

A Two-Population Delayed Evolutionary Game with Bayesian Learning for Cyber Defense

Sina Keyhanian

s.keyhanian@aut.ac.ir

Supreme National Defense University

Research Article

Keywords: Evolutionary Game Theory, Advanced Persistent Threats, Delay Differential Equations, Bayesian Filtering, Hopf Bifurcation, Cybersecurity

Posted Date: October 16th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11253630/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

A Two-Population Delayed Evolutionary Game with Bayesian Learning for Cyber Defense

Sina Keyhanian^{1,2*}

¹Assistant Professor of Strategic Management of Cyberspace, Supreme
National Defense University, Tehran, Iran.

²Visiting Lecturer at Department of Industrial Engineering &
Management Systems, Amirkabir University of Tehran, Tehran, Iran.

Corresponding author. E-mail: s.keyhanian@aut.ac.ir;

Abstract

Advanced Persistent Threats (APTs) exploit detection latency, yet evolutionary models of cyber defense typically assume instantaneous adaptation. We formulate a two-population evolutionary game where defenders observe the attacker mix (aggressive vs. stealthy) only through a delayed Poisson alert stream, filtered via a Bayesian approach. The mean-field limit of this Kalman–Bucy filter is a belief equation whose gain κ grows with sensor coverage. This model couples two replicator equations and the filter through a detection delay τ_d and a remediation delay τ_r . For $\tau_r = 0$, we prove a closed-form detection-delay threshold $\tau_d^*(\kappa)$ where the interior Nash equilibrium undergoes a Hopf bifurcation, supercritical throughout the parameter ranges examined, leading to recurrent attacker–defender cycles. When attacker adaptivity exceeds congestion damping, stability requires a minimum learning speed regardless of latency. With both delays present, a moderate remediation lag can non-additively enlarge the admissible detection delay. A semi-empirical calibration on 303,220 incidents from Microsoft’s GUIDE dataset estimates the learning speed at **0.07** per day for a single organization, rising to **3** per day for pooled telemetry, multiplying the admissible detection delay by **1.3–2.4** when $m/b^2 = 0.6$. Incident-management logs (ServiceNow, BPIC 2013) show congestion relaxation takes 3–6 days, placing incident queues close to the remediation threshold when accounting for heavy-tailed resolution times. The classical threshold is recovered as the frozen-attacker limit.

Keywords: Evolutionary Game Theory, Advanced Persistent Threats, Delay Differential Equations, Bayesian Filtering, Hopf Bifurcation, Cybersecurity

1 Introduction

Advanced Persistent Threats (APTs) and other targeted intrusions are fundamentally contests over time. On the defender’s side of the clock, frontline investigations report a global median dwell time of 14 days in 2025, up from 11 days in the previous year, with latency significantly exacerbated to a median of 122 days for cyber-espionage and insider-style IT-worker incidents, and extending to nearly 400 days for backdoors implanted on network edge devices that lack endpoint telemetry [Mandiant \(2026\)](#). The adversary’s operational timeline, conversely, has compressed sharply. The median time between initial access and the hand-off to a follow-on intrusion team plummeted from over eight hours in 2022 to 22 seconds in 2025, while the estimated mean time to exploit a vulnerability has become negative (-7 days), indicating that exploitation typically precedes official patch releases [Mandiant \(2026\)](#). Across 600 breached organizations, IBM measured a mean breach lifecycle of 241 days in 2025, split into 181 days to identify the breach and 60 days to contain it [IBM Security and Ponemon Institute \(2025\)](#).

Documented corporate breaches consistently exhibit this two-stage latency, characterized by an extended detection delay followed by a protracted response lag. During the 2017 Equifax breach, attackers exploited a web-application vulnerability for approximately 76 days before suspicious traffic was noticed, although a patch had been available since March [U.S. Government Accountability Office \(2018\)](#). Similarly, the 2013 Target intrusion generated malware alerts that ultimately went uninvestigated in time [U.S. Senate Committee on Commerce, Science, and Transportation \(2014\)](#), while unauthorized access in the Marriott–Starwood case persisted undetected from 2014 until 2018 [Information Commissioner’s Office \(2020\)](#). Operational data suggest where defender time goes: in Microsoft’s GUIDE dataset of real incidents triaged by Security Operations Center (SOC) analysts [Freitas, Kalajdjieski, Gharib, and McCann \(2024\)](#), only about one graded incident in five is a true positive (Section 5.6), meaning analyst capacity is consumed largely by benign and false alerts. Defense is therefore not a binary state of prevention but a time-sensitive pipeline of sensing, triage, validation, and containment [Dash and Acken \(2026\)](#). In that pipeline, detection delay, triage congestion, and the speed at which defenders learn how attackers are behaving jointly determine outcomes.

Despite the inherently temporal nature of these threats, the prevailing mathematical frameworks in cyber defense—particularly static and simultaneous-move game-theoretic models—routinely rely on the assumption of instantaneous payoffs. This assumption compresses the protracted reality of reconnaissance, lateral movement, and privilege escalation into a single payoff event [Saint-Hilaire, Neal, Cuppens, et al. \(2025\)](#); [Yang, Zhang, and Song \(2025\)](#). While Evolutionary Game Theory (EGT) provides a robust alternative for modeling population-level strategic adaptation [Bashir, Shamszaman, Song, and Han \(2026\)](#), current cyber EGT approaches largely fail to explicitly parameterize operational delays (e.g., detection latency and remediation lag) or account for the incomplete information defenders have about how the attacker population is currently behaving [Tan, Jin, Zhang, et al. \(2023\)](#).

Concurrently, cyber-epidemic and CPS models have shown that time delays act as stability boundaries: in delayed malware-propagation models, exceeding a critical delay

destabilizes the malware-free or endemic equilibrium through a Hopf bifurcation whose direction can be obtained by normal-form theory [Nithya, Madhusudanan, Murthy, et al. \(2024\)](#), and delay-dependent oscillatory outbreaks have been reported numerically in medical Internet of Things (IoT) networks [Kumar, Rai, and Mishra \(2026\)](#). In these models, however, the attacker is a compartment rather than a strategic player. Conversely, the threshold delay of two-strategy replicator dynamics is classical [Alboszta and Miękisz \(2004\)](#); [Wesson, Rand, and Rand \(2016\)](#); [Yi and Zuwang \(1997\)](#), and the connection between sequential Bayesian inference and replicator–mutator dynamics has been made rigorous [Pathiraja and Wacker \(2025\)](#). What is missing is a model in which a Bayesian filter of a strategically adapting attacker population feeds the defenders’ fitness through an explicit information delay, and in which that coupling, rather than a single aggregated lag, determines stability.

This paper supplies such a model. The main contributions are:

1. A two-population evolutionary game between defending subsystems (proactive hunting vs. passive monitoring) and intrusion campaigns (aggressive vs. stealthy), in which defenders evaluate payoffs through the posterior mean of the attacker mix, obtained as the mean-field limit of a Kalman–Bucy filter of a delayed Poisson alert stream. The filter gain κ is tied explicitly to alert discriminability and sensor coverage.
2. A closed-form detection-delay threshold $\tau_d^*(\kappa)$ for the coupled system (Theorem 3), a minimum learning-speed condition $\kappa_{\min}$ that no latency reduction can circumvent, and the limiting tolerances for slow and instantaneous learning (Corollary 4). The classical single-population threshold is recovered as the frozen-attacker limit (Theorem 6).
3. The direction of the bifurcation via the first Lyapunov coefficient (Proposition 5): the bifurcation is supercritical throughout the parameter ranges examined, and the predicted amplitude matches simulation.
4. A two-delay stability chart showing that attacker adaptation lowers the remediation-delay threshold by 29% relative to the single-population benchmark, and that detection and remediation delays are not additive. A single Hopf curve separates stable from oscillatory regimes, and along it the destabilizing mode changes continuously from a slow belief-loop oscillation to a fast congestion oscillation.
5. A finite-population validation with explicit Poisson alerts, confirming the mean-field limit with fluctuations of order $N^{-1/2}$.
6. A semi-empirical calibration of the learning speed on 303,220 analyst-triaged incidents from Microsoft’s GUIDE dataset. The estimator of κ does not require the unobservable alert rates of each tactic. It shows that pooled telemetry raises κ from about 0.07 to about 3 per day and multiplies the admissible detection delay by 1.3–2.4 for $m/b^2 = 0.6$, and by considerably more for $m/b^2 = 1.2$ (Section 5.6).
7. A calibration of the remediation channel on two public incident-management logs. It gives a congestion relaxation time of three to six days and shows that the heavy tail of resolution times, rather than the typical case, decides the stability of the remediation channel (Section 5.7).

2 Related Work

2.1 Dwell Time and the Limits of Static Security Games

Recent literature consistently emphasizes that delay is intrinsic to the APT lifecycle. Dwell time and detection latency are not merely post-incident metrics, but active components of attacker advantage [Dash and Acken \(2026\)](#). Delayed containment expands exploit windows and can push networks into unrecoverable ransomware or cascading cyber-physical failure regimes [Davies and MacFarlane \(2026\)](#); [Napel, van Eeten, and Parkin \(2025\)](#).

Static cyber games nevertheless typically model interactions with instantaneous, one-shot payoffs. Recent critiques argue that static matrices ignore the sequential progression of kill-chains, the decay of forensic evidence and the cumulative nature of data loss [Saint-Hilaire et al. \(2025\)](#); [Yang et al. \(2025\)](#). Partially observable Markov decision processes and stochastic games provide dynamic alternatives [Njoku \(2026\)](#), but they prioritize equilibrium computation for a few agents over population-level adaptation.

2.2 Bayesian Learning and Evolutionary Adaptation

Dynamic Bayesian games handle hidden attacker types through posterior updates across stages and target perfect Bayesian equilibria [Huang and Zhu \(2020\)](#); [Tian, Dong, and Wang \(2021\)](#); their adaptation mechanism is equilibrium-centric rather than population-dynamic. Cyber EGT models capture bounded rationality through replicator dynamics, including a two-population attacker–defender game calibrated on incident data [Bashir et al. \(2026\)](#); [Tan et al. \(2023\)](#), but without belief states. Outside security, several lines combine Bayesian reasoning with evolutionary dynamics:

- Bayes-based strategy updating under noisy signals, which alters equilibrium selection [Zhao and Li \(2020\)](#);
- Bayesian sampling equilibria with global stability results [Sawa and Wu \(2023\)](#);
- a decaying Bayesian update linked to replicator–mutator dynamics on graphs [Dragicevic \(2024\)](#);
- Bayesian inferential strategies in repeated games [Patra, Sengupta, Paul, and Chakraborty \(2024\)](#);
- incomplete information within eco-evolutionary replicator dynamics [Patra, Das Bairagya, and Chakraborty \(2025\)](#);
- a rigorous identification of the Kushner–Stratonovich filtering equation with a replicator–mutator equation [Pathiraja and Wacker \(2025\)](#).

The union of these ingredients is therefore not new in itself. The present contribution is narrower: an explicit filter of a *strategically evolving* opponent population, subject to an information delay, whose output enters the defenders’ fitness and whose stability is analyzed jointly with both replicator equations.

2.3 Delay-Induced Instability

Delay-induced loss of stability of interior equilibria in two-strategy replicator dynamics, with threshold $\pi/(2K)$, where K is the delayed self-damping coefficient, and Hopf bifurcation, is classical [Alboszta and Miękisz \(2004\)](#); [Iijima \(2011\)](#); [Wesson et al. \(2016\)](#); [Yi and Zuwang \(1997\)](#). Later work separates delays on costs and benefits [Wettergren \(2023\)](#) and introduces strategy-dependent delays [Mohamadichamgavi and Bodnar \(2025\)](#). In cybersecurity, delay-differential malware models establish Hopf bifurcations, in some cases with direction and stability of the periodic orbits [Nithya et al. \(2024\)](#); [Zhang and Li \(2026\)](#), while other studies report delay-driven oscillations numerically [Kumar et al. \(2026\)](#). Related work on delayed defense includes impulsive differential games for APT defense outsourcing [Qin, Yang, Yang, and Huang \(2024\)](#), optimal repair on time-varying networks [Wang, Li, Wang, and Meng \(2023\)](#), hypergames for deception [Wan, Cho, Zhu, and Singh \(2022\)](#) and logit dynamics in multiplayer trust games [Mangalindan, Hota, and Srivastava \(2025\)](#). Table 1 summarizes the positioning.

Table 1 Positioning relative to representative prior work (“–”: not modeled).

Study	Strategic attacker	Bayesian filter	Evolutionary dynamics	Delay	Hopf analysis	Cyber
Huang & Zhu Huang and Zhu (2020)	✓	✓ (stagewise)	–	–	–	✓
Bashir et al. Bashir et al. (2026)	✓	–	✓	–	–	✓
Nithya et al. Nithya et al. (2024)	–	–	–	✓	✓	✓
Kumar et al. Kumar et al. (2026)	–	–	–	✓	–	✓
Pathiraja & Wacker Pathiraja and Wacker (2025)	–	✓	✓	–	–	–
Delayed replicator Alboszta and Miękisz (2004) ; Wesson et al. (2016)	–	–	✓	✓	✓	–
This study	✓	✓	✓	✓ (two)	✓	✓

3 Model Formulation

We model APT defense as a two-population evolutionary game with one-sided incomplete information: a population of defending subsystems (host agents, sensors, analyst teams) and a population of intrusion campaigns adapt their strategies, while the

defenders observe the attacker population only through a delayed, noisy alert stream that they filter in a Bayesian manner. The term “Bayesian” refers only to this information-processing layer: the model is not a Harsanyi game with privately informed attacker types. Attackers are a strategic population whose tactic mix is unobserved by defenders, and the solution concept is the rest point of the coupled population dynamics.

3.1 Players, Strategies and Payoffs

Defenders choose between Proactive Hunting (D_P) and Passive Monitoring (D_M); $x(t) \in [0, 1]$ is the share playing D_P . Intrusion campaigns choose between an *Aggressive* tactic (A : rapid privilege escalation and lateral movement, which hunting can intercept) and a *Stealthy* tactic (S : low-and-slow persistence); $y(t) \in [0, 1]$ is the share playing A . Pairwise payoffs (row player first) are

$$\text{Defender: } \begin{array}{c|cc} & A & S \\ \hline D_P & \alpha - C_p & -C_p \\ D_M & 0 & 0 \end{array} \quad \text{Attacker: } \begin{array}{c|cc} & D_P & D_M \\ \hline A & g - \delta & g \\ S & 0 & 0 \end{array} \quad (1)$$

where $\alpha > 0$ is the loss averted when a hunter meets an aggressive campaign, $C_p > 0$ the friction cost of hunting, $g > 0$ the tempo gain of aggressive tactics and $\delta > 0$ the expected loss of an aggressive campaign intercepted by a hunter (payoffs of D_M and S are normalized to zero). In addition, hunters share finite analyst capacity: the hunting load committed $\tau_r \geq 0$ time units earlier, still queued in triage and remediation, imposes the congestion cost $\beta x(t - \tau_r)$, $\beta > 0$. This cost is a population-level externality (as in congestion games), not a pairwise payoff, and is therefore kept outside the matrices. Under complete information the fitness differentials are $\Delta U_D = \alpha y - C_p - \beta x(t - \tau_r)$ and $\Delta U_A = g - \delta x$. Information is one-sided. Campaigns are assumed to observe the defensive posture $x(t)$ without delay, because hunting activity is directly visible to an intruder operating inside the network. Defenders observe y only through alerts (Section 3.2). Relaxing the attackers’ observation assumption adds a third delay (Section 6).

3.2 Information Structure and Bayesian Filtering

Defenders do not observe y . They observe alerts from N monitored campaigns: each aggressive (stealthy) campaign triggers alerts at Poisson rate λ_A ($\lambda_S < \lambda_A$), and an alert surfaces only after a detection delay $\tau_d \geq 0$ (dwell before first detectable footprint plus sensor-to-triage latency). The alert count N_t thus has intensity $N[\lambda_S + (\lambda_A - \lambda_S)y(t - \tau_d)]$.

Remark 1 (Why a static hidden type cannot carry the model) If instead a single campaign had a static hidden type $\theta \in \{A, S\}$, the exact posterior would be the jump filter $\log \frac{p_t}{1-p_t} = \log \frac{p_0}{1-p_0} - (\lambda_A - \lambda_S)t + N_t \log \frac{\lambda_A}{\lambda_S}$ [Snyder \(1972\)](#): belief in A *decreases* between alerts and jumps upward at each alert. Hence (i) a smooth drift $\dot{p} \propto p(1-p)(\lambda_A - \lambda_S)$ is not a Bayesian filter; (ii) p_t is a bounded martingale converging almost surely to $\mathbf{1}\{\theta = A\}$, so no

interior belief persists; and (iii) $\mathbb{E}[p_t] = p_0$, so learning leaves history-averaged linear payoffs unchanged. A non-degenerate belief dynamic requires a latent state that itself evolves; the attacker population $y(t)$ supplies it.

Let $Z_t = (N_t/N - \lambda_S t)/(\lambda_A - \lambda_S)$ be the normalized alert record. When $N\bar{\lambda}$ is large, with $\bar{\lambda} = \lambda_S + (\lambda_A - \lambda_S)y^*$, the diffusion approximation of the counting process gives $dZ_t = y(t - \tau_d)dt + \sigma_N dW_t$ with $\sigma_N^2 = \bar{\lambda}/[N(\lambda_A - \lambda_S)^2]$. Defenders do not know the attackers' adaptation law and adopt the agnostic random-walk prior $dy = \sigma_y dB_t$; under this (deliberately misspecified [Pathiraja and Wacker \(2025\)](#)) prior, $\mathbb{E}[y(t) | \mathcal{F}_t] = \mathbb{E}[y(t - \tau_d) | \mathcal{F}_t]$, and the posterior mean $p(t)$ is given by the stationary Kalman–Bucy filter [Kalman and Bucy \(1961\)](#)

$$dp = \kappa (dZ_t - p dt), \quad \kappa = \frac{\sigma_y}{\sigma_N} = \sigma_y(\lambda_A - \lambda_S)\sqrt{N/\bar{\lambda}}.$$

Conditional on the attacker path, $dp = \kappa[y(t - \tau_d) - p]dt + \sigma_y dW_t$, whose tracking-error variance $\sigma_y\sigma_N/2$ vanishes in the joint limit $N \rightarrow \infty$ with κ fixed. The mean-field belief law is therefore

$$\dot{p}(t) = \kappa [y(t - \tau_d) - p(t)]. \quad (2)$$

The learning speed κ increases with alert discriminability $\lambda_A - \lambda_S$ and sensor coverage N , which makes it an investment variable. In the terminology of [Pathiraja and Wacker \(2025\)](#), Eq. (2) is the Gaussian closure of a replicator-type filtering equation; the new element here is that its output drives, and is driven by, strategic dynamics.

3.3 Delay-Coupled Dynamics

Because payoffs are linear in y , the defenders' posterior-expected fitness differential is exactly $\alpha p - C_p - \beta x(t - \tau_r)$ (certainty equivalence). Replicator dynamics for both populations and Eq. (2) yield

$$\begin{aligned} \dot{x}(t) &= x(t)(1 - x(t)) [\alpha p(t) - \beta x(t - \tau_r) - C_p], \\ \dot{y}(t) &= y(t)(1 - y(t)) [g - \delta x(t)], \\ \dot{p}(t) &= \kappa [y(t - \tau_d) - p(t)]. \end{aligned} \quad (3)$$

Why replicator dynamics. The replicator law is the mean-field limit of pairwise proportional imitation [Schlag \(1998\)](#): a unit occasionally samples a peer and switches to the peer's strategy with probability proportional to the payoff difference. On the defender side, this corresponds to the diffusion of hunting playbooks across teams and subsystems as they prove their worth. On the attacker side, it corresponds to the copying of successful tradecraft within shared tooling and affiliate ecosystems. Neither side is assumed to optimize; both adapt by imitation, which is the bounded-rationality premise of the evolutionary approach. The revision rate is absorbed into the payoff units.

Figure 1 summarizes the feedback architecture of Eq. (3).




Fig. 1 Feedback architecture of the model. Intrusion campaigns adapt to the hunting pressure they observe. Defenders see the campaign mix only through a delayed alert stream, which a Bayesian filter converts into the belief p . Hunting decisions based on p feed back onto the attackers and, through analyst congestion with lag τ_r , onto the defenders themselves. Stability depends on how the delays τ_d and τ_r interact with the learning speed κ (Section 4).

Remark 2 (Detection and remediation delays are not interchangeable) τ_d acts on the information channel (attacker $\rightarrow$ belief) and τ_r on the congestion feedback (defender $\rightarrow$ defender). They enter the characteristic equation (5) through different factors, so the stability region is not a function of $\tau_d + \tau_r$ (Section 5.5).

4 Equilibrium and Stability Analysis

4.1 Interior Equilibrium

Proposition 1 (Existence and uniqueness) *If $0 < g < \delta$ and $C_p + \beta g/\delta < \alpha$, system (3) has the unique interior equilibrium*

$$E^* = (x^*, y^*, p^*), \quad x^* = \frac{g}{\delta}, \quad y^* = p^* = \frac{C_p + \beta x^*}{\alpha}, \quad (4)$$

which is the unique interior Nash equilibrium of the complete-information population game with congestion.

Proof In the interior, $\dot{y} = 0 \iff x = g/\delta$, $\dot{p} = 0 \iff p = y$, and then $\dot{x} = 0 \iff \alpha y = C_p + \beta x^*$. The two inequalities are equivalent to $x^*, y^* \in (0, 1)$. At E^* both populations are indifferent between their pure strategies, which characterizes an interior Nash equilibrium. $\square$

E^* is an equilibrium allocation, not a welfare optimum: no planner objective is involved. Stability of E^* is a dynamic property and must not be conflated with evolutionary stability of the static game; in asymmetric games an interior mixed equilibrium is never an Evolutionarily Stable Strategy (ESS) Selten (1980), and without congestion ($\beta = 0$), filtering lag and delays, E^* is a neutrally stable center of the bimatrix

replicator dynamics [Hofbauer and Sigmund \(1998\)](#). Congestion supplies the damping; learning lag and delays erode it. Moreover, since the belief law (2) is unbiased at rest ($p^* = y^*$), beliefs are correct at E^* . Incomplete information therefore does not move the equilibrium; it acts only through the transient belief lag, and hence only on stability.

Lemma 2 (Well-posedness, invariance and boundary rest points) *Under the conditions of Proposition 1, for every continuous initial history with values in $[0, 1]^3$, system (3) has a unique global solution, and this solution remains in $[0, 1]^3$. The rest points on the boundary are $(0, 0, 0)$, $(1, 0, 0)$, $(0, 1, 1)$, $(1, 1, 1)$ and, if $\hat{x} = (\alpha - C_p)/\beta < 1$, also $(\hat{x}, 1, 1)$. Each of them has a real characteristic root $\lambda > 0$ that does not depend on τ_d or τ_r , and is therefore unstable for all delays.*

Proof The right-hand side is Lipschitz on bounded sets of histories, so the method of steps gives local existence and uniqueness [Hale and Verduyn Lunel \(1993\)](#). Invariance follows from three observations. First, $x(t) = x(0) \exp\{\int_0^t (1-x)F ds\}$, where F is the bracket, so x remains nonnegative; the same argument applied to $1-x$ bounds x by 1, and likewise for y . Second, $p(t) = e^{-\kappa t}p(0) + \int_0^t \kappa e^{-\kappa(t-s)}y(s-\tau_d) ds$ is a convex combination of values in $[0, 1]$. Third, boundedness then gives global existence.

At a boundary rest point, the linearization is block-triangular, and the transverse direction enters only through the undelayed factor $x(1-x)$ or $y(1-y)$. The transverse roots are therefore:

- $\lambda = g$ in the y -direction at $(0, 0, 0)$;
- $\lambda = \beta + C_p$ in the x -direction at $(1, 0, 0)$;
- $\lambda = \alpha - C_p$ in the x -direction at $(0, 1, 1)$;
- $\lambda = \delta - g$ in the y -direction at $(1, 1, 1)$;
- $\lambda = \delta\hat{x} - g$ in the y -direction at $(\hat{x}, 1, 1)$.

All of these are positive under the hypotheses: $\alpha - C_p > \beta g/\delta > 0$, and $y^* < 1$ is equivalent to $\hat{x} > x^* = g/\delta$. $\square$

When $\hat{x} \geq 1$, the four vertices are connected on the boundary by a heteroclinic cycle of matching-pennies type: hunting deters aggression, aggression rewards hunting. E^* is the only rest point that can attract interior trajectories. When it loses stability, the interior limit cycle can approach this boundary cycle, which is what the near-total attacker dominance at large τ_d in Section 5.5 reflects.

4.2 Characteristic Equation

Linearizing Eq. (3) at E^* with $u = x - x^*$, $v = y - y^*$, $\pi = p - p^*$ gives $\dot{u} = x^*(1-x^*)[\alpha\pi - \beta u(t-\tau_r)]$, $\dot{v} = -\delta y^*(1-y^*)u$, $\dot{\pi} = \kappa[v(t-\tau_d) - \pi]$. Eliminating v and π yields

$$\begin{aligned} \Delta(\lambda) &= \lambda(\lambda + \kappa)(\lambda + b e^{-\lambda\tau_r}) + \kappa m e^{-\lambda\tau_d} = 0, \\ b &= \beta x^*(1-x^*) > 0, \quad m = \alpha\delta x^*(1-x^*)y^*(1-y^*) > 0. \end{aligned} \tag{5}$$

Here b is the congestion damping and m the gain of the attacker-defender loop. Since $\Delta(0) = \kappa m > 0$, $\lambda = 0$ is never a root.

Theorem 3 (Detection-delay threshold of the coupled system) *Let $\tau_r = 0$ and let $\omega^* > 0$ be defined by $(\omega^*)^2 = z^*$, the unique positive root of*

$$F(z) = z(z + \kappa^2)(z + b^2) - \kappa^2 m^2. \quad (6)$$

(i) *If $m > b(b + \kappa)$, E^* is unstable for every $\tau_d \geq 0$.*

(ii) *If $m < b(b + \kappa)$, E^* is locally asymptotically stable if and only if $\tau_d < \tau_d^*$, where*

$$\tau_d^* = \frac{1}{\omega^*} \left[\frac{\pi}{2} - \arctan \frac{\omega^*}{\kappa} - \arctan \frac{\omega^*}{b} \right] > 0. \quad (7)$$

At $\tau_d = \tau_d^*$ a simple pair $\pm i\omega^*$ crosses the imaginary axis with $\text{Re } d\lambda/d\tau_d > 0$, and a Hopf bifurcation of periodic solutions with period close to $2\pi/\omega^*$ occurs; its direction is given by Proposition 5. No stability switch back occurs for $\tau_d > \tau_d^*$.

Proof For $\tau_d = 0$, $\Delta(\lambda) = \lambda^3 + (\kappa + b)\lambda^2 + \kappa b\lambda + \kappa m$ has positive coefficients, and by the Routh–Hurwitz criterion all roots lie in $\text{Re } \lambda < 0$ iff $(\kappa + b)\kappa b > \kappa m$, i.e. $m < b(b + \kappa)$. Write $\Delta = P(\lambda) + Qe^{-\lambda\tau_d}$ with $P(\lambda) = \lambda(\lambda + \kappa)(\lambda + b)$ and $Q = \kappa m$. A root $\lambda = i\omega$, $\omega > 0$, requires $|P(i\omega)| = Q$, i.e. $F(\omega^2) = 0$. Since $F(0) < 0$ and F is strictly increasing on $[0, \infty)$, ω^* is unique. The phase condition $e^{-i\omega^*\tau_d} = -P(i\omega^*)/Q$ with $\arg[-P(i\omega)] = \frac{3\pi}{2} + \varphi(\omega)$, $\varphi = \arctan(\omega/\kappa) + \arctan(\omega/b) \in (0, \pi)$, gives the crossing delays $\tau_{d,j} = [\frac{\pi}{2} - \varphi(\omega^*) + 2j\pi]/\omega^* \geq 0$, $j \in \mathbb{Z}$. By the criterion of Cooke and van den Driessche [Cooke and van den Driessche \(1986\)](#), $\text{sign Re } (d\lambda/d\tau_d)|_{i\omega^*} = \text{sign } F'(z^*) = +1$, so every crossing is from left to right and the number of roots in $\text{Re } \lambda > 0$ is nondecreasing in τ_d (roots depend continuously on τ_d for this retarded equation [Hale and Verduyn Lunel \(1993\)](#)). Finally, $\varphi(\omega^*) < \pi/2 \iff (\omega^*)^2 < \kappa b \iff F(\kappa b) > 0 \iff \kappa^2 b^2 (\kappa + b)^2 > \kappa^2 m^2 \iff m < b(b + \kappa)$. Under (ii) the first crossing is therefore $\tau_{d,0} = \tau_d^* > 0$ and the equilibrium is stable before it and unstable after it; under (i) the equilibrium is already unstable at $\tau_d = 0$ and remains so. Simplicity of $\pm i\omega^*$, the absence of other imaginary roots and transversality give the Hopf bifurcation [Hale and Verduyn Lunel \(1993\)](#). $\square$

Corollary 4 (Minimum learning speed and limiting tolerances) (a) *E^* is stable for sufficiently small detection delay iff $\kappa > \kappa_{\min} := \max\{0, m/b - b\}$.* (b) *As $\kappa \rightarrow \infty$ (instantaneous learning), $\tau_d^* \rightarrow \arctan(b/\omega_\infty)/\omega_\infty$ with $\omega_\infty^2 = \frac{1}{2}(\sqrt{b^4 + 4m^2} - b^2)$.* (c) *If $m < b^2$, then $\tau_d^* \rightarrow (b^2 - m)/(mb)$ as $\kappa \rightarrow 0$.*

Proof (a) restates Theorem 3. (b) $F(z)/\kappa^2 \rightarrow z(z + b^2) - m^2$ and $\arctan(\omega^*/\kappa) \rightarrow 0$. (c) Setting $z = \kappa\zeta$ gives $F/\kappa^2 \rightarrow \zeta^2 b^2 - m^2$, so $(\omega^*)^2 = \kappa m/b + o(\kappa)$; expanding $\arctan(\omega^*/\kappa) = \frac{\pi}{2} - \kappa/\omega^* + o(\kappa^{1/2})$ and $\arctan(\omega^*/b) = \omega^*/b + o(\kappa^{1/2})$ in Eq. (7) gives $\tau_d^* = \kappa/(\omega^*)^2 - 1/b + o(1)$. $\square$

Corollary 4(a) is a design condition absent from single-population delay models: when attackers adapt strongly relative to congestion damping ($m > b^2$), even zero detection delay cannot stabilize defense unless learning is fast enough, because the filter lag $1/\kappa$ acts as a distributed delay in the attacker–defender loop.

Figure 2 visualizes Theorem 3 with an independent computation. The rightmost characteristic roots are obtained by Chebyshev pseudospectral collocation of the linearized system with 40 nodes [Breda, Maset, and Vermiglio \(2005\)](#). As τ_d grows,

infinitely many delay-induced roots enter from $\text{Re } \lambda = -\infty$, but only the belief-loop pair reaches the imaginary axis. It does so at $\pm 0.48554i$ when $\tau_d = 1.3733$, matching ω^* and τ_d^* of Eq. (7) to five decimals. The spectral abscissa in panel (b) crosses zero exactly at $\tau_d^*(\kappa)$ for each κ . Faster learning widens the stability margin at small delays, but that margin is then consumed more steeply as delay increases.




Fig. 2 (a) Rightmost characteristic roots of Eq. (5) for $\tau_r = 0$ and $\kappa = 1$ as τ_d increases from 0 to 2.6 (color), computed by pseudospectral collocation. Stars mark the crossing $\pm i\omega^*$ predicted by Theorem 3. (b) Spectral abscissa $\max \text{Re } \lambda$ against τ_d for three learning speeds; circles mark $\tau_d^*(\kappa)$ from Eq. (7).

Proposition 5 (Direction of the Hopf bifurcation) *Write the linearization of Eq. (3) as $\dot{X} = A_0 X(t) + A_r X(t - \tau_r) + A_d X(t - \tau_d)$, with $X = (u, v, \pi)^\top$ and characteristic matrix $\Delta(z) = zI - A_0 - A_r e^{-z\tau_r} - A_d e^{-z\tau_d}$, where*

$$A_0 = \begin{pmatrix} 0 & 0 & \alpha x^*(1 - x^*) \\ -\delta y^*(1 - y^*) & 0 & 0 \\ 0 & 0 & -\kappa \end{pmatrix},$$

$$A_r = \begin{pmatrix} -b & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad A_d = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & \kappa & 0 \end{pmatrix}.$$

At a crossing $\lambda = i\omega$, let $\Delta(i\omega)q = 0$ and $r\Delta'(i\omega)q = 1$. For a vector $w \in \mathbb{C}^3$ and an exponent z , let $\hat{w}(z) = (w_1, w_2, w_3, w_1 e^{-z\tau_r}, w_2 e^{-z\tau_d})$ collect current and delayed components $(\cdot, \cdot, \cdot, \cdot^r, \cdot^d)$. The second- and third-order terms of Eq. (3) at E^* are the symmetric forms

$$B_1(u, w) = (1 - 2x^*)[\alpha(u_1 w_3 + u_3 w_1) - \beta(u_1 w_1^r + u_1^r w_1)],$$

$$B_2(u, w) = -\delta(1 - 2y^*)(u_1 w_2 + u_2 w_1),$$

$$C_1(u, w, s) = 2\beta S[u_1 w_1 s_1^r] - 2\alpha S[u_1 w_1 s_3],$$

$$C_2(u, w, s) = 2\delta S[u_1 w_2 s_2],$$

with $B_3 = C_3 = 0$, where $S[\cdot]$ denotes the sum over the three placements of the distinguished factor. Let $\phi = \hat{q}(i\omega)$, $h_{20} = \hat{H}_{20}(2i\omega)$ with $H_{20} = \Delta(2i\omega)^{-1}B(\phi, \phi)$, and $h_{11} = \hat{H}_{11}(0)$ with $H_{11} = \Delta(0)^{-1}B(\phi, \bar{\phi})$. Then the first Lyapunov coefficient is $\ell_1 = \text{Re } c_1/\omega$ with

$$c_1 = \frac{1}{2} r [C(\phi, \phi, \bar{\phi}) + B(\bar{\phi}, h_{20}) + 2B(\phi, h_{11})].$$

If $\ell_1 < 0$, the bifurcation is supercritical: orbitally stable periodic solutions exist for $\tau_d > \tau_d^*$, and the x -component has half-range $\approx 2|q_1| \sqrt{-\mu \text{Re } \lambda'(\tau_d^*)/\text{Re } c_1}$ for $\mu = \tau_d - \tau_d^*$.

Proof This is the center-manifold reduction of Faria and Magalhães [Faria and Magalhães \(1995\)](#) written in terms of the characteristic matrix, which reduces to Kuznetsov's formula for Ordinary Differential Equations (ODEs) when all delays vanish [Kuznetsov \(2004\)](#). Since the nonlinear terms depend on the history only through $x(t - \tau_r)$ and $y(t - \tau_d)$, and the relevant functions are $\phi(\theta) = qe^{i\omega\theta}$, $h_{20}(\theta) = H_{20}e^{2i\omega\theta}$ and $h_{11}(\theta) = H_{11}$, all forms reduce to the evaluations $\hat{w}(z)$. The amplitude follows from the truncated normal form $\dot{\zeta} = (i\omega + \mu\lambda'(\tau_d^*))\zeta + c_1\zeta|\zeta|^2$ and $u \approx \zeta q + \bar{\zeta}\bar{q}$. $\square$

For the baseline of Section 5.5, $\ell_1 = -0.468$, $\text{Re } c_1 = -0.227$ and $\text{Re } \lambda'(\tau_d^*) = 0.081$. The coefficient ℓ_1 remains negative:

- for all $\kappa \in [10^{-2}, 10^2]$ and $m \in \{0.3, 0.6, 1.2\}$ (Fig. 10a);
- at every sampled point of the two-delay stability boundary, where it ranges from -1.52 to -0.30 .

The bifurcation is therefore supercritical throughout the ranges examined. A sign proof valid for all parameters is not attempted.

Theorem 6 (Frozen-attacker benchmark) *If attackers do not adapt ($g = \delta = 0$, $y \equiv \bar{p}$), then $p(t) \rightarrow \bar{p}$ and, on the invariant set $\{y = p = \bar{p}\}$, Eq. (3) reduces to $\dot{x} = x(1 - x)[\alpha\bar{p} - \beta x(t - \tau) - C_p]$ with $\tau = \tau_r$ and interior equilibrium*

$$x^* = \frac{\alpha\bar{p} - C_p}{\beta}. \quad (8)$$

Its local stability is lost at the classical threshold

$$\tau^* = \frac{\pi}{2\beta x^*(1 - x^*)}, \quad (9)$$

where a Hopf bifurcation occurs.

Proof With $m = 0$, Eq. (5) factorizes as $\lambda(\lambda + \kappa)(\lambda + be^{-\lambda\tau_r})$; the root $\lambda = 0$ is the neutral direction along the continuum $y \equiv \text{const}$, $\lambda = -\kappa$ is the filter mode, and $\lambda + be^{-\lambda\tau_r} = 0$ is Hayes' equation [Hayes \(1950\)](#), whose roots leave the left half-plane at $\tau_r = \pi/(2b)$ with positive speed. $\square$

Theorem 6 is the known threshold of delayed two-strategy replicator dynamics [Alboszta and Miękisz \(2004\)](#); [Wesson et al. \(2016\)](#); [Yi and Zuwang \(1997\)](#); we retain it only as the benchmark against which the coupled results are measured.

4.3 Operational Implications

Theorem 3 and Corollary 4 give three quantitative design rules for a SOC: (1) a *latency budget*, $\tau_d < \tau_d^*(\kappa)$, beyond which hunting effort and attacker aggressiveness enter self-sustained cycles of period $\approx 2\pi/\omega^*$, i.e. recurrent breach waves; (2) a *minimum learning speed* $\kappa_{\min}$, below which no reduction of detection delay helps; and (3) an *investment trade-off*: raising κ (coverage N , alert discriminability $\lambda_A - \lambda_S$) buys delay tolerance, but with saturation at the instantaneous-learning limit of Corollary 4(b).

In what sense a tipping point. We use the term for bifurcation-induced tipping [Ashwin, Wieczorek, Vitolo, and Cox \(2012\)](#): a parameter threshold at which the attracting state changes qualitatively, here from a stable equilibrium to self-sustained oscillation. Because the Hopf bifurcation is supercritical (Proposition 5), the transition is non-hysteretic. Bringing τ_d back below τ_d^* restores stability, unlike the irreversible transitions caused by saddle-node or subcritical bifurcations. The transition is nevertheless operationally abrupt. By Proposition 5, the half-range of x grows as the square root of the excess delay, approximately $0.72\sqrt{\tau_d - \tau_d^*}$ at the baseline, so its slope is unbounded at the threshold. Exceeding the threshold by only 1% therefore already produces swings of about ± 0.08 in the hunting share. As the threshold is approached, the system also shows the generic early-warning signals of critical slowing down [Scheffer et al. \(2009\)](#) (Section 5.5).

Scaling and calibration. Substituting $\lambda = bs$ in Eq. (5) shows that stability depends only on the dimensionless groups $b\tau_d$, $b\tau_r$, κ/b and m/b^2 . The natural time unit is therefore the congestion relaxation time $1/b$, and all delays in Section 5 are expressed in this unit. Calibrating the model requires three estimates:

- b , from the decay of analyst backlog after a hunting surge;
- κ , from alert rates and coverage via Section 3.2;
- m , from the rate at which campaign tactics shift in response to hunting pressure.

Section 5.6 calibrates the learning speed κ on real SOC telemetry. Section 5.7 estimates $1/b \approx 3$ –6 days and the distribution of τ_r from IT-service-management logs; m is not identified by public data and is treated by sensitivity analysis. At $1/b = 5.5$ days and the illustrative learning speed $\kappa = b$, the baseline threshold $\tau_d^* = 1.373/b$ corresponds to about 7.5 days, well below the 2025 global median dwell time of 14 days [Mandiant \(2026\)](#).

5 Numerical Validation

Sections 5.1–5.4 examine the frozen-attacker benchmark of Theorem 6 ($g = \delta = 0$, $y \equiv \bar{p}$, single delay $\tau \equiv \tau_r$); Section 5.5 examines the full coupled system (3) and verifies Theorem 3. Section 5.6 calibrates the learning speed on real SOC data, and Section 5.7 calibrates the remediation channel on incident-management logs.

5.1 Parameter Configuration

We calibrate the system parameters to represent a high-stakes critical infrastructure network. In the benchmark the posterior settles at the stationary attacker mix, $p(t) \rightarrow$

$\bar{p}$ (Theorem 6), which we take to reflect high confidence in an aggressive APT presence at the operational horizon of interest, setting $\bar{p} = 0.8$. The vulnerability penalty is set to $\alpha = 5.0$, and the proactive hunting cost is $C_p = 1.0$. The negative externality parameter is set to $\beta = 4.0$, representing the severe network congestion and SOC alert fatigue caused by widespread proactive isolation.

Using Eq. (8), the interior evolutionary equilibrium is calculated as:

$$x^* = \frac{5.0(0.8) - 1.0}{4.0} = 0.75$$

This indicates that at equilibrium, 75% of the subsystems should be engaged in proactive hunting. The damping coefficient $K = \beta x^*(1 - x^*)$ (the benchmark counterpart of b) evaluates to $K = 4.0 \times 0.75 \times 0.25 = 0.75$. Applying Theorem 6, the critical operational delay threshold is:

$$\tau^* = \frac{\pi}{2(0.75)} \approx 2.094$$

5.2 Simulation Results: The Tipping Point

All integrations in this section use the second-order Heun (explicit trapezoidal) scheme on a grid with $\tau/dt \in \mathbb{N}$, so that delayed values are read exactly at grid points, and a constant initial history. We simulate two latency scenarios: an agile defense regime ($\tau < \tau^*$) and a delayed remediation regime ($\tau > \tau^*$).




Fig. 3 Time series evolution of the proactive defense strategy $x(t)$. When the operational delay ($\tau = 1.0$) is below the critical threshold τ^* , the ecosystem asymptotically converges to the interior equilibrium $x^* = 0.75$. When the delay exceeds the threshold ($\tau = 3.0$), the system destabilizes into sustained oscillations.

As depicted in Figure 3, when the operational delay is tightly controlled ($\tau = 1.0$), the strategic population smoothly dampens initial perturbations and converges to

x^* . Conversely, when the delay expands past the tipping point ($\tau = 3.0$), asymptotic stability is lost. The delayed feedback loop causes defenders to overcommit to proactive hunting long after the threat footprint has shifted, followed by a drastic under-commitment that leaves the system highly vulnerable.




Fig. 4 Delay-plane projection $(x(t), x(t - \tau))$ of the trajectory for $\tau = 3.0 > \tau^*$, after discarding the initial transient. The trajectory settles onto a closed orbit encircling the unstable equilibrium point x^* .

Figure 4 illustrates the qualitative signature of the delay-induced instability: since the DDE state space is infinite-dimensional, this panel shows the projection of the trajectory onto the current-state/delayed-state plane $(x(t), x(t - \tau))$ rather than a true phase portrait. For $\tau = 3.0$, the projected trajectory converges to a closed orbit around the unstable equilibrium x^* , numerically consistent with the periodic solution predicted by Theorem 6.

Operationally, this limit cycle is consistent with the cyclical failure observed in organizations suffering from extended APT dwell times: an initial breach triggers a massive, delayed remediation effort which depletes resources and induces alert fatigue, forcing a subsequent relaxation of security postures that paves the way for the attacker’s secondary lateral movement [Dash and Acken \(2026\)](#).

Furthermore, to evaluate the sensitivity of the system to the operational delay, we swept τ over a fine grid, integrated the unmodified Delay Differential Equation (DDE) with a sufficiently small step size, discarded an initial transient, and extracted the *local* extrema of the resulting steady-state trajectory using a peak-detection procedure (rather than the global min/max of the tail, which would conflate transient behavior with the true oscillation amplitude near τ^*).

As shown in Figure 5, the steady-state bounds of the population strategy are plotted against varying values of τ . For $\tau < 2.094$, the local maxima and minima




Fig. 5 Bifurcation diagram illustrating the steady-state amplitude of the proactive defense strategy against operational delay (τ). The solid lines represent the mean local maxima and minima of $x(t)$ over the retained steady-state window. For $\tau < \tau^*$, the amplitude collapses to the equilibrium $x^* = 0.75$. Beyond τ^* , the oscillation amplitude grows with the delay while remaining bounded within $(0, 1)$.

converge to the equilibrium $x^* = 0.75$. Immediately upon crossing τ^* , the oscillation amplitude grows continuously with the delay, consistent with the periodic solution predicted by Theorem 6, while remaining bounded within the $(0, 1)$ domain; we note that trajectories for τ close to τ^* exhibit critical slowing down, so a longer transient window is retained near the threshold to avoid conflating residual transient decay with genuine oscillation amplitude.

5.3 Numerical Robustness

To confirm that the results above are not artifacts of the discretization or of the assumed initial history, we perform two checks at $\tau = 3.0$: a step-size refinement and a sweep over the constant initial history x_0 . Table 2 reports the resulting post-transient amplitude (mean local maximum minus mean local minimum over $t \in [60, 150]$). It varies by less than 0.02% between $dt = 0.02$ and $dt = 0.0025$, and it is identical to four decimals across $x_0 \in \{0.2, 0.5, 0.85\}$. In the stable regime ($\tau = 1.0$), all three histories converge to $x^* = 0.75$ to six decimals by $t = 150$.

5.4 Sensitivity of the Tipping Point to Threat Belief

Because $\tau^* = \pi/[2\beta x^*(1 - x^*)]$ depends on $\bar{p}$ only through $x^* = (\alpha\bar{p} - C_p)/\beta$, the tipping point can be evaluated in closed form across the $(\bar{p}, \beta)$ plane without re-integrating the DDE. Figure 6 maps $\tau^*(\bar{p}, \beta)$ over the region where an interior equilibrium exists. For fixed β , the tolerance is lowest when $x^* = 1/2$, where $K = \beta x^*(1 - x^*)$ is maximal, and it grows without bound as $\bar{p}$ pushes x^* toward 0 or 1. For fixed $\bar{p}$, $K = a(1 - a/\beta)$ with $a = \alpha\bar{p} - C_p$ is increasing in β , so τ^* decreases monotonically in β toward the finite limit $\pi/(2a)$. Stronger congestion therefore always reduces delay tolerance in the benchmark.

Table 2 Sensitivity of the steady-state amplitude at $\tau = 3.0$ to step size dt and initial history x_0 .

Check	Varying parameter	Amplitude
Step size ($x_0 = 0.85$)	$dt = 0.02$	0.6771
	$dt = 0.01$	0.6771
	$dt = 0.005$	0.6771
	$dt = 0.0025$	0.6772
Initial history ($dt = 0.005$)	$x_0 = 0.20$	0.6771
	$x_0 = 0.50$	0.6771
	$x_0 = 0.85$	0.6771




Fig. 6 Critical delay threshold $\tau^*(\bar{p}, \beta)$ over the region admitting an interior equilibrium ($C_p < \alpha\bar{p} < C_p + \beta$), with $\alpha = 5.0$, $C_p = 1.0$ fixed. Darker colors denote smaller τ^* (the color scale is capped at the 90th percentile); white contours mark $\tau^* \in \{0.5, 1, 2, 4\}$; the blank region admits no interior equilibrium.

5.5 Coupled Attacker–Defender–Belief Dynamics

We now integrate the full system (3) with $\alpha = 5$, $\beta = 4$, $C_p = 1$ as above, $g = 1$, $\delta = 2$ and $\kappa = 1$. These give $x^* = 0.5$, $y^* = p^* = 0.6$, $b = 1$ and $m = 0.6$. Since $m < b^2$, $\kappa_{\min} = 0$, and Theorem 3 gives $\omega^* = 0.486$ and $\tau_d^* = 1.373$ (period ≈ 12.9). As in Section 5.2, we use the Heun scheme on a grid with $\tau/dt \in \mathbb{N}$ and constant initial histories; the scripts are supplied as supplementary material. Independently of the proof, the number of characteristic roots with $\text{Re } \lambda > 0$ was computed from the argument increment of $\Delta(i\omega)$ (Mikhailov criterion Stépán (1989)). It is 0 at $0.95 \tau_d^*$ and 2 at $1.05 \tau_d^*$ for every $\kappa \in \{0.2, 0.5, 1, 2, 5, 20, 10^3\}$ tested.




Fig. 7 Coupled dynamics with $\tau_r = 0$. (a) Hunting share x and aggressive-campaign share y for $\tau_d = 1.0 < \tau_d^*$ (solid, convergence to E^*) and $\tau_d = 1.6 > \tau_d^*$ (dashed, sustained attacker-defender cycles). (b) The same trajectories in the (y, x) plane after $t = 20$. (c) Steady-state range of x (last 300 time units of $t \in [0, 2000]$) against τ_d , with the normal-form prediction of Proposition 5 (solid). The small nonzero range just below τ_d^* is a slow transient.

Figure 7 confirms Theorem 3 and Proposition 5 in the nonlinear system. The oscillation emerges from zero amplitude at τ_d^* and follows the square-root law of a supercritical Hopf bifurcation, although steeply: its range reaches 0.38 at $\tau_d = 1.45$. At $\tau_d = 1.35$ (1.7% below τ_d^*), both a near-equilibrium history ($x_0 = 0.51, y_0 = p_0 = 0.6$) and a distant one ($x_0 = 0.95, y_0 = p_0 = 0.1$) decay to a range below 10^{-4} by $t = 5600$, so no coexisting large-amplitude cycle exists there.

Figure 8 gives the global picture.

- Panel (a) shows the three-dimensional geometry. Below the threshold, trajectories spiral into E^* . Above it, they settle on a tilted limit cycle along which belief p trails the attacker mix y : this is the mechanism of the recurrent breach waves.
- Panel (b) maps the peak aggressive share on the attractor over the (τ_d, κ) plane from 1,845 nonlinear simulations. The simulated regime agrees with Theorem 3 in 98% of grid cells lying more than 0.1 from the predicted boundary; the remaining cells are slow transients near the threshold.
- Beyond the threshold, attacker dominance escalates quickly. For $\kappa = 1$, the peak aggressive share rises from the equilibrium value 0.6 to 0.86 at $\tau_d = 1.6$ and to 0.98 at $\tau_d = 2.4$. Once the latency budget is exceeded, aggressive tactics periodically take over the campaign mix almost entirely, as the cycle approaches the boundary heteroclinic cycle of Lemma 2.

Figure 9(a) shows that τ_d^* increases with κ and saturates, matching both limits of Corollary 4. For $m = 0.6$ it runs from $0.667 = (b^2 - m)/(mb)$ as $\kappa \rightarrow 0$ to 2.044 as $\kappa \rightarrow \infty$. For $m = 1.2$, no detection delay is admissible when $\kappa \leq \kappa_{\min} = 0.2$.

Figure 9(b) quantifies Remark 2, with three findings.

Attacker adaptation shrinks the remediation budget. With $\tau_d = 0$, stability is lost at $\tau_r = 1.116$, which is 29% below the frozen-attacker value $\pi/(2b) = 1.571$.

The two delays are not additive. Along the stability boundary, the admissible detection delay rises from 1.373 ($\tau_r = 0$) to 1.465, 1.536, 1.574 and 1.533 for $\tau_r = 0.25, 0.5, 0.75$ and 1.0. A moderate remediation lag thus *enlarges* the admissible detection




Fig. 8 (a) Trajectories in the (x, y, p) phase space for $\tau_d = 1.0$ (spiral into E^* , star) and $\tau_d = 1.6$ (limit cycle), with their projections on the floor; color encodes time. (b) Peak aggressive share $\max_t y(t)$ over the last 200 time units of $t \in [0, 700]$ on a 41×45 grid of (τ_d, κ) (vectorized Heun integration, $dt = 0.025$, $\tau_r = 0$). Dashed: $\tau_d^*(\kappa)$ from Theorem 3.




Fig. 9 (a) Detection-delay tolerance τ_d^* (Theorem 3) against learning speed κ for three loop gains m . δ is varied with $g = \delta x^*$ so that $x^* = 0.5$ stays fixed; the dotted line marks $\kappa_{\min}$ of Corollary 4 for $m = 1.2 > b^2$. (b) Stability chart in the (τ_d, τ_r) plane for $\kappa = 1$. Shading shows the Mikhailov root count on a 121×106 grid; the colored curve is the exact crossing curve $\Delta(i\omega) = 0$, colored by the crossing frequency ω . Dashed: frozen-attacker threshold $\pi/(2b)$ of Theorem 6. Dotted: the boundary an additive model $\tau_d + \tau_r = \tau_d^*$ would predict. Cross and circle: $(1.45, 0)$ oscillates, $(1.45, 0.5)$ converges.

delay, and nonlinear integration confirms this: at $\tau_d = 1.45$ the system oscillates for $\tau_r = 0$ (range 0.38) but converges for $\tau_r = 0.5$.

The destabilizing mode changes continuously. The stability region is bounded by a single Hopf curve. Along it, the crossing frequency increases monotonically from 0.486 at the pure detection-delay end, a slow belief-loop oscillation, to 1.18 at the pure remediation-delay end, a fast congestion oscillation. The sharp turn near $(\tau_d, \tau_r) \approx (1.45, 1.10)$, where $\omega \approx 0.68$, is not a double-Hopf point, since a single pair of roots crosses there. $\ell_1 < 0$ at all sampled boundary points.

Aggregating both latencies into a single τ , as single-population models do, therefore mispredicts both the size and the shape of the stability region.




Fig. 10 (a) First Lyapunov coefficient at τ_d^* (Proposition 5) against κ for the three loop gains of Fig. 9(a); it is negative throughout. (b) Half-range of x against $\mu = \tau_d - \tau_d^*$: Heun integration ($dt = 0.005$, $t \in [0, 4000]$) versus the normal-form prediction. (c) Standard deviation of x around x^* ($\tau_d = 1.0$) with explicit Poisson alerts ($\lambda_A = 2$, $\lambda_S = 0.5$) and a Kalman–Bucy update at fixed $\kappa = 1$, against the number of monitored campaigns N . (d) Stochastic ($N = 10^4$) and mean-field trajectories for $\tau_d = 1.0$ and 1.6 .

Figure 10 validates the two approximations behind the analysis.

First, the normal-form amplitude agrees with simulation to within 5% for $\mu \in \{0.01, 0.02, 0.04\}$: the half-ranges are 0.076, 0.105 and 0.143 in simulation versus 0.072, 0.102 and 0.144 predicted. At $\mu = 0.005$ the gap is 8%, owing to slow convergence of the transient.

Second, replacing the mean-field belief law by the actual filter driven by Poisson counts leaves the qualitative picture unchanged. Fluctuations around E^* scale as $N^{-1/2}$: the standard deviation of x is 0.038, 0.012, 0.0034 and 0.0010 for $N = 10^3, \dots, 10^6$. Above τ_d^* the stochastic trajectories follow the mean-field cycle. For small coverage N and τ_d close to τ_d^* , the weakly damped focus converts alert noise into small irregular oscillations around E^* (Fig. 10d), whose statistics announce the threshold. With $N = 10^4$ and τ_d rising from 0.4 through 0.8, 1.0, 1.2 and 1.3 to 1.35:

- the standard deviation of x rises from 0.007 through 0.009, 0.011, 0.016 and 0.025 to 0.048;
- its lag-one autocorrelation (one time unit) rises from 0.72 to 0.88.

These are the signatures of critical slowing down [Scheffer et al. \(2009\)](#). In principle, they could be monitored in SOC telemetry as a leading indicator that a latency budget is being exhausted.

5.6 Semi-empirical Calibration on Microsoft GUIDE

Data. We use the public test split of GUIDE [Freitas et al. \(2024\)](#), a collection of real Microsoft Defender Extended Detection and Response (XDR) incidents graded by customer SOC analysts (CDLA-2.0 license). The split contains 4,147,992 evidence/alert rows from 4,400 organizations and 5,970 detectors. Because the triage grade is an incident-level label, we aggregate rows into 303,220 incidents keyed by (OrgId, IncidentId), each dated by its first alert. The grades are 62,540 true positives (TP) (20.6%), 148,215 benign positives (BP) (48.9%) and 92,465 false positives (FP) (30.5%), with 13.7 rows per incident on average.

First-alert dates span 22 November 2023 to 17 June 2024, but 97.7% of incidents fall in two complete weeks, 3–16 June 2024 (Fig. 11a). Full-span daily averages (e.g. 413 true positives per day, against a median of one) therefore mix a sparse tail with the dense sample, and we do not use them. We restrict the analysis to the two-week window. Within it, we use increments between consecutive weekdays only, because weekend volume is only about 40% of weekday volume.

Identification of κ . In the model, the defender observes the aggregate alert count with intensity $N[\lambda_0 + \lambda_S + (\lambda_A - \lambda_S)y(t - \tau_d)]$. Here benign and false positives form a background rate λ_0 that does not depend on y : they are not produced by stealthy campaigns, so they add noise but no information. Let r_t denote true-positive incidents per active organization-day. Under the random-walk prior of Section 3.2, day-to-day increments of r_t have variance $(\lambda_A - \lambda_S)^2\sigma_y^2$ plus Poisson sampling noise $r_t/O_t + r_{t-1}/O_{t-1}$, where O_t is the number of active organizations. Combined with $\sigma_N^2 = \bar{\lambda}/[N(\lambda_A - \lambda_S)^2]$, this gives

$$\hat{\kappa}(N) = \sqrt{\widehat{\text{Var}}(\Delta r) - \widehat{\mathbb{E}}[\text{Poisson variance}]} \sqrt{N/\bar{\lambda}}, \quad (10)$$

where $\bar{\lambda}$ is the total (pre-triage) incident rate per organization-day; this extends the definition of Section 3.2 by the background rate λ_0 . The unobservable quantities λ_A , λ_S and y^* cancel. Three caveats apply:

- Increments also reflect changes in which organizations are active, so Eq. (10) is an upper bound on attacker-mix volatility.
- Using the pre-triage rate $\bar{\lambda}$ is conservative.
- With eight weekday increments, the uncertainty is substantial; we report 95% bootstrap Confidence Intervals (CI) (20,000 resamples).

Table 3 shows that a single organization learns the attacker mix on a time scale $1/\kappa \approx 14$ days, whereas pooled telemetry across organizations learns it in about 8 hours. Pooling therefore raises κ by the factor $\sqrt{N} \approx 43$.

Two interpretive assumptions underlie this comparison.

Table 3 Belief-channel quantities estimated from GUIDE (3–16 June 2024, weekdays).

Quantity	Estimate [95% CI]
Total incident rate $\bar{\lambda}$ (per organization-day)	13.43
True-positive rate $\bar{r}$ (per organization-day)	2.66
Active organizations per weekday, N	1,890
$\text{sd}(\Delta r)$: raw / Poisson component	0.266 / 0.053
$(\lambda_A - \lambda_S)\sigma_y$ [$\text{day}^{-3/2}$]	0.261 [0.077, 0.420]
κ , single organization ($N = 1$) [day^{-1}]	0.071 [0.021, 0.115]
κ , pooled telemetry ($N \approx 1,890$) [day^{-1}]	3.09 [0.92, 4.98]

Organizations as sensors. Treating each active organization as an independent sensor of a common attacker population presumes that organizations face a shared campaign mix. Sector-specific targeting reduces the effective N to the organizations exposed to the same mix, so the pooled value is an upper bound.

Accuracy of the mean-field law. The stationary tracking error of the filter has standard deviation $\sigma_y/\sqrt{2\kappa} = \hat{\sigma}_{\text{ex}}/[(\lambda_A - \lambda_S)\sqrt{2\kappa}]$, where $\hat{\sigma}_{\text{ex}} = 0.261$ is the excess volatility estimated above. It is therefore $\sqrt{43.5} \approx 6.6$ times larger for a single organization than for pooled telemetry. If, for instance, all true positives are attributed to aggressive campaigns, so that $\lambda_A - \lambda_S \approx \bar{r}/y^* \approx 4.4$ at $y^* = 0.6$, the error is about 0.16 for $N = 1$ but 0.024 when pooled. The deterministic threshold is thus accurate for pooled telemetry, but for a single organization it is only a first-order guide, around which alert noise generates sizeable quasi-cycles (Fig. 10).

GUIDE does not identify b or m . It contains no triage-completion timestamps and no attacker-side data. The within-incident alert span (median 0, mean 1.06 hours) measures alert clustering, not remediation, and the lag-1 autocorrelation of daily volume over the full span (0.91) reflects the sampling step rather than congestion dynamics. We therefore take b from the ITSM logs of Section 5.7 ($1/b \approx 3$ –6 days, shaded in Fig. 11b), show the full range $1/b \in [1, 28]$ days for context, treat $m/b^2 \in \{0.6, 1.2\}$ as scenarios, and compute τ_d^* in days from Theorem 3 in the dimensionless form of Section 4.3.

Figure 11(b) gives five results. Dwell time, the interval from compromise to detection of an intrusion, bounds from above the time it takes an intrusion to register in the alert stream. We therefore use the 14-day median as a conservative proxy for τ_d .

- With $1/b = 7$ days and $m/b^2 = 0.6$, the tolerance is 7.8 days for a single organization and 14.0 days with pooled telemetry.
- Over $1/b \in [1, 28]$ days, pooling multiplies the admissible detection delay by 1.3–2.4 when $m/b^2 = 0.6$.
- The effect is much larger when attackers adapt strongly. For $m/b^2 = 1.2$ and $1/b = 7$ days, the tolerance is 1.2 days for a single organization versus 6.3 days with pooling. For $1/b < 2.8$ days, a single organization falls below the minimum learning speed $\kappa_{\min}$ of Corollary 4(a) and cannot be stabilized at any detection delay.




Fig. 11 (a) Daily incident arrivals in the GUIDE test split by first-alert date (log scale). The shaded two-week window contains 97.7% of incidents and is used for calibration. (b) Detection-delay tolerance τ_d^* in days against the congestion relaxation time $1/b$, for the calibrated learning speeds of a single organization and of pooled telemetry (bands: 95% bootstrap intervals for κ). Solid: $m/b^2 = 0.6$; dashed: $m/b^2 = 1.2$. Dotted: 2025 global median dwell time [Mandiant \(2026\)](#).

- The pooled value is already close to the instantaneous-learning ceiling $2.04/b$ of Corollary 4(b). Beyond this level of information sharing, further gains must come from shorter delays or stronger congestion damping, not from faster learning.
- Measured against the 2025 global median dwell time of 14 days, a single organization with $m/b^2 = 0.6$ stays within its tolerance only if $1/b \gtrsim 11$ days, whereas pooled telemetry does so for $1/b \gtrsim 7$ days.

Because Theorem 3 assumes $\tau_r = 0$, these budgets concern the detection channel alone. The remediation channel imposes a separate ceiling. For the calibrated learning speeds, the remediation-only threshold is $\tau_r^* \approx 1.08/b$ (pooled) to $1.32/b$ (single organization), i.e. about 6–7 days at the measured $1/b \approx 5.5$ days. Section 5.7 shows that observed incident queues operate close to this ceiling once their long tail is taken into account. Reported mean containment times of about two months [IBM Security and Ponemon Institute \(2025\)](#) lie far beyond it. In practice, the remediation delay may therefore be the binding constraint. These statements remain conditional on the scenario values of m ; they illustrate how the theory converts measurable telemetry into latency budgets rather than providing a validated forecast.

5.7 Remediation Delay and Congestion from ITSM Process-Mining Logs

GUIDE identifies the information channel but not the remediation channel. For the latter, we use two public IT Service Management (ITSM) event logs that record the full life cycle of real incidents:

- the ServiceNow incident log of an IT company [Amaral, Fantinato, and Peres \(2018\)](#): 24,918 incidents, 2016–2017;
- the Volvo IT incident and problem-management logs of Business Process Intelligence Challenge (BPIC) 2013 [Steeman \(2013a, 2013b\)](#).

These are service-desk rather than security incidents. We use them as a proxy for the analyst-queue dynamics that follow a hunting decision, not as a description of SOC workflows.

Estimation. For each incident, the remediation delay is the time from opening to resolution. Administrative closure is excluded: in the ServiceNow log it follows resolution by a median of 5.0 days. As with GUIDE, first-event dates contain a sparse early tail. We therefore restrict each log to incidents whose opening falls between the 10th and 90th percentiles of opening dates.

We estimate b as the relaxation rate of the backlog $B(t)$ of open incidents, using two estimators:

- Little’s law, $b_L = \lambda/\bar{B} = 1/\bar{T}$, where λ is the arrival rate and $\bar{T} = \bar{B}/\lambda$ the mean time in system over the calibration window (it differs from the mean resolution time in Table 4, which is computed per incident over all incidents opened in the window);
- the dynamic estimator $b_{AR} = -\ln \rho_1$, where ρ_1 is the lag-one autocorrelation of the daily backlog after weekday effects and a 28-day trend are removed.

The dynamic estimator requires high volume. We apply it only when $\lambda \geq 5$ incidents per day. Confidence intervals come from a moving-block bootstrap with 7-day blocks.

Table 4 Remediation delay and congestion rate from ITSM logs (calibration windows as in Fig. 12b).

	ServiceNow (UCI)	BPIC 2013 incidents	BPIC 2013 problems
Incidents in window	19,106	5,529	1,188
Resolution time: median / mean / p90 [days]	0.87 / 7.08 / 14.98	0.25 / 3.21 / 11.07	80.2 / 127.9 / 347.4
Share resolved within one day	0.52	0.59	0.08
Arrivals λ [per day] / mean backlog $\bar{B}$	258.8 / 1,421	425.0 / 1,417	2.1 / 387
b_L [day^{-1}] (95% CI)	0.182 (0.158, 0.195)	0.30 (indicative)	0.0054
b_{AR} [day^{-1}] (95% CI)	0.191 (0.131, 0.394)	—	—
$b \tau_r$ with $\tau_r = \text{median} / \text{mean}$	0.17 / 1.35	0.07 / 0.96	0.43 / 0.68

Results. Table 4 and Fig. 12 give four findings.

The congestion time scale is short and well identified in the high-volume log. In the ServiceNow log, the two independent estimators agree ($b_L = 0.182$, $b_{AR} = 0.191$ per day). Little’s law is robust to excluding the first two weeks of warm-up (0.187); the dynamic estimator is noisier (0.32 in that case, within its interval). The congestion relaxation time is therefore $1/b \approx 5.2$ –5.5 days. For BPIC 2013 the window covers only 16 days, during which the backlog first builds up and then drains (Fig. 12b). Its value $1/b \approx 3.3$ days is therefore only indicative. Taken together, the logs place $1/b$ between about 3 and 6 days, which is the band shaded in Fig. 11(b).

Remediation delays are extremely right-skewed. More than half of all incidents are resolved within a day, but the mean exceeds the median by a factor of 8–13 and the 90th percentile lies at 11–15 days. Root-cause (problem) remediation is two orders of magnitude slower (median 80 days).




Fig. 12 (a) Empirical Cumulative Distribution Function (CDF) of incident resolution times (log scale); dotted lines mark medians, dash-dot lines means. (b) Daily backlog of open incidents, aligned at the start of each calibration window (shaded). (c) Dimensionless remediation delay $b\tau_r$, using b_{AR} for ServiceNow and b_L for BPIC 2013, against the remediation-only threshold $b\tau_r^* \approx 1.08$ – 1.32 of Section 5.6.

The remediation channel operates close to its stability boundary once the tail is taken into account. With the median as the representative delay, $b\tau_r \approx 0.07$ – 0.17 , far inside the threshold 1.08 – 1.32 . With the mean, $b\tau_r \approx 0.96$ (BPIC 2013) and 1.35 (ServiceNow, with b_{AR} ; interval 0.93 – 2.79), which straddles the threshold. Because $b_L = 1/\bar{T}$ is itself built from the mean sojourn time, a mean-based comparison with b_L is not an independent test. That is why the ServiceNow bar in Fig. 12(c) uses the dynamic estimator. The substantive conclusion is structural. A queue whose congestion signal lags by its own mean sojourn time sits at $b\tau_r \approx 1$, only 7–24% below the Hopf threshold of the point-delay model. Stability of the remediation channel therefore depends on how quickly the long tail of slow incidents is cleared, not on the typical case.

Consequences for the detection-delay budget. At the measured $1/b \approx 5.5$ days and $m/b^2 = 0.6$, Theorem 3 gives $\tau_d^* \approx 5.7$ days for a single organization and 10.9 days with pooled telemetry (Fig. 11b). Both are below the 2025 global median dwell time of 14 days Mandiant (2026). Over the measured band $1/b \in [3.3, 6.3]$ days, the tolerance ranges from 3.0 to 6.8 days for a single organization and from 6.5 to 12.6 days with pooling. Under the scenario $m/b^2 = 0.6$, therefore, the median intrusion dwells longer than the belief loop can tolerate even with pooled telemetry. Closing that gap requires faster detection, not only faster learning.

Two caveats apply. First, the logs describe IT service desks; security incident response involves forensic work and may be slower, which would lengthen τ_r . Second, the heavy tails mean that a point delay is a simplification. Replacing $\beta x(t - \tau_r)$ by a distributed delay whose kernel is fitted to Fig. 12(a) is the natural extension; the ITSM logs already provide the data to fit it.

6 Conclusion

APT campaigns exploit the time defenders need to see them and to act. We modeled this as a two-population evolutionary game in which defenders filter a delayed Poisson alert stream to track a strategically adapting attacker population, with separate detection and remediation delays.

The analysis gives four results.

- A closed-form detection-delay threshold $\tau_d^*(\kappa)$ at which the interior Nash equilibrium undergoes a Hopf bifurcation into recurrent attacker–defender cycles; this bifurcation is supercritical throughout the ranges examined.
- A minimum learning speed $\kappa_{\min}$ below which no latency reduction stabilizes defense.
- Explicit limits of the delay tolerance for slow and instantaneous learning.
- Numerical evidence that attacker adaptation shrinks the remediation budget and that detection and remediation delays interact non-additively. In particular, a moderate remediation lag can widen the admissible detection delay.

On real SOC telemetry (Microsoft GUIDE), the learning speed of a single organization is about 0.07 per day, whereas telemetry pooled across organizations reaches about 3 per day, close to the instantaneous-learning ceiling. Information sharing therefore buys delay tolerance, but only up to that ceiling. Incident-management logs place the congestion relaxation time at three to six days and show remediation delays so right-skewed that incident queues operate close to the remediation threshold once the tail is counted. At the measured congestion time scale, even pooled telemetry tolerates less detection delay than the 14-day median dwell time. The classical threshold $\pi/(2\beta x^*(1-x^*))$ of delayed replicator dynamics is recovered only in the frozen-attacker limit.

Operationally, the results suggest that SOC performance should be managed as a joint budget over learning speed, detection delay and remediation delay, rather than as a single mean-time-to-respond target.

Five extensions are natural:

- identification of m , which requires paired defender and attacker time series, and re-estimation of b and τ_r on security-specific (rather than IT-service) incident workflows;
- a distributed-delay version of the congestion term, with kernel fitted to observed resolution-time distributions;
- a proof of the sign of ℓ_1 over the whole parameter space;
- attackers that also observe defenders with delay, which adds a third delay;
- finite-population and networked versions, in which alert noise and topology interact with the thresholds derived here.

Statements and Declarations

Funding: The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

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

Author Contributions: The author confirms sole responsibility for the following: study conception and design, data collection, mathematical analysis, interpretation of results, and manuscript preparation.

Data Availability: All data are public: GUIDE on Kaggle (CDLA-2.0) [Freitas et al. \(2024\)](#), the ServiceNow incident log in the UCI Machine Learning Repository (CC BY 4.0) [Amaral et al. \(2018\)](#), and the BPIC 2013 logs on 4TU.ResearchData [Steeman \(2013a, 2013b\)](#).

Code Availability: All scripts reproducing every figure and Tables 2–4 (`fig4_heatmap.py`, `coupled_analysis.py`, `fig5_fig6_coupled.py`, `guide_calibration.py`, `itsm_calibration.py`) are provided as supplementary material.

References

- Alboszta, J., & Miękisz, J. (2004). Stability of evolutionarily stable strategies in discrete replicator dynamics with time delay. *Journal of Theoretical Biology*, *231*(2), 175–179. 10.1016/j.jtbi.2004.06.012
- Amaral, C., Fantinato, M., & Peres, S. (2018). *Incident management process enriched event log*. UCI Machine Learning Repository. 10.24432/C57S4H
- Ashwin, P., Wieczorek, S., Vitolo, R., & Cox, P. (2012). Tipping points in open systems: Bifurcation, noise-induced and rate-dependent examples in the climate system. *Philosophical Transactions of the Royal Society A*, *370*(1962), 1166–1184. 10.1098/rsta.2011.0306
- Bashir, A., Shamszaman, Z. U., Song, Z., & Han, T. A. (2026). Co-evolutionary dynamics of attack and defence in cybersecurity. *Knowledge-Based Systems*. 10.1016/j.knosys.2026.115750
- Breda, D., Maset, S., & Vermiglio, R. (2005). Pseudospectral differencing methods for characteristic roots of delay differential equations. *SIAM Journal on Scientific Computing*, *27*(2), 482–495.
- Cooke, K. L., & van den Driessche, P. (1986). On zeroes of some transcendental equations. *Funkcialaj Ekvacioj*, *29*, 77–90.
- Dash, S., & Acken, J. M. (2026). Intrusion detection latency: the neglected metric. *Cybersecurity*. 10.1186/s42400-026-00574-7
- Davies, S. R., & MacFarlane, R. (2026). Comprehensive performance benchmarking and comparative analysis of active ransomware threats. *Computers and Electrical Engineering*. 10.1016/j.compeleceng.2026.110963
- Dragicevic, A. Z. (2024). The unification of evolutionary dynamics through the

- Bayesian decay factor in a game on a graph. *Bulletin of Mathematical Biology*. 10.1007/s11538-024-01299-9
- Faria, T., & Magalhães, L. T. (1995). Normal forms for retarded functional differential equations with parameters and applications to Hopf bifurcation. *Journal of Differential Equations*, 122(2), 181–200. 10.1006/jdeq.1995.1144
- Freitas, S., Kalajdjieski, J., Gharib, A., & McCann, R. (2024). *AI-driven guided response for security operation centers with Microsoft Copilot for security*. arXiv:2407.09017. (GUIDE dataset: <https://www.kaggle.com/datasets/Microsoft/microsoft-security-incident-prediction>) 10.48550/arXiv.2407.09017
- Hale, J. K., & Verduyn Lunel, S. M. (1993). *Introduction to functional differential equations*. New York: Springer-Verlag. 10.1007/978-1-4612-4342-7
- Hayes, N. D. (1950). Roots of the transcendental equation associated with a certain difference-differential equation. *Journal of the London Mathematical Society*, s1-25(3), 226–232. 10.1112/jlms/s1-25.3.226
- Hofbauer, J., & Sigmund, K. (1998). *Evolutionary games and population dynamics*. Cambridge: Cambridge University Press. 10.1017/CBO9781139173179
- Huang, L., & Zhu, Q. (2020). A dynamic games approach to proactive defense strategies against advanced persistent threats in cyber-physical systems. *Computers and Security*. 10.1016/j.cose.2019.101660
- IBM Security, & Ponemon Institute. (2025). *Cost of a data breach report 2025* (Tech. Rep.). IBM Corporation. (Covers 600 organizations breached between March 2024 and February 2025)
- Iijima, R. (2011). Heterogeneous information lags and evolutionary stability. *Mathematical Social Sciences*. 10.1016/j.mathsocsci.2010.12.002
- Information Commissioner’s Office. (2020, October). *Penalty notice: Marriott international inc.* (Tech. Rep.). Wilmslow, UK: ICO.
- Kalman, R. E., & Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1), 95–108. 10.1115/1.3658902
- Kumar, P., Rai, P., & Mishra, B. K. (2026). Securing medical iot devices: A time-delay mathematical approach for cyber-attack detection and mitigation. *Boletim da Sociedade Paranaense de Matematica*. 10.5269/bspm.82697
- Kuznetsov, Y. A. (2004). *Elements of applied bifurcation theory* (3rd ed.). New York: Springer. 10.1007/978-1-4757-3978-7
- Mandiant. (2026, March). *M-Trends 2026: Data, insights, and strategies from the*

- frontlines* (Tech. Rep.). Google Cloud. Retrieved from <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026/> (Published 24 March 2026)
- Mangalindan, D. H., Hota, A. R., & Srivastava, V. (2025). On logit dynamics for multiplayer trust game. In *Proceedings of the american control conference*. 10.23919/ACC63710.2025.11107502
- Mohamadichamgavi, J., & Bodnar, M. (2025). Bifurcation analysis of replicator dynamics with logistic growth and strategy-dependent time delays in snowdrift game. *Dynamic Games and Applications*. 10.1007/s13235-025-00671-1
- Napel, G. T., van Eeten, M., & Parkin, S. (2025). Speedrunning the maze: Meeting regulatory patching deadlines in a large enterprise environment. In *Proceedings - ieee symposium on security and privacy*. 10.1109/SP61157.2025.00081
- Nithya, D., Madhusudanan, V., Murthy, B. S. N., et al. (2024). Delayed dynamics analysis of sei2rs malware propagation models in cyber-physical systems. *Computer Networks*. 10.1016/j.comnet.2024.110481
- Njoku, J. (2026). Asymmetric partially observable markov decision process (pomdp)-based stochastic games for modeling stealthy apt propagation. *Journal of Computer Virology and Hacking Techniques*. 10.1007/s11416-026-00638-9
- Pathiraja, S., & Wacker, P. (2025). Connections between sequential bayesian inference and evolutionary dynamics. *Philosophical Transactions of the Royal Society A*. 10.1098/rsta.2024.0241
- Patra, A., Das Bairagya, J., & Chakraborty, S. (2025). Bayesian ecoevolutionary game dynamics. *Physical Review E*, 111, 044401. 10.1103/PhysRevE.111.044401
- Patra, A., Sengupta, S., Paul, A., & Chakraborty, S. (2024). Inferring to cooperate: Evolutionary games with Bayesian inferential strategies. *New Journal of Physics*. 10.1088/1367-2630/ad4e5e
- Qin, Y., Yang, X., Yang, L.-X., & Huang, K. (2024). Modeling and study of defense outsourcing against advanced persistent threat through impulsive differential game approach. *Computers & Security*, 141, 104003. 10.1016/j.cose.2024.104003
- Saint-Hilaire, K. A., Neal, C., Cuppens, F., et al. (2025). A real-time automated attack-defense graph generation approach. *Journal of Information Security and Applications*. 10.1016/j.jisa.2025.104266
- Sawa, R., & Wu, J. (2023). Statistical inference in evolutionary dynamics. *Games and Economic Behavior*. 10.1016/j.geb.2022.11.008

- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., ... Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature*, 461, 53–59. 10.1038/nature08227
- Schlag, K. H. (1998). Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78(1), 130–156. 10.1006/jeth.1997.2347
- Selten, R. (1980). A note on evolutionarily stable strategies in asymmetric animal conflicts. *Journal of Theoretical Biology*, 84(1), 93–101. 10.1016/S0022-5193(80)81038-1
- Snyder, D. L. (1972). Filtering and detection for doubly stochastic Poisson processes. *IEEE Transactions on Information Theory*, 18(1), 91–102. 10.1109/TIT.1972.1054756
- Steeman, W. (2013a). *Bpi challenge 2013, closed problems*. 4TU.ResearchData. Retrieved from <https://doi.org/10.4121/uuid:c2c3b154-ab26-4b31-a0e8-8f2350ddac11> 10.4121/uuid:c2c3b154-ab26-4b31-a0e8-8f2350ddac11
- Steeman, W. (2013b). *Bpi challenge 2013, incidents*. 4TU.ResearchData. Retrieved from <https://doi.org/10.4121/uuid:500573e6-ac9e-4b0f-802c-47aa80172e27> 10.4121/uuid:500573e6-ac9e-4b0f-802c-47aa80172e27
- Stépán, G. (1989). *Retarded dynamical systems: Stability and characteristic functions*. Harlow: Longman Scientific & Technical.
- Tan, J., Jin, H., Zhang, H., et al. (2023). A survey: When moving target defense meets game theory. *Computer Science Review*, 50, 100544. 10.1016/j.cosrev.2023.100544
- Tian, M., Dong, Z., & Wang, X. (2021). Analysis of false data injection attacks in power systems: A dynamic bayesian game-theoretic approach. *ISA Transactions*. 10.1016/j.isatra.2021.01.011
- U.S. Government Accountability Office. (2018). *Data protection: Actions taken by Equifax and federal agencies in response to the 2017 breach* (Tech. Rep. No. GAO-18-559). GAO.
- U.S. Senate Committee on Commerce, Science, and Transportation. (2014, March). *A “kill chain” analysis of the 2013 Target data breach* (Tech. Rep.). Majority Staff Report.
- Wan, Z., Cho, J.-H., Zhu, M., & Singh, M. P. (2022). Foureye: Defensive deception against advanced persistent threats via hypergame theory. *IEEE Transactions on Network and Service Management*, 19(1), 370–384. 10.1109/TNSM.2021.3117698

- Wang, Z., Li, J., Wang, Y., & Meng, W. (2023). Optimal repair strategy against advanced persistent threats under time-varying networks. *IEEE Transactions on Information Forensics and Security*, 18, 5844–5857. 10.1109/TIFS.2023.3318954
- Wesson, E., Rand, R., & Rand, D. (2016). Hopf bifurcations in two-strategy delayed replicator dynamics. *International Journal of Bifurcation and Chaos*, 26(1), 1650006. 10.1142/S0218127416500061
- Wettergren, T. A. (2023). Replicator dynamics of evolutionary games with different delays on costs and benefits. *Applied Mathematics and Computation*, 128228. 10.1016/j.amc.2023.128228
- Yang, Y., Zhang, W., & Song, S. (2025). Distributed fusion estimator for cyber-physical system with multi-channel network under dos attacks. *Digital Signal Processing*. 10.1016/j.dsp.2025.105192
- Yi, T., & Zuwang, W. (1997). Effect of time delay and evolutionarily stable strategy. *Journal of Theoretical Biology*, 187(1), 111–116. 10.1006/jtbi.1997.0427
- Zhang, S., & Li, C. (2026). Asymmetric higher-order interactions in a fractional-order delayed malware propagation model. *Communications in Nonlinear Science and Numerical Simulation*. 10.1016/j.cnsns.2026.110605
- Zhao, C., & Li, J. (2020). Equilibrium selection under the Bayes-based strategy updating rules. *Symmetry*, 12(5), 739. 10.3390/sym12050739

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplementarymaterial.rar](#)