Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

State-Aware Attack Timing and Resilience Strategy Selection in Cyber-Kinetic Freight Disruption

Willow Yang

willowyang@shsu.edu

Sam Houston State University

Pamela Zelbst

Sam Houston State University

Research Article

Keywords: cyber-kinetic disruption, state-aware attack timing, transportation resilience, attacker–defender modeling, inventory buffering, freight dispersion

Posted Date: November 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11294907/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

State-Aware Attack Timing and Resilience Strategy Selection in Cyber-Kinetic Freight Disruption

Willow Yang
(**Corresponding Author**)
Sam Houston State University
College of Business Administration
Department of Analytics, Information Systems and Supply Chain
Huntsville, Texas, USA
willowyang@shsu.edu
ORCID: 0000-0002-3873-0629

Pamela J. Zelbst
Sam Houston State University
College of Business Administration
Department of Analytics, Information Systems and Supply Chain
Huntsville, Texas, USA
zelbst@shsu.edu

State-Aware Attack Timing and Resilience Strategy Selection in Cyber-Kinetic Freight Disruption

Abstract

The increasing digitalization of freight transportation creates opportunities for cyber compromise to produce physical freight movement disruptions and downstream operational consequences. This study examines whether a defender's resilience strategy should change when an attacker can observe evolving freight and inventory conditions and adapt disruption timing accordingly. We develop a sequential defender-attacker model of a time-sensitive freight system in which the defender selects an ex-ante resilience strategy before operational uncertainty is realized. State-aware attack timing is compared with stochastic disruption and a precommitted strategic benchmark while holding the attacker's disruption capability constant. The analysis considers two resilience mechanisms: inventory buffering and structural freight resilience through freight exposure dispersion or alternative transportation capability. Results show that state-aware timing can, under certain conditions, reverse the defender's preferred resilience strategy, with inventory buffering preferred under stochastic and pre-committed timing but freight dispersion becoming preferable when the attacker can condition timing on the realized operational state. In our computational analysis, selecting inventory buffering under a stochastic or pre-committed threat model when the threat is state-aware increases expected loss by 19% relative to the optimal state-aware resilience strategy. Sensitivity analysis identifies when this reversal occurs, including the roles of operational uncertainty, attacker information accuracy, freight-dispersion effectiveness, and the cost of operational interruption. Importantly, perfect operational visibility is not required for the resilience choice to change. These findings show that accounting for an adversary's information and ability to adapt disruption timing can be important when selecting resilience investments against cyber-kinetic freight disruption.

Keywords: cyber-kinetic disruption; state-aware attack timing; transportation resilience; attacker–defender modeling; inventory buffering; freight dispersion

1. Introduction

In 2026, components from an Australian F-35 in transit to the United States were unexpectedly diverted to Hong Kong, prompting an investigation into how the shipment was rerouted. Although the incident has not been

established as a deliberate cyberattack, it highlights the potential consequences of losing control over the movement of critical freight. More broadly, the growing digitalization of freight transportation creates new security risks in which cyber access can produce direct physical disruptions. In connected and automated freight systems, malicious commands can stop, disable, misroute, or otherwise interfere with vehicles carrying critical materials. Such a cyber-kinetic event can affect both freight already in transit and the transportation capacity needed for subsequent replenishment. For downstream manufacturers operating with limited inventories and time-sensitive material flows, these effects can propagate into material shortages and production interruption. The severity of those consequences depends not only on the physical disruption itself, but also on the operational state of the freight and inventory system when the disruption occurs: the same attack can cause limited disruption when inventory is sufficient but become far more damaging after inventory has been depleted or replenishment conditions have deteriorated. When an adversary can observe those evolving conditions and choose when to activate an existing cyber capability, the timing of disruption becomes state-aware; that is, the attack is conditioned on the observed freight and inventory state. This raises a central resilience question for the defender: whether and when an attacker's ability to adapt disruption timing to the defender's operational state changes the resilience strategy the defender should employ in advance.

State-aware attack timing differs from strategic timing that is determined in advance. A precommitted strategic attacker selects the attack time before uncertain freight and inventory conditions are realized and cannot revise that decision as the system evolves. A state-aware attacker, in contrast, can wait, observe the realized freight and inventory state, and then decide when to activate the same cyber capability. The key difference is therefore not whether the attacker behaves strategically; both attackers do. Rather, it is whether attack timing is fixed before the operating state is known or can be adapted to the state that actually materializes. When future operating conditions are uncertain, conditioning attack timing on the realized state can make the same cyber capability more consequential than one activated according to a preselected schedule. This matters for resilience planning because different strategies may not provide protection in the same way over time. For example, safety stock provides an inventory buffer against temporary freight disruption, but its protective value depends on how much inventory remains when the disruption occurs. Freight dispersion or alternative transportation capability protects the system differently by reducing reliance on any single freight movement and limiting the exposure created by its disruption. Its effectiveness is therefore less dependent on the inventory state at the specific time of attack. In this study, we

examine whether, and under what conditions, accounting for state-aware attack timing changes the defender's resilience decision. Specifically, when does an attacker's ability to observe evolving freight and inventory conditions and adapt the timing of a cyber-kinetic freight disruption alter the defender's optimal resilience strategy relative to stochastic and precommitted adversarial threat models?

To address this question, we develop a sequential defender-attacker model of a time-sensitive freight supply system exposed to a cyber-kinetic disruption of automated freight transportation. The defender makes an ex ante resilience investment before operational uncertainty is realized. We focus on two conceptually different resilience mechanisms: an inventory-buffering strategy, represented by safety stock, and a structural freight-resilience strategy, represented by dispersion of freight exposure or alternative transportation capability. The model compares state-aware attack timing against stochastic disruption and a precommitted strategic benchmark to determine whether access to realized freight and inventory information changes the defender's preferred resilience strategy.

The study contributes to the transportation security and resilience planning literature in three ways. First, it isolates the effect of state-aware attack timing by distinguishing an attacker that must commit to disruption timing before operational uncertainty is resolved from one that can observe the realized freight and inventory state and adapt timing accordingly. Second, the analysis shows that state awareness can affect more than the magnitude of disruption losses: under certain conditions, it can change the resilience strategy the defender should select in advance. Thus, a strategy that is optimal under stochastic or precommitted timing may become suboptimal when disruption timing can be conditioned on the evolving operational state. Third, the study identifies the conditions under which accounting for state-aware timing materially affects the defender's resilience decision and when simpler threat representations lead to the same reslience choice.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on cyber-physical risks in connected and automated transportation, intentional transportation and supply-chain disruption, and supply-chain resilience. In Section 3, we introduce the problem setting and sequential attacker–defender model, followed by the experimental design and threat-model comparisons in Section 4. Section 5 reports computational results and sensitivity analyses. Finally, Section 6 discusses the implications for transportation security and managerial decision-making, along with the study's limitations and opportunities for further research.

2. Literature Review

The growing connectivity and automation of transportation systems have increased their exposure to cyber threats capable of producing physical and operational consequences. Research on connected vehicles and transportation networks has documented vulnerabilities in vehicle communications and control, as well as the potential for compromised vehicles or infrastructure to disrupt traffic flows and network connectivity (Vivek et al. 2019; Ghosal et al. 2021). Cyberattacks on transportation-control systems can also generate broader network effects through congestion and service degradation (Laszka et al. 2019). For freight-dependent operations, these effects can extend beyond the transportation system itself. Disruptions involving shipment delays or quantity losses can propagate downstream through inventory shortages, production interruption, and subsequent recovery decisions (Paul et al. 2019).

There is a substantial body of attacker–defender research that examines intentional disruption and strategic interactions between defenders and adversaries. Sequential formulations commonly allow a defender to make protection or resource-allocation decisions while anticipating the subsequent actions of a strategic attacker (Brown et al. 2006). Hunt and Zhuang (2024), in a systematic review of 127 attacker–defender studies, document the broad use of such models across security and infrastructure-protection settings. Recent work has extended sequential attacker–defender formulations to transportation and cyber-physical resilience, although cybersecurity remains a relatively new application (Hausken et al. 2024). Hager and Reuter-Oppermann (2026) develop a stochastic defender–attacker–defender model for transportation-infrastructure protection and relief-train deployment, while Yin et al. (2026) formulate a stochastic defender–attacker–defender model for resilience investment against hybrid cyber-physical risks. Related supply-chain studies use interdiction and attacker–defender formulations to determine facility protection, backup capacity, and recovery-resource allocation under intentional disruption (Parajuli et al. 2017; Ahmadpour et al. 2025).

A relavant line of research examines adversarial decisions that adapt to evolving system states or information. Zhuang et al. (2010) show that information revealed over multiple periods can affect attacker behavior and the relative value of alternative defensive expenditures in a signaling game. Sefair and Smith (2016) study dynamic shortest-path interdiction in which the attacker can interdict arcs as the user moves through the network, and Bochkarev and Smith (2024) extend this line of research using a game-tree framework in which attack and

movement decisions evolve during the interaction. Wu et al. (2026) examine the timing of attacks in networks with capital or operating buffers and identify conditions under which delaying an attack can increase system losses. Their complete-information models optimize attack schedules against known system dynamics, while their partial-information formulation selects an initial strategy rather than conditioning subsequent attack timing on newly observed freight or inventory states. These models are closely related to the present study in their treatment of timing, information, or state-contingent adversarial action, but they do not examine whether observing the realized operating state of a freight-dependent system changes the resilience investment selected by the defender ex ante.

Disruption-management research has also examined alternative ways to protect operations against supply interruption. Inventory buffering can sustain production during temporary supply disruptions, while sourcing flexibility, contingency capacity, and network redundancy provide alternative means of maintaining supply (Tomlin 2006; Peng et al. 2011; Parajuli et al. 2017). Lewis et al. (2013), for example, examine inventory decisions when port conditions evolve according to an observed but exogenous stochastic process.

In this study, we examine a question that has yet to be resolved in the existing literature: whether the information available to an attacker when disruption timing is selected can change the resilience strategy the defender should choose in advance. Comparing stochastic disruption with strategic attack alone cannot isolate this effect because an attacker that commits in advance may already exploit predictable differences in vulnerability across periods. Likewise, a multi-period or sequential attacker–defender model is not necessarily state-aware if the attack decision cannot be conditioned on operational information revealed as the system evolves. This distinction is particularly relevant when alternative resilience strategies differ in how their effectiveness depends on the realized operational state. To isolate the effect of state-aware timing, we compare stochastic disruption timing, precommitted strategic timing, and state-aware timing while holding the attacker's underlying disruption capability and the defender's feasible resilience alternatives constant. The analysis focuses on whether and when adapting disruption timing to realized freight and inventory conditions changes the resilience strategy the defender should select ex ante.

3. Modeling Framework

3.1 Decision setting

We consider a freight-dependent operation that requires periodic replenishment to satisfy demand over a finite

planning horizon ($t = 1, \dots, T$). The operation may represent, for example, a manufacturing facility, distribution center, or other critical node that depends on inbound freight. A cyber-kinetic disruption can interfere with the physical movement of this freight, creating an immediate supply shortage or delay and a spillover effect on inbound delivery schedules in subsequent periods. We focus the model on the role of attack timing and operational information. Rather than reproducing a specific cyberattack mechanism or a detailed transportation network, we represent a cyber-kinetic event through its operational consequence: a disruption to freight availability and subsequent replenishment capacity. This abstraction allows us to isolate the mechanism central to the study: whether an attacker's ability to observe evolving freight and inventory conditions and strategically time the disruption changes the defender's preferred resilience strategy.

Before operations begin, the defender selects a resilience strategy x from a feasible set X. In the stylized cases considered in this study, $X = \{S, D\}$ contains two alternatives that address disruption through different mechanisms: inventory buffering S, which increases the inventory available before or during the planning horizon; and freight dispersion D, which reduces exposure to concentrated transportation disruption by distributing inbound freight across alternative shipments, routes, or transportation resources. Inventory buffering protects the operation by absorbing temporary supply shortfalls after they occur, whereas freight dispersion reduces the amount of replenishment exposed to a single disruption. The comparison therefore captures a basic resilience tradeoff between absorbing disruption consequences and reducing transportation exposure. The defender commits to the resilience decision before the operational state is realized. The attacker subsequently acts under one of three information structures described in Section 3.3.

3.2 Freight and inventory dynamics

Let q_t denote the baseline replenishment scheduled to arrive during period t and let d_t denote realized demand. In the absence of disruption, the available inventory evolves according to the standard balance dynamics between inbound freight and demand. Let $I_0(x)$ denote the initial inventory under resilience strategy x. Inventory buffering increases initial inventory by the specified safety-stock amount, whereas freight dispersion retains the baseline initial inventory.

Let

$$A_t = \begin{cases} 1, & \text{if the cyber-kinetic disruption is initiated in period } t, \\ 0, & \text{otherwise.} \end{cases}$$

The attacker has a common disruption budget across all threat models and may initiate the disruption at most once during the planning horizon:

$$\sum_{t=1}^{T} A_t \leq 1.$$

To capture both an immediate freight effect and a short carryover effect, realized replenishment can be represented as

$$R_t = q_t(1 - \gamma_x A_t)(1 - \delta_x A_{t-1}),$$

where γ_x represents the fraction of current-period replenishment lost or delayed when the disruption is initiated and δ_x represents a spillover effect that carries into the next period. The resilience strategy x changes the system's exposure. For example, freight dispersion can reduce the fraction of inbound freight affected by a single disruption. This formulation does not imply that every cyberattack simultaneously destroys cargo and closes transportation capacity. Instead, γ_x and δ_x provide a parsimonious representation of a cyber-kinetic event whose physical consequences may extend beyond the instant at which the attack is activated. Alternative disruption profiles can be examined through sensitivity analysis without changing the underlying information structure.

Let I_t denote ending inventory and U_t unmet demand in period t. Inventory evolves as

$$U_t = \max\{0, d_t - I_{t-1} - R_t\},$$

and

$$I_t = \max\{0, I_{t-1} + R_t - d_t\}.$$

An inventory-buffering investment increases the inventory initially available to the system or otherwise adds safety stock that can be drawn down following a transportation disruption. Freight dispersion instead reduces the operational impact of the disrupted shipment without necessarily increasing inventory. For a resilience strategy x and attack policy A, total system cost is

$$J(x, A, \omega) = C_x + \sum_{t=1}^{T} (c_u U_t + K Z_t),$$

where C_x is the cost of the resilience strategy x, c_u is the unit cost of unmet demand, Z_t equals one when a stockout occurs in period t, K represents the fixed operational consequence associated with a stockout, and ω denotes the realization of uncertain operating conditions (e.g. stochastic demand $d_t(\omega)$ or baseline replenishment $q_t(\omega)$), which in turn dictate the realized shortage U_t and stockout indicator Z_t. Including both a variable shortage cost and a fixed stockout consequence is useful because the operational significance of freight disruption may not be proportional only to the quantity short. Crossing an inventory threshold can trigger production interruption, emergency transportation, service failure, etc. The model can nevertheless be evaluated without K when only quantity-based shortage cost is appropriate.

3.3 Disruption-selection processes

To isolate the effect of attack timing, we use the same physical system and feasible defender strategies across all three threat models: stochastic timing, precommitted strategic timing, and state-aware timing. The models differ in how disruption timing is determined and, in particular, whether the attacker can condition its selection on the realized operational state.

Recall that A denotes the disruption policy over the horizon. Let $A(\tau)$ denote the specific policy where the disruption is initiated in period τ, such that $A_\tau = 1$ and $A_t = 0$ for all $t \neq \tau$. For notational convenience, we write the resulting system cost as $J(x, \tau, \omega)$. The three threat models differ in how τ is determined.

3.3.1 Stochastic disruption

Under stochastic planning, disruption timing is exogenous. Let π_t denote the probability that disruption occurs in period t. The defender selects the resilience strategy to minimize expected system cost:

$$x^{ST} = \arg\min_{x \in X} E_\omega \left[\sum_{t=1}^{T} \pi_t J(x, t, \omega) \right]$$

where τ is drawn from the specified probability distribution.

This case represents conventional planning under uncertain disruption timing. Although the disruption may occur at different points in the planning horizon, its timing is independent of the realized inventory and freight conditions.

3.3.2 Precommitted adversarial disruption

This model introduces a strategic attacker that selects the disruption period before uncertain operating conditions are realized. After observing the defender's resilience decision and knowing the system structure and probability distributions, the attacker evaluates each candidate disruption period $\tau \in \{1, \dots, T\}$ and selects the period that maximizes expected system cost:

$$\max_{\tau \in \{1,\dots,T\}} E_{\omega}\left[J(x, \tau, \omega)\right]$$

Anticipating this timing decision, the defender selects the resilience strategy $x \in X$ that minimizes the resulting worst case expected system cost:

$$x^{PC} = \operatorname{argmin}_{x \in X} \max_{\tau \in \{1,\dots,T\}} E_{\omega}[J(x, \tau, \omega)].$$

We refer to this case as precommitted strategic planning. Unlike the stochastic benchmark, the disruption period is selected strategically rather than drawn from a probability distribution. However, the attacker must commit to a period before the realized operational state is observed. This distinction allows the comparison with the state-aware model to isolate the effect of conditioning attack timing on evolving freight and inventory conditions.

3.3.3 State-aware adaptive disruption

In this model, the attacker observes the evolving operational state and decides whether to initiate the disruption or preserve the attack opportunity for a later period. Let $s_t = (I_{t-1}, d_t, q_t)$ denote the information available to the attacker at decision time t. We assume that the attacker observes current period demand and nominal replenishment before deciding whether to disrupt inbound freight but does not observe future demand or future random realizations.

Let $\mathcal{F} = \{\mathcal{F}_t\}_{t=1}^{T}$ denote the filtration generated by the observable history, where $\mathcal{F}_t = \sigma(s_1, \dots, s_t)$. At each eligible period, the attacker chooses between initiating the disruption and preserving the attack opportunity for a later period.

The attack time τ is therefore a stopping time adapted to $\mathcal{F}$. Let $\mathcal{T}(\mathcal{F})$ denote the set of feasible stopping times, including the option not to attack ($\tau = \infty$). The defender solves

$$x^{SA} = \arg\min_{x \in X} \max_{\tau \in \mathcal{T}(\mathcal{F})} E_{\omega}\,[J(x, \tau, \omega)].$$

The precommitted and state-aware attackers have the same physical disruption capability and attack opportunity. They differ in when the timing decision is made relative to the realization of operational conditions. The precommitted attacker selects the disruption period before those conditions are observed, whereas the state-aware attacker can condition the decision on freight and inventory information revealed during operations. Thus, state awareness changes the information available for selecting the attack time rather than the attacker's underlying disruption capability.

The distinction between precommitted and state-aware timing becomes consequential when uncertain operating conditions create variation in the system's vulnerability over time. If the operational trajectory is known in advance, observing the realized state provides no additional timing advantage. Our computational analysis therefore examines whether, and under what conditions, state-aware timing changes the defender's resilience decision.

3.4 Evaluating measures

The principal outcome of interest is whether the threat model changes the defender's resilience decision. Let x^{ST}, x^{PC}, and x^{SA} denote the resilience strategies selected under stochastic, precommitted strategic, and state-aware planning, respectively. A **resilience-choice reversal** occurs when $x^{SA} \neq x^{PC}$ or $x^{SA} \neq x^{ST}$. Such a reversal indicates that accounting for state-aware attack timing changes the preferred resilience strategy.

We also measure the consequence of selecting a resilience strategy under a benchmark threat model and subsequently evaluating that strategy against a state-aware attacker. Let x^{m} denote the strategy selected under benchmark model m, where $m \in \{ST, PC\}$. Its decision regret under state-aware timing is

$$Regret_{m} = V^{SA}(x^{m}) - V^{SA}(x^{SA}),$$

where $V^{SA}(x) = \max_{\tau \in \mathcal{T}(\mathcal{F})} E_{\omega}[J(x, \tau, \omega)]$ is the expected system cost of resilience strategy x when evaluated against the state-aware attacker.

Relative regret is defined as

$$RRegret_m = \frac{V^{SA}(x^m) - V^{SA}(x^{SA})}{V^{SA}(x^{SA})}.$$

These measures distinguish changes in disruption consequences from changes in resilience choice. State-aware timing may increase expected system cost without changing the preferred resilience strategy. Conversely, even a modest change in expected cost may be important if it changes the preferred strategy or produces substantial regret when a strategy selected under a benchmark threat model is evaluated against a state-aware attacker.

4. Illustrative Computational Analysis

We use a small finite-horizon system to examine whether state-aware attack timing can change the defender's preferred resilience strategy. The analysis considers a single freight-dependent operation, uncertain operating conditions, one cyber-kinetic disruption opportunity, and the two resilience strategies defined in Section 3. The physical system and feasible defender strategies remain the same across the three threat models. We first compare the resilience decisions under stochastic and precommitted strategic timing and then examine whether allowing the attacker to condition disruption timing on the realized operational state changes the preferred resilience strategy. The resulting differences in system cost, resilience choice, and decision regret provide the basis for the computational comparison.

4.1 Baseline case

The baseline case considers a four-period planning horizon with nominal replenishment of $q_t = 8$ units in each period. Demand uncertainty occurs in the first period, with $d_1 = 4$ with probability 0.40 and $d_1 = 12$ with probability 0.60. Subsequent demand is $d_2 = 10$, $d_3 = 8$, and $d_4 = 6$. The first-period demand realization affects inventory carried into subsequent periods, creating different operational states before the attacker selects the disruption time. The two resilience strategies provide protection through different mechanisms. Under inventory buffering S, five units of initial inventory are added as a buffer against subsequent freight disruption. A disruption reduces inbound freight by 50% in the disruption period and reduces replenishment by 50% in the following period. Under freight dispersion D, the system receives no additional inventory buffer but is less exposed to the

transportation disruption: inbound freight is reduced by 15% in the disruption period and by 10% in the following period.

The attacker has a single disruption opportunity and may initiate the disruption in period 2 or period 3. Unmet demand incurs a unit cost of 1, and each period with unmet demand incurs an additional fixed interruption cost of $K = 20$. The three threat models are applied to this baseline. Under stochastic timing, disruption occurs in periods 2 or 3 with equal probability, independent of the realized operational state. Under precommitted timing, the attacker selects the disruption period before first-period demand is realized. Under state-aware timing, the attacker observes the operational state resulting from the first-period demand realization and can condition the disruption time on that state.

Table 1 reports system cost for each combination of resilience strategy, first-period demand realization, and disruption time. As shown in the table, freight dispersion is less sensitive to disruption timing in the baseline case. Neither disruption time causes a loss following low first-period demand, but under high demand the corresponding costs are 68.0 and 67.2. For inventory buffering, however, the most damaging disruption time depends on the first-period demand realization. Following low demand, disruption in period 2 produces a system cost of 21, compared with zero in period 3. Following high demand, disruption in period 3 produces a cost of 67, compared with 49 in period 2. Observing the realized demand state therefore allows a state-aware attacker to select the more damaging disruption time for each state. The fundamental question is whether this timing advantage is sufficient to change the resilience strategy selected by the defender.

Table 1. System loss by resilience strategy, demand state, and attack time

Resilience strategy	First-period demand	Attack in (t=2)	Attack in (t=3)
Inventory buffering (S)	4	21.0	0.0
Inventory buffering (S)	12	49.0	67.0
Freight dispersion (D)	4	0.0	0.0
Freight dispersion (D)	12	68.0	67.2

4.2 Resilience decisions across threat models

Table 2 compares expected system cost and the resulting resilience decision under the three threat models. Under stochastic timing, the two resilience strategies perform similarly, with expected costs of 39.00 for inventory buffering and 40.56 for freight dispersion. Inventory buffering is therefore selected, although its cost advantage is small. Allowing the attacker to select the disruption time strategically in advance slightly increases expected cost. The precommitted attacker selects period 3 against inventory buffering, producing an expected cost of 40.20, and period 2 against freight dispersion, producing an expected cost of 40.80. Despite this strategic timing, the resilience decision remains unchanged, with inventory buffering retaining only a small cost advantage.

Table 2. Expected system loss under alternative threat models

Threat model	Inventory buffering	Freight dispersion	Preferred strategy
Stochastic	**39.00**	40.56	Inventory buffering
Precommitted	**40.20**	40.80	Inventory buffering
State-aware	48.60	**40.80**	Freight dispersion

Under state-aware timing, however, the expected cost for inventory buffering increases to 48.60, while freight dispersion remains at 40.80. This changes the resilience strategy selected by the defender to freight dispersion. Inventory buffering is more sensitive to attack timing because its effectiveness depends on inventory carried across periods, creating variation in vulnerability as the operational state evolves. Freight dispersion, by contrast, reduces exposure to the disruption itself and is less sensitive to disruption timing. This result does not imply that freight dispersion is generally superior to inventory buffering; rather, it shows that resilience strategies can differ in their sensitivity to state-aware attack timing.

The resilience-choice reversal also has a meaningful decision consequence. Both stochastic and precommitted strategic timings lead the defender to select inventory buffering, which produces an expected system cost of 48.60 when evaluated under state-aware timing. Selecting freight dispersion instead reduces the expected cost to 40.80. The resulting decision regret is (48.60-40.80=7.80), or approximately 19.1% relative to the cost of the strategy selected under state-aware timing. The reverse penalty is considerably smaller. If freight dispersion is selected but evaluated under stochastic or precommitted strategic timing, the corresponding cost penalties are only 1.56 and 0.60,

respectively. Thus, the baseline case reveals an asymmetric consequence: inventory buffering incurs a substantial penalty when evaluated under state-aware timing, whereas freight dispersion remains close to the minimum-cost strategy under the other two threat models.

5. Sensitivity Analysis and Boundary Conditions

The baseline case in Section 4 demonstrates that state-aware timing can alter the resilience decision when operational states create different vulnerability patterns across resilience strategies. In this section, we examine under what conditions the attacker's access to operational state information is sufficiently consequential to change that decision. We consider four dimensions: (1) the probability of the high demand state, (2) the accuracy of the attacker's operational information, (3) the effectiveness of freight dispersion in reducing disruption exposure, and (4) the fixed cost associated with a stockout. Unless otherwise stated, all other parameters retain their baseline values from Section 4.

5.1 Demand uncertainty

The baseline case assumes that the high first-period demand state occurs with probability p_H=0.60. We vary p_H to examine how the likelihood of this state affects the resilience decision under the three threat models. Under stochastic timing, inventory buffering is preferred when p_H exceeds approximately 0.522. Precommitted strategic timing yields a similar threshold, with inventory buffering preferred when $p_H > 0.525$. Under state-aware timing, the threshold differs substantially. The expected system costs for inventory buffering and freight dispersion are, respectively, $V^{SA}(S) = 21 + 46p_H$ and $V^{SA}(D) = 68p_H$. Equating these costs yields a threshold of $p_H < \frac{21}{22} \approx 0.955$, below which freight dispersion is preferred under state-aware timing. The resulting resilience-choice reversal therefore occurs over the range: $0.525 < p_H < 0.955$. Within this interval, stochastic and precommitted strategic timings lead to inventory buffering, whereas state-aware timing leads to freight dispersion.

These thresholds also identify an important boundary condition for the value of operational-state information. At relatively low values of p_H, freight dispersion is already preferred under stochastic and precommitted strategic timing, so observing the realized state does not alter the resilience decision. As p_H approaches one, the high-demand state becomes increasingly predictable before operations begin, reducing the incremental value of observing the

realized state. When p_H=1, the relevant operational state is known in advance, allowing the precommitted attacker to select the same disruption time as the state-aware attacker.

The resilience-choice reversal thus occurs when the high-demand state is sufficiently likely for inventory buffering to be preferred under stochastic and precommitted strategic timing, but not so likely that inventory buffering remains preferred even when disruption timing is conditioned on the realized operational state. State-aware timing is therefore most consequential for the resilience decision over an intermediate range of demand-state probabilities.

5.2 Accuracy of attacker information

The baseline state-aware model assumes that the attacker observes the relevant operational state correctly. To examine the effect of imperfect information, we replace perfect state observation with a binary signal. Let η denote signal accuracy. With probability η, the signal correctly identifies whether the system is in the high or low demand state; with probability $1-\eta$, it indicates the wrong state. Thus, $\eta = 0.50$ represents an uninformative signal, whereas $\eta = 1$ corresponds to perfect state observation. After receiving the signal, the attacker selects the disruption time that maximizes expected system loss conditional on the information conveyed by that signal.

Table 3 summarizes selected results. When the signal is weak, it does not alter the attacker's preferred disruption time, and the result coincides with precommitted strategic timing. For inventory buffering, the attacker begins conditioning disruption timing on the signal when η >0.5625. Beyond this point, expected system cost increases with signal accuracy as the attacker becomes more effective at identifying the more damaging disruption time for the realized state.

Table 3. Effect of attacker information accuracy

Signal accuracy η	Expected loss: Inventory buffering	Expected loss: Freight dispersion	Preferred strategy
0.50	40.20	40.80	Inventory buffering
0.60	40.92	40.80	Freight dispersion
0.70	42.84	40.80	Freight dispersion

Signal accuracy η	Expected loss: Inventory buffering	Expected loss: Freight dispersion	Preferred strategy
0.80	44.76	40.80	Freight dispersion
0.90	46.68	40.80	Freight dispersion
1.00	48.60	40.80	Freight dispersion

These results identify two distinct information thresholds: one at which signal accuracy begins to alter the attacker's timing decision, and a higher threshold at which it changes the defender's preferred resilience strategy. Inventory buffering remains preferred until its expected cost exceeds the 40.80 cost of freight dispersion, which occurs at $\eta = 0.5937$. Thus, perfect state observation is not required to produce a resilience-choice reversal. In the baseline case, moderately accurate information is sufficient to change the preferred resilience strategy. As information accuracy increases further, the cost disadvantage of inventory buffering becomes progressively larger, reaching 7.80 under perfect state observation.

5.3 Effectiveness of freight dispersion

The resilience-choice reversal also depends on the effectiveness of freight dispersion in limiting exposure to disruption. To examine this effect, we scale both disruption parameters under freight dispersion by an exposure multiplier r: $\gamma_D = 0.15r, \delta_D = 0.10r$. The baseline corresponds to $r = 1$. Values below one represent more effective dispersion, with a smaller fraction of inbound freight exposed to disruption, whereas values above one represent less effective dispersion. Three regions emerge. When dispersion is highly effective $r \leq 0.50$, freight dispersion is preferred under both precommitted strategic and state-aware timing. At $r = 0.50$, freight dispersion and inventory buffering have equal expected cost under precommitted strategic timing. In this region, state-aware timing does not produce the resilience-choice reversal observed in the baseline case because freight dispersion is already competitive under precommitted timing.

For $0.50 < r \leq 1$, the resilience decisions diverge. Inventory buffering is preferred under precommitted strategic timing, whereas freight dispersion is preferred under state-aware timing. This interval therefore represents the principal reversal region with respect to dispersion effectiveness.

When $r > 1$, the reduced effectiveness of dispersion causes an additional stockout in the low-demand state following a period-2 disruption. Because the model includes a fixed cost whenever a stockout occurs, this additional stockout produces a sharp increase in expected system cost. Inventory buffering consequently becomes preferred under both precommitted strategic and state-aware timing.

These results identify another boundary condition for the effect of state-aware timing. When freight dispersion is sufficiently effective, it is already preferred under precommitted strategic timing; when it is insufficiently effective, inventory buffering is preferred even under state-aware timing. The resilience-choice reversal occurs between these regions, where the two strategies remain competitive but differ in their sensitivity to state-aware disruption timing.

5.4 Operational interruption cost

The fixed stockout cost K captures settings in which even a modest material shortage can interrupt production, delay critical service, or require costly emergency responses. The baseline analysis uses $K = 20$. When K is small, freight dispersion is preferred because the total quantity of unmet demand dominates and freight dispersion tends to generate less expected shortage. As K increases, avoiding the occurrence and duration of stockouts becomes increasingly important.

Under stochastic timing, inventory buffering becomes preferred when $K > 4.4$. Precommitted strategic timing yields a slightly higher threshold, with inventory buffering preferred when $K > 5$. State-aware timing produces a different result: freight dispersion remains preferred across the range of nonnegative values of K. Consequently, when $K > 5$, both stochastic and precommitted strategic timing lead to inventory buffering, whereas state-aware timing leads to freight dispersion, same as the resilience-choice reversal observed in the baseline case.

Inventory buffering can prevent stockouts under some disruption timings, making it increasingly attractive as the cost of an interruption rises. Under state-aware timing, however, the attacker can condition the disruption time on the realized operational state and exploit periods in which the available buffer is insufficient to prevent or limit a stockout. Freight dispersion is less sensitive to this mechanism because it reduces the portion of inbound freight exposed to the disruption. If the inventory buffer were sufficiently large to prevent stockouts across the relevant operational states and disruption times, this timing advantage would diminish or disappear.

5.5 Boundary conditions and interpretation

Table 4 summarizes the parameter regions in which state-aware timing changes the resilience decision relative to stochastic and precommitted strategic timing. The sensitivity results show that the effect of state-aware timing depends on both the information available to the attacker and the operational conditions under which that information is used. Demand uncertainty creates variation in the system's vulnerability over time, while information accuracy determines whether the attacker can identify and exploit that variation. The effectiveness of freight dispersion determines whether the alternative resilience strategies remain sufficiently competitive for their ranking to change, and the fixed interruption cost determines the consequence of disruption pushing the system into a stockout state. The resilience-choice reversal therefore occurs when operational states remain uncertain and observable, the competing resilience strategies differ in their sensitivity to disruption timing, and the consequences of exploiting those differences are large enough to alter the preferred resilience decision.

Table 4. Summary of boundary conditions for the resilience-choice reversal

Dimension	Region without the principal reversal	Resilience-choice reversal region
Probability of high-demand state p_H	$p_H \leq 0.525$ or $p_H \geq 0.955$	$0.525 < p_H < 0.955$
Information accuracy η	$\eta \leq 0.5937$	$\eta > 0.5937$
Dispersion exposure multiplier r	$r \leq 0.5$ or $r > 1$	$0.5 < r \leq 1$
Fixed interruption cost K	$K \leq 5$	$K > 5$

6. Conclusions and Future Research

6.1 Key findings

This study examines whether the resilience strategy selected by a defender should change when an attacker can observe evolving freight and inventory conditions before deciding when to disrupt freight movement. A sequential attacker–defender framework is developed to compare three threat models: stochastic disruption timing, strategic timing selected before operational conditions are realized, and state-aware timing conditioned on the realized operational state. By holding the attacker's underlying disruption capability constant across these settings, the analysis isolates the effect of operational information and adaptive timing on the defender's resilience decision.

The results demonstrate that state-aware attack timing can reverse the defender's preferred resilience strategy even when the attacker's physical disruption capability remains unchanged. In the baseline case, inventory buffering is preferred under both stochastic and precommitted attack timing, whereas freight dispersion becomes preferable when the attacker can observe the operational state and adapt the timing of disruption. The underlying mechanism is the difference in how the two strategies provide protection over time. Inventory buffering carries protection across periods, but its effectiveness depends on how much inventory remains after earlier demand has been realized. A state-aware attacker can exploit this temporal variation by waiting until the buffer has been depleted and the system is more vulnerable. Freight dispersion instead reduces the portion of inbound freight exposed to disruption and is therefore less dependent on the contemporaneous inventory state. As a result, selecting inventory buffering based on a stochastic or precommitted threat model when the relevant threat is state-aware produces decision regret of approximately 19%.

The sensitivity analysis shows that this resilience-choice reversal is not universal. It depends on demand uncertainty, the accuracy of the attacker's operational information, the effectiveness of freight dispersion, and the cost of operational interruption. State-aware timing is most consequential when operational conditions vary sufficiently for vulnerable states to emerge, and the attacker has enough information to distinguish those states. Perfect operational visibility by the attacker is not required; even imperfect information can be sufficient to alter the resilience strategy the defender should select. Conversely, accounting explicitly for state-aware timing is less likely to change the resilience decision when operational conditions are highly predictable, the attacker has limited information about the evolving system state, or one resilience strategy remains strongly dominant across states.

6.2 Implications for transportation security and resilience planning

These findings have important implications for cyber-kinetic freight security. Connected freight systems increasingly rely on transportation management systems, vehicle connectivity, inventory platforms, sensors, and other digital sources of operational information. Such technologies improve visibility and coordination for legitimate users, but compromised access can also reveal when freight-dependent operations are particularly vulnerable to disruption. The security problem therefore extends beyond preventing unauthorized control of transportation assets or protecting the physical movement of freight. It also includes protecting operational information that could allow an adversary to identify when disruption would have the greatest effect. Limiting unnecessary access to real-time

freight and inventory information can itself become part of a broader transportation-resilience strategy.

The results also suggest that resilience investments should be evaluated not only by their expected performance under disruption, but also by how their effectiveness varies across operational states that may be observable to an adversary. A strategy that performs well against stochastic or precommitted disruption may become less attractive if its protection deteriorates in particular inventory or freight states that an attacker can identify. Conversely, strategies whose effectiveness is less sensitive to temporary operational conditions may provide greater robustness to state-aware timing. This does not imply that freight dispersion is generally preferable to inventory buffering. Rather, resilience planners should consider both the physical disruption capability of the adversary and the operational information available to that adversary when comparing alternative resilience investments.

6.3 Future research

Future research can extend the present framework in several directions. The current model uses a stylized finite-horizon freight system to isolate the effect of state-aware timing. Future work could examine transportation networks with multiple origins, destinations, facilities, routes, and inventory locations, allowing attack timing to interact with rerouting, congestion, network redundancy, and spatial variation in freight exposure. The resilience decision could also be expanded from mutually exclusive alternatives to portfolios combining inventory buffering, transportation diversification, alternative sourcing, cyber hardening, detection, recovery capacity, and emergency logistics. Future research could also model the information environment in greater detail. The present analysis captures imperfect observation through signal accuracy, but operational information may also be delayed, incomplete, strategically protected, or deliberately distorted. Incorporating information access, detection risk, deception, and defensive information protection would allow the value of operational visibility itself to become part of the defender–attacker interaction.

ACKNOWLEDGEMENTS

The authors acknowledge the support provided by the Institute for Homeland Security at Sam Houston State University.

Statements and Declarations

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Competing Interests

The authors have no relevant financial or non-financial interests to disclose.

Author Contributions

Willow Yang and Pamela Zelbst contributed to the study conception and design. Model development and analysis were performed by Willow Yang and Pamela Zelbst. The first draft of the manuscript was written by Willow Yang, and both authors commented on previous versions of the manuscript. All authors read and approved the final manuscript.

Data Availability

The illustrative dataset generated and analyzed during the current study is included within the article.

Ethical Approval

Not applicable.

Informed Consent

Not applicable.

References

Ahmadpour, L. Sahraeian, R. & Eshghi, K. (2025) Resilient downstream oil supply chain under interdiction and recovery resources allocation using a Column and Constraint Generation algorithm. *Computers & Operations Research, 177*, 106978.

Bochkarev, A. A. & Smith, J. C. (2024) Monte Carlo tree search for dynamic shortest-path interdiction. *Networks, 84*(4), 398–419.

Brown, G. Carlyle, M. Salmerón, J. & Wood, K. (2006) Defending Critical Infrastructure. *Interfaces, 36*(6), 530–544. DOI: 10.1287/inte.1060.0252.

Ghosal, A. Sagong, S. U. Halder, S. Sahabandu, K. Conti, M. Poovendran, R. & Bushnell, L. (2021) “Truck platoon security: State-of-the-art and road ahead. *Computer Networks, 185*, 107658.

Hager, F. & Reuter-Oppermann, M. (2026) A stochastic tri-level interdiction model for relief train location and infrastructure protection. *Transportation Research Part E: Logistics and Transportation Review, 213*, 104963. DOI: 10.1016/j.tre.2026.104963.

Hausken, K. Welburn, J. W. & Zhuang, J. (2024) A Review of Attacker–Defender Games and Cyber Security. *Games, 15*(4), Article 28. DOI: 10.3390/g15040028.

Hunt, K. & Zhuang, J. (2024) A review of attacker-defender games: Current state and paths forward. *European Journal of Operational Research, 313*(2), 401–417. DOI: 10.1016/j.ejor.2023.04.009.

Laszka, A. Abbas, W. Vorobeychik, Y. & Koutsoukos, X. (2019) Detection and mitigation of attacks on transportation networks as a multi-stage security game. *Computers & Security, 87*, 101576.

Lewis, B. M. Erera, A. L. Nowak, M. A. & White, C. C. III. (2013) Managing inventory in global supply chains facing port-of-entry disruption risks. *Transportation Science, 47*(2), 162–180.

Parajuli, A. Kuzgunkaya, O. & Vidyarthi, N. (2017) Responsive contingency planning of capacitated supply networks under disruption risks. *Transportation Research Part E: Logistics and Transportation Review, 102*, 13–37.

Paul, S. K. Asian, S. Goh, M. & Torabi, S. A. (2019) Managing sudden transportation disruptions in supply chains under delivery delay and quantity loss. *Annals of Operations Research, 273*, 783–814.

Peng, P. Snyder, L. V. Lim, A. & Liu, Z. (2011) Reliable logistics networks design with facility disruptions. *Transportation Research Part B: Methodological, 45*(8), 1190–1211.

Sefair, J. A. & Smith, J. C. (2016) Dynamic shortest-path interdiction. *Networks, 68*, 315–330.

Tomlin, B. (2006) On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science, 52*(5), 639–657.

Vivek, S. Yanni, D. Yunker, P. J. & Silverberg, J. L. (2019) Cyberphysical risks of hacked internet-connected vehicles. *Physical Review E, 100*, 012316.

Wu, K. Amini, H. & Sefair, J. A. (2026) *Strategic attacks, network losses, and cascading failures.* SSRN preprint.

Yin, Z. Fang, C. Fang, Y. & Xie, M. (2026) A stochastic defender-attacker-defender model for smart system resilience enhancement against hybrid cyber-physical risks. *IISE Transactions*.

Zhuang, J. Bier, V. M. & Alagoz, O. (2010) Modeling secrecy and deception in a multiple-period attacker–defender signaling game. *European Journal of Operational Research, 203*(2), 409–418.