ranges below force $\alpha' = \alpha$ and $\beta' = \beta$, and there is nothing to prove. Otherwise, let $\alpha \leq \alpha' \leq 1 - \beta$ and $\theta = (\alpha' - \alpha)/(1 - \alpha - \beta)$. Given W_n , let $\tilde{W}_n = 1$ with probability $1 - \theta\beta$ if $W_n = 1$ and $\theta(1 - \beta)$ if $W_n = 0$, independently across periods; a direct calculation gives alarm probabilities $1 - \beta$ when the campaign is visible and α' otherwise, so $\tilde{W}$ has the distribution of the signals of (L, α', β) and $V(L, \alpha', \beta) \leq V(L, \alpha, \beta)$. For $\beta \leq \beta' \leq 1 - \alpha$ the corruption with $\theta = (\beta' - \beta)/(1 - \alpha - \beta)$, $1 - \theta(1 - \alpha)$ and $\theta\alpha$ does the same.

Part (b). When $\alpha + \beta = 1$ the alarm probability in (3.6) is α for every countdown, so the signals are independent of the countdown process and (3.7) leaves every belief unchanged. The informed defender's problem is then the uninformed defender's problem.

Part (a). The informed defender can ignore the signals, so $V \geq 0$. Part (c), applied first to α and then to β , gives $V(L, \alpha, \beta) \leq V(L, 0, 0)$. The clairvoyant defender sees K_n , and can draw $\tilde{W}_n$ from (3.6) for any detector, so $g_2^* \leq g_1^*$ and $V(L, 0, 0) \leq \bar{V}$.

Part (d). Let $L + 1 \leq M$, and let W_n be the signals of a detector with lead time $L + 1$. In the first open period after a reopening, let $\tilde{W}_n$ be an independent alarm with probability α , and in every other open period let $\tilde{W}_n = W_{n-1}$. We show that $\tilde{W}$ has the distribution of the signals of a detector with lead time L . In the first open period after a reopening $K_n = 0$, the campaign is visible to neither detector, and the alarm probability is α , as required. In any other open period n , period $n - 1$ was open, so $K_{n-1} \neq 1$. If $K_{n-1} \geq 2$ then $K_n = K_{n-1} - 1$, and the campaign is visible to the detector with lead time $L + 1$ in period $n - 1$ if and only if it is visible to the detector with lead time L in period n , since the two conditions are $K_{n-1} = L + 1$ and $K_n = L$. If $K_{n-1} = 0$ then $K_n \in \{0, M\}$, and since $M \geq L + 1$ the campaign is visible to neither detector in either period. In every case $\tilde{W}_n$ is an alarm with probability $1 - \beta$ when the campaign is visible to the detector with lead time L and α otherwise, and the $\tilde{W}_n$ are independent given the countdowns because the W_n

are. Finally, $\tilde{W}_n$ uses only signals seen before period n , so the defender can act on it. Therefore g^* is non-increasing in L , and since g_0^* does not depend on L , V is non-decreasing in L . $\square$

Proof of Theorem 5

The cost-to-go discounted at γ is convex in the net inventory, by the argument of Federgruen and Zipkin (1986b). The period cost (3.3) is jointly convex in (x, q) , the next inventory $x + q - 1$ is linear, G is convex by Lemma 2, and the minimum over $q \in [0, C]$ of a jointly convex function is convex in x , so convexity propagates through value iteration, with no edge to disturb it because the inventory is unbounded. Under discounting, the costs of the closed periods and the value after a closure carry powers of γ . These weights are nonnegative, so the discounted closure term is convex for the same reason G is. On an unbounded lattice of step δ , with 1 and C multiples of δ , the same steps apply with convexity as non-decreasing forward differences. Fix a belief and write $\Phi(y)$ for the convex cost-to-go from end-of-period inventory y , the expected cost from that point on net of the premium of the current order. The order minimizes $p[y - x]^+ + \Phi(y)$ over $y \in [x - 1, x + C - 1]$. Both minimizers below exist. As $y \rightarrow \infty$, $\Phi(y) \geq hy$. As $y \rightarrow -\infty$, the inventory t periods on is at most $y + (C - 1)t$, so $\Phi(y) \geq b \sum_{t \geq 0} \gamma^t [-y - (C - 1)t]^+ \geq -\frac{by}{1-\gamma} - \frac{b(C-1)\gamma}{(1-\gamma)^2}$, and $py + \Phi(y) \rightarrow \infty$ because $p < b/(1 - \gamma)$. Without that condition, the surge branch can be empty. When $p > b/(1 - \gamma)$, removing η from any order above one unit saves $p\eta$ at once and costs at most $b\eta$ in each later period, $b\eta/(1 - \gamma)$ in all, so no optimal order exceeds one unit and no finite lower level exists. With S_{hi} the minimizer of Φ and $S_{lo} \leq S_{hi}$ the minimizer of $py + \Phi(y)$, the objective is minimized at $\max\{x, S_{lo}\}$ for $y \geq x$ and at $\min\{x, S_{hi}\}$ for $y \leq x$, which is the policy of the theorem once the surge capacity and the demand truncate the move. When $p = 0$ the two minimizers coincide, and we take the same one when Φ has several. We do not carry the theorem to the average cost, for the reasons given after its statement. $\square$

Proof of Lemma 6

Let $t \geq 1$ count the periods of an open spell, with $K_1 = 0$. For $K_t = k \geq 1$ the campaign must have been initiated in period $s = t - M + k$, which requires $s \geq 2$, and the periods $1, \dots, s - 1$ must have been quiet, so $P\{K_t = k\} = \lambda(1 - \lambda)^{s-2} = \lambda(1 - \lambda)^{t-M+k-2}$ for $k \geq \max\{1, M - t + 2\}$, and $P\{K_t = 0\} = (1 - \lambda)^{t-1}$. The gateway is open in period t if and only if K_t is one of these values, so for $t \geq M + 1$

$$P\{\text{open in } t\} = (1 - \lambda)^{t-1} \left[1 + \lambda \sum_{k=1}^M (1 - \lambda)^{k-M-1} \right] = (1 - \lambda)^{t-1} (1 - \lambda)^{-M},$$

since $\lambda \sum_{k=1}^M (1 - \lambda)^{k-M-1} = (1 - \lambda)^{-M} - 1$. Dividing gives (4.1), with no dependence on t . For $t \leq M$ the smallest possible countdown is $M - t + 2 \geq 2$, so $K_t \neq 1$ and the gateway cannot close in period $t + 1$. Hence it is open with certainty in periods $1, \dots, M + 1$. For $t \geq M + 1$ the probability of a closure in period $t + 1$ is $P\{K_t = 1 \mid \text{open in } t\} = \lambda$ by (4.1). $\square$

Proof of Proposition 7

We account for one attack cycle under the quiet inventory policy, from the first risky period of one open spell to the first risky period of the next, assuming that the recovery and the restoration are complete before the first risky period, so that the closure begins at net inventory S_0 and the cycle begins and ends there.

Premium. Over a cycle, the net inventory returns to S_0 , so the orders sum to the periods, $\sum_n q_n = \sum_n 1$. The policy never orders less than one unit in an open period, so $\sum_n [q_n - 1]^+ = \sum_n (q_n - 1)$ over open periods, which equals the number of closed periods, T . The premium is pT .

Risky periods. The order in period $M + 1$ brings the net inventory to S_0 , and each open period $t \geq M + 1$ costs hS_0 . The number of such periods is geometric with mean $1/\lambda$ by Lemma 6, giving hS_0/λ .

Closure. With the drawdown continuous, the inventory falls from S_0 at one unit a period for the $U = \min\{T, S_0\}$ periods it lasts, at a holding cost of $hE \int_0^U (S_0 - t) dt = hE[S_0U - U^2/2]$.

The backlog then grows at one unit a period for the remaining $B = [T - S_0]^+$ periods, at a cost of $b \mathbb{E} \int_0^B t dt = b \mathbb{E}[B^2]/2$. These are the continuous counterparts of the two parts of G in Lemma 2.

Recovery. The backlog B falls by $C - 1$ per period and is cleared after $B/(C - 1)$ periods, at a cost of $b \mathbb{E} \int_0^{B/(C-1)} (B - (C - 1)t) dt = b \mathbb{E}[B^2]/(2(C - 1))$. With the closure's $b \mathbb{E}[B^2]/2$, the backlog costs $b \mathbb{E}[B^2] C/(2(C - 1))$.

Restoration. The quiet inventory is acquired at $C - 1$ per period over the last $S_0/(C - 1)$ periods of the safe window, so the holding is $h \int_0^{S_0/(C-1)} (C - 1)t dt = h S_0^2/(2(C - 1))$. When $T < S_0$ the surviving $S_0 - T$ units are held through the safe window as well. We neglect that term, which is of the order $h (M + 1) \mathbb{E} [S_0 - T]^+$ and is small when S_0 will not last for the typical duration. Adding the five parts gives (4.2). $\square$

Proof of Proposition 8

On real levels, the derivative of (4.2) is

$$J'(S_0) = \frac{h}{\lambda} + h \mathbb{E}[\min\{T, S_0\}] + \frac{h S_0}{C - 1} - \frac{b C}{C - 1} \mathbb{E} [T - S_0]^+,$$

since raising S_0 raises $S_0 U - U^2/2$ at the rate $U = \min\{T, S_0\}$ and lowers B^2 at the rate $2B$. Each term is continuous and non-decreasing in S_0 , so (4.2) is convex, and it is minimized at the smallest $S_0 \geq 0$ at which the derivative is non-negative, which is (4.3). At $S_0 = 0$ the condition reads $b C \mathbb{E}[T]/(C - 1) \leq h/\lambda$, so the optimal level is positive if and only if $\lambda b C \mathbb{E}[T] > h (C - 1)$, and the derivative then vanishes at S^* , where (4.3) holds with equality. $\square$

Proof of Corollary 9

For the geometric duration, $\mathbb{E} [T - S]^+ = \sum_{j>S} (1-r)^{j-1} = (1-r)^S/r$ at integer levels S , and setting $(1-r)^S/r$ equal to the right side of (4.5) gives $(1-r)^S = r h (C - 1)/(\lambda b C) = r h (1-\rho)/(\lambda b)$. $\square$

Proof of Proposition 10

The first two terms of (5.4) are convex quadratics in Q . The last is convex because $\tilde{G}$ is. Its derivative (5.3) is non-decreasing in S , since $E[\min\{T, S\}]$ rises with S and $E[T - S]^+$ falls. The derivative of (5.4) is $hQ/(C - 1) + (1 - \omega)(p + hQ) + \omega \tilde{G}'(S_0 + Q)$, and (5.5) is the condition that it be non-positive. The minimizer of a convex function on $[0, (C - 1)L]$ is the largest point at which its derivative is non-positive, and 0 when there is none. $\square$

Proof of Proposition 11

Write $X = \alpha(E[\Omega] - 1)(p + hQC/(C - 1))$ for the marginal cost of false alarms, with $Q = S_A - S_0$. The derivative of (5.6) in S_0 at fixed S_A is

$$\frac{h}{\lambda} + \frac{hS_0}{C - 1} + \beta \tilde{G}'(S_0) - (1 - \beta) \frac{hQ}{C - 1} - X,$$

and a quiet unit is worth holding if and only if this is non-positive. Its derivative in S_A at fixed S_0 is $(1 - \beta)(hQ/(C - 1) + \tilde{G}'(S_A)) + X$. Their sum,

$$\frac{h}{\lambda} + \frac{hS_0}{C - 1} + \beta \tilde{G}'(S_0) + (1 - \beta) \tilde{G}'(S_A),$$

is the derivative along the direction that raises both levels together and leaves the hedge unchanged. That move never meets the ceiling $Q \leq (C - 1)L$, since the hedge stays the same. When $S_0^* > 0$ the levels can move together either way, so the sum vanishes at the joint minimizer. When $S_0^* = 0$ they can only rise, so the sum is non-negative there. Since $\tilde{G}$ is convex and $S_A^* \geq S_0^*$, $\tilde{G}'(S_A^*) \geq \tilde{G}'(S_0^*)$. The informed objective with no hedge, $J_1(S, S)$, is the uninformed cost $J(S)$ of (4.2), whose derivative is $J'(S) = h/\lambda + hS/(C - 1) + \tilde{G}'(S)$ and whose minimizer is S^* . When $S_0^* > 0$, replacing $\tilde{G}'(S_A^*)$ in the vanishing sum by the smaller $\tilde{G}'(S_0^*)$ shows that $J'(S_0^*) \leq 0$, so by the convexity of J , $S_0^* \leq S^*$. When $S_0^* = 0$ there is nothing to prove. Replacing $\tilde{G}'(S_0^*)$ by the larger $\tilde{G}'(S_A^*)$ instead shows that $\tilde{G}'(S_A^*) \geq -(h/\lambda + hS_0^*/(C - 1))$, so that $J'(S_A^*) \geq h(S_A^* - S_0^*)/(C - 1) \geq 0$

and $S_A^* \geq S^*$. When $S_0^* = 0$, the sum is non-negative rather than zero, and the same replacement gives the same inequality, since raising a term of a non-negative sum keeps it non-negative. The two conditions that place a unit in a tier each weigh a benefit that falls with the unit index s (through $E[T - s]^+$) against holding costs that rise with s . If a unit meets a condition, every lower unit meets it too, so the three tiers are intervals. $\square$

Proof of Proposition 12

By (5.7), V_1 is the minimum cost with no hedge, which does not depend on L , less the constrained minimum of $J_1(S_0, S_A)$ over $S_0 \leq S_A \leq S_0 + (C - 1)L$, divided by $E[\Lambda]$. The lead time enters only through the constraint $Q \leq (C - 1)L$, so each of the proposition's three claims is a claim about how that minimum moves with L . First, relaxing the constraint cannot raise the minimum, so V_1 is non-decreasing in L . Second, an integer $L \geq L^* = \lceil Q^*/(C - 1) \rceil$ gives $(C - 1)L \geq Q^*$. The unconstrained minimizer (S_0^*, S_A^*) is then feasible, and the constrained minimum equals the unconstrained one. Lead time beyond L^* would let the surge acquire more than the hedge Q^* the defender wants on an alarm, so it changes nothing, and V_1 is constant. Third, V_1 is concave in L . Write (5.6) in terms of S_0 and the hedge $Q = S_A - S_0$. It is jointly convex in (S_0, Q) . Its terms are of two kinds. Some are linear or quadratic in S_0 and Q , with positive coefficients on the squares, so they are convex. The others are $\tilde{G}(S_0)$ and $\tilde{G}(S_0 + Q)$, convex too, because the derivative (5.3) of $\tilde{G}$ is non-decreasing. The minimum of a jointly convex function over the convex set $\{S_0 \geq 0, 0 \leq Q \leq c\}$ is a convex function of the bound $c = (C - 1)L$, so the informed cost is convex in L and its negative, the gain, is concave. $\square$

Proof of Proposition 13

By the proof of Proposition 12, (5.6) is jointly convex in (S_0, S_A) , so (S^*, S^*) minimizes it if and only if no feasible direction lowers the cost there. A feasible direction raises S_A at least as much as S_0 , and it may lower S_0 only when $S^* > 0$. Along the diagonal, where the hedge stays zero, the cost is (4.2), whose derivative is zero at S^* when $S^* > 0$ and non-negative when $S^* = 0$ (Proposition 8). Raising

S_A alone changes the cost at the rate $(1 - \beta) \tilde{G}'(S^*) + \alpha (E[\Omega] - 1) p$, since the terms in Q^2 have zero derivative at $Q = 0$. The hedge is therefore zero if and only if $\alpha (E[\Omega] - 1) p \geq (1 - \beta)(-\tilde{G}'(S^*))$. When $p > 0$ this is $\alpha \geq \alpha^*$. When $p = 0$ and $\beta < 1$ it fails for every α , since $\tilde{G}'(S^*) < 0$. At $S^* > 0$ the derivative of (4.2) vanishes, so $-\tilde{G}'(S^*) = h/\lambda + hS^*/(C - 1)$, and at $S^* = 0$, $-\tilde{G}'(0) = b C E[T]/(C - 1)$ by (5.3). For (5.9), (3.5) gives $E[\Omega] - 1 = 1/\lambda + M - 1$, so that

$$\alpha^* = (1 - \beta) \frac{h}{p} \cdot \frac{1 + \lambda S^*/(C - 1)}{1 + \lambda(M - 1)},$$

which is close to $(1 - \beta)h/p$ when $\lambda S^*/(C - 1)$ and $\lambda(M - 1)$ are small. When $S^* = 0$, $\alpha^* = (1 - \beta) b C E[T]/((C - 1) p (1/\lambda + M - 1))$, which is nearly proportional to λ when $\lambda(M - 1)$ is small. $\square$

Proof of Proposition 14

With the tiers fixed, (5.6) is affine in $1 - \beta$ and in α . Since $J_1(S_0, S_0)$ involves neither,

$$J_1(S_0, S_A) = J_1(S_0, S_0) - (1 - \beta) \left(\tilde{G}(S_0) - \tilde{G}(S_A) - \frac{hQ^2}{2(C - 1)} \right) + \alpha (E[\Omega] - 1) \left(pQ + \frac{hQ^2 C}{2(C - 1)} \right),$$

a constant plus a multiple of $1 - \beta$ and a multiple of α . The coefficient of α is the numerator of (5.10), and the bracket multiplying $1 - \beta$ is its denominator, so along a curve of constant V_1 the two increments cancel. Each pair of tiers gives a cost that is affine in $(\alpha, 1 - \beta)$, and the minimum in (5.7) is the lower envelope of this family. The holding terms of (5.6) rise without bound in the levels while the gain from a closure is at most $\tilde{G}(0)$, so the minimum is attained on a compact set of tiers. Wherever a single pair attains the minimum, the envelope is differentiable and has that pair's slopes in α and in $1 - \beta$. This is the envelope theorem, in the form of Danskin (1966). The optimal tiers move with $(\alpha, 1 - \beta)$, but their movement does not enter the derivative, and the slope of the curve is (5.10) at the optimal tiers. $\square$

References

- Adams, N., Chisnall, R., Pickering, C., Schauer, S., Company Peris, R., and Papagiannopoulos, I. (2021), “Guidance for ports: security and safety against physical, cyber and hybrid threats,” *Journal of Transportation Security*, 14, 197–225.
- Altuğ, M. S. and Muharremoglu, A. (2011), “Inventory Management with Advance Supply Information,” *International Journal of Production Economics*, 129, 302–313.
- Arifoğlu, K. and Özekici, S. (2010), “Optimal Policies for Inventory Systems with Finite Capacity and Partially Observed Markov-Modulated Demand and Supply Processes,” *European Journal of Operational Research*, 204, 421–438.
- (2011), “Inventory Management with Random Supply and Imperfect Information: A Hidden Markov Model,” *International Journal of Production Economics*, 134, 123–137.
- Arreola-Risa, A. and DeCroix, G. A. (1998), “Inventory Management under Random Supply Disruptions and Partial Backorders,” *Naval Research Logistics*, 45, 687–703.
- Axelsson, S. (2000), “The base-rate fallacy and the difficulty of intrusion detection,” *ACM Transactions on Information and System Security*, 3, 186–205.
- Benjaafar, S., Cooper, W. L., and Mardan, S. (2011), “Production-inventory systems with imperfect advance demand information and updating,” *Naval Research Logistics*, 58, 88–106.
- Betts, R. K. (1982), *Surprise Attack: Lessons for Defense Planning*, Brookings Institution.
- Blackwell, D. (1953), “Equivalent Comparisons of Experiments,” *Annals of Mathematical Statistics*, 24, 265–272.
- Breznitz, S. (1984), *Cry Wolf: The Psychology of False Alarms*, Lawrence Erlbaum Associates.
- Brown, G., Carlyle, M., Salmerón, J., and Wood, K. (2006), “Defending Critical Infrastructure,” *Interfaces*, 36, 530–544.
- C4ADS (Center for Advanced Defense Studies) (2019), “Above Us Only Stars: Exposing GPS Spoofing in Russia and Syria,” C4ADS, Mar. 2019.
- Cavusoglu, H., Mishra, B., and Raghunathan, S. (2005), “The value of intrusion detection systems in information technology security architecture,” *Information Systems Research*, 16, 28–46.
- Coogan, S. and Bosch, H. (2024), “After Hack, Sea-Tac Airport Website Back Up in Time for Thanksgiving Travel,” MyNorthwest, 25 Nov. 2024.
- Cybersecurity and Infrastructure Security Agency (2023), “The Attack on Colonial Pipeline: What We’ve Learned & What We’ve Done Over the Past Two Years,” CISA.
- Danskin, J. M. (1966), “The Theory of Max-Min, with Applications,” *SIAM Journal on Applied Mathematics*, 14, 641–664.

- Dzhongova, E., Anderson, D., de Ruiter, J., Novakovic, V., and Ruiz Oses, M. (2017), “False alarm rates of liquid explosives detection systems,” *Journal of Transportation Security*, 10, 145–169.
- ERR News (2024), “Finnair Suspends Flights to Tartu for 1 Month to Seek GPS Jamming Solution,” ERR News (Estonian Public Broadcasting), Apr. 2024.
- Federgruen, A. and Zipkin, P. H. (1986a), “An Inventory Model with Limited Production Capacity and Uncertain Demands I: The Average Cost Criterion,” *Mathematics of Operations Research*, 11, 193–207.
- (1986b), “An Inventory Model with Limited Production Capacity and Uncertain Demands II: The Discounted Cost Criterion,” *Mathematics of Operations Research*, 11, 208–215.
- Gaffney, J. E. and Ulvila, J. W. (2001), “Evaluation of intrusion detectors: a decision theory approach,” in *Proceedings 2001 IEEE Symposium on Security and Privacy*, pp. 50–61.
- Galindo Pacheco, G. and Batta, R. (2016), “Forecast-driven model for prepositioning supplies in preparation for a foreseen hurricane,” *Journal of the Operational Research Society*, 67, 98–113.
- Gallego, G. and Özer, Ö. (2001), “Integrating Replenishment Decisions with Advance Demand Information,” *Management Science*, 47, 1344–1360.
- Gao, L., Yang, N., Zhang, R., and Luo, T. (2017), “Dynamic Supply Risk Management with Signal-Based Forecast, Multi-Sourcing, and Discretionary Selling,” *Production and Operations Management*, 26, 1399–1415.
- Gayon, J.-P., Benjaafar, S., and de Véricourt, F. (2009), “Using Imperfect Advance Demand Information in Production-Inventory Systems with Multiple Customer Classes,” *Manufacturing and Service Operations Management*, 11, 128–143.
- Gkonis, K. G. and Psaraftis, H. N. (2010), “Container transportation as an interdependent security problem,” *Journal of Transportation Security*, 3, 197–211.
- Grabo, C. M. (2002), *Anticipating Surprise: Analysis for Strategic Warning*, Joint Military Intelligence College, Center for Strategic Intelligence Research, reissued, University Press of America, 2004, ISBN 9780761829522.
- Greenberg, A. (2018), “The Untold Story of NotPetya, the Most Devastating Cyberattack in History,” *Wired*, 22 Aug. 2018.
- Gürler, Ü. and Parlar, M. (1997), “An Inventory Problem with Two Randomly Available Suppliers,” *Operations Research*, 45, 904–918.
- Hariharan, R. and Zipkin, P. (1995), “Customer-Order Information, Leadtimes, and Inventories,” *Management Science*, 41, 1599–1607.
- Howard, R. A. (1966), “Information Value Theory,” *IEEE Transactions on Systems Science and Cybernetics*, 2, 22–26.

- Jakšič, M., Fransoo, J. C., Tan, T., de Kok, A. G., and Rusjan, B. (2011), “Inventory management with advance capacity information,” *Naval Research Logistics*, 58, 355–369.
- Karahalios, H. (2025), “Cybersecurity threats of remote autonomous ships while approaching ports,” *Journal of Transportation Security*, 18, Article 2.
- Kiefer, J. and Wolfowitz, J. (1956), “On the Characteristics of the General Queueing Process, with Applications to Random Walk,” *The Annals of Mathematical Statistics*, 27, 147–161.
- Krishnamurthy, V. (2016), *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*, Cambridge, UK: Cambridge University Press.
- Lindley, D. V. (1952), “The theory of queues with a single server,” *Mathematical Proceedings of the Cambridge Philosophical Society*, 48, 277–289.
- Lodree, Jr., E. J. and Taskin, S. (2009), “Supply chain planning for hurricane response with wind speed information updates,” *Computers & Operations Research*, 36, 2–15.
- Maersk (2018), “Annual Report 2017,” A.P. Møller-Mærsk A/S, Copenhagen.
- Mandiant (2025), “M-Trends 2025 Report,” Google Cloud Security.
- Muharremoglu, A. and Tsitsiklis, J. N. (2008), “A Single-Unit Decomposition Approach to Multi-echelon Inventory Systems,” *Operations Research*, 56, 1089–1103.
- Murphy, A. H. (1977), “The value of climatological, categorical and probabilistic forecasts in the cost-loss ratio situation,” *Monthly Weather Review*, 105, 803–816.
- Oka, M. E., Hromada, M., and Jasek, R. (2026), “Cybersecurity in critical rail infrastructure: sector specific adoption of the NIS 2 directive in Europe,” *Journal of Transportation Security*, 19, Article 25.
- O’Kelly, M. E. (2025), “Transportation security at hubs: addressing key challenges across modes of transport,” *Journal of Transportation Security*, 18, Article 4.
- Özer, Ö. and Wei, W. (2004), “Inventory Control with Limited Capacity and Advance Demand Information,” *Operations Research*, 52, 988–1000.
- Parlar, M. and Berkin, D. (1991), “Future Supply Uncertainty in EOQ Models,” *Naval Research Logistics*, 38, 107–121.
- Parlar, M., Wang, Y., and Gerchak, Y. (1995), “A Periodic Review Inventory Model with Markovian Supply Availability,” *International Journal of Production Economics*, 42, 131–136.
- Port Authority of New York and New Jersey (2012), “Port of New York/New Jersey to Close at Midnight in Anticipation of Hurricane Sandy’s Arrival,” Press release.
- Port of Seattle (2024), “Port Cyberattack Archive,” Port of Seattle, dated updates from 24 Aug. to 22 Nov. 2024.

- Provost, F. and Fawcett, T. (2001), “Robust classification for imprecise environments,” *Machine Learning*, 42, 203–231.
- Puterman, M. L. (1994), *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, New York: John Wiley and Sons.
- Raiffa, H. and Schlaifer, R. (1961), *Applied Statistical Decision Theory*, Boston: Division of Research, Graduate School of Business Administration, Harvard University.
- Regnier, E. (2008), “Public evacuation decisions and hurricane track uncertainty,” *Management Science*, 54, 16–28.
- Regnier, E. and Harr, P. A. (2006), “A dynamic decision model applied to hurricane landfall,” *Weather and Forecasting*, 21, 764–780.
- Reuters (2015), “Hackers Ground 1,400 LOT Passengers at Warsaw Airport,” ABC News (Australia), carrying the Reuters report, 22 June 2015.
- Richardson, D. S. (2000), “Skill and relative economic value of the ECMWF ensemble prediction system,” *Quarterly Journal of the Royal Meteorological Society*, 126, 649–667.
- Ross, S. M. (1996), *Stochastic Processes*, New York: John Wiley and Sons, 2nd ed.
- Saghafian, S. and Van Oyen, M. P. (2012), “The Value of Flexible Backup Suppliers and Disruption Risk Information: Newsvendor Analysis with Recourse,” *IIE Transactions*, 44, 834–867.
- (2016), “Compensating for Dynamic Supply Disruptions: Backup Flexibility Design,” *Operations Research*, 64, 390–405.
- Schmitt, A. J., Snyder, L. V., and Shen, Z.-J. M. (2010), “Inventory Systems with Stochastic Demand and Supply: Properties and Approximations,” *European Journal of Operational Research*, 206, 313–328.
- Smallwood, R. D. and Sondik, E. J. (1973), “The Optimal Control of Partially Observable Markov Processes over a Finite Horizon,” *Operations Research*, 21, 1071–1088.
- Snyder, L. V., Atan, Z., Peng, P., Rong, Y., Schmitt, A. J., and Sinsoysal, B. (2016), “OR/MS Models for Supply Chain Disruptions: A Review,” *IIE Transactions*, 48, 89–109.
- Song, J. and Zipkin, P. H. (1996), “Inventory Control with Information About Supply Conditions,” *Management Science*, 42, 1409–1419.
- Tambe, M. (2011), *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*, Cambridge University Press.
- Tan, T., Güllü, R., and Erkip, N. (2007), “Modelling Imperfect Advance Demand Information and Analysis of Optimal Inventory Policies,” *European Journal of Operational Research*, 177, 897–923.

- Taskin, S. and Lodree, Jr., E. J. (2011), “A Bayesian decision model with hurricane forecast updates for emergency supplies inventory management,” *Journal of the Operational Research Society*, 62, 1098–1108.
- Tomlin, B. (2006), “On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks,” *Management Science*, 52, 639–657.
- (2009), “Impact of Supply Learning When Suppliers Are Unreliable,” *Manufacturing & Service Operations Management*, 11, 192–209.
- Tomlin, B. T. and Snyder, L. V. (2007), “On the Value of a Threat Advisory System for Managing Supply Chain Disruptions,” Working paper, University of North Carolina and Lehigh University, version of 26 Apr. 2007.
- Treharne, J. T. and Sox, C. R. (2002), “Adaptive Inventory Control for Nonstationary Demand and Partial Information,” *Management Science*, 48, 607–624.
- Universal Cargo (2012), “Hurricane Sandy Aftermath: NY/NJ Ports Shipping Amidst Wreckage,” Blog post, November 8, 2012.
- U.S. Department of Energy, Office of Electricity Delivery and Energy Reliability (2012), “Hurricane Sandy Situation Report #15,” November 4, 2012.
- Véronneau, S. and Yoho, K. D. (2020), “Introduction to the special issue on maritime security: current and future challenges,” *Journal of Transportation Security*, 13, 133–139.
- Yolmeh, A. and Baykal-Gürsoy, M. (2018), “Urban rail patrolling: a game theoretic approach,” *Journal of Transportation Security*, 11, 23–40.