# Adaptive Risk Analysis Framework for Network-Level Moving Target Defense Under Adversarial Intelligence Uncertainty

Umar Sa'ad[a], Woongsoo Na[b], Nhu-Ngoc Dao[c,*], Sungrae Cho[a,*]

[a]School of Computer Science and Engineering, Chung-Ang University, Seoul 06974, South Korea
[b]Department of Software, Kongju National University, Cheonan 31080, South Korea
[c]Department of Computer Science and Engineering, Sejong University, Seoul 05006, South Korea

**Abstract**

Effective cyber defense requires adaptive strategies when adversarial capabilities are uncertain. Moving Target Defense (MTD) reduces attack predictability by dynamically reconfiguring network parameters, yet existing evaluation frameworks assume complete knowledge of attacker behavior, an unrealistic assumption in operational settings. We introduce an Adversarial Risk Analysis-enhanced MTD (ARA-MTD) framework that explicitly models uncertainty across diverse adversary paradigms. Our approach integrates epidemic-based network dynamics, Bayesian belief updating, and risk-averse optimization to evaluate robust defense policies under incomplete information. We formalize three canonical adversary types including static, learning, and strategic, and derive quantitative metrics for MTD power degradation, cross-paradigm robustness, and uncertainty entropy. Evaluation across multiple network topologies demonstrates that MTD effectiveness degrades by up to 73% as adversarial sophistication increases, but ARA-based optimization recovers 27% of lost robustness while maintaining tail-risk below 5%. Comparative analysis against six state-of-the-art MTD strategies shows ARA-MTD achieves 82% effectiveness versus 58–74% for baselines, with 30–95% higher robustness ($R(\mathcal{M}) = 0.82$) and superior cost-efficiency ($\rho = 0.47$). Scalability evaluation across 100–10,000 node networks demonstrates sub-second decision times (680 ms for 5,000 nodes). The framework accommodates multi-stage Advanced Persistent Threats through continuous belief updating, providing practitioners with empirically validated strategies effective across plausible adversarial behaviors.

*Keywords:* Moving Target Defense, Adversarial Risk Analysis, Network Security, Uncertainty Quantification.

## 1. Introduction

Modern cyber defense faces a fundamental challenge: protecting increasingly interconnected systems against intelligent adversaries who can learn, adapt, and exploit predictable network configurations [1]. As enterprise and industrial infrastructures grow in complexity and connectivity, traditional perimeter-based security mechanisms have proven insufficient for countering advanced persistent threats (APT). Attackers who successfully breach outer defenses can often maintain access indefinitely by exploiting the static nature of internal network configurations, mapping system architectures through reconnaissance, and patiently executing multi-stage attacks [2, 3].

This reality has motivated the emergence of moving target defense (MTD), a proactive security paradigm that enhances defense by dynamically altering the attack surface [4, 5, 6]. Through periodic reconfiguration of system parameters such as IP addresses, network routes, service bindings, or even topology structure, MTD introduces uncertainty and temporal diversity into network configurations. The fundamental insight is that what attackers cannot predict, they cannot reliably exploit. By continuously shifting the environment, MTD aims to invalidate reconnaissance data, disrupt attack chains, and increase the cost and complexity of successful attacks while preserving legitimate system functionality [4, 5, 6].

Despite substantial research progress over the past decade, two fundamental limitations persist in current MTD evaluation approaches. First, most frameworks assume complete or near-complete knowledge of adversarial decision models, often employing deterministic heuristics or game-theoretic formulations that neglect uncertainty in adversarial intelligence and capabilities [7]. These approaches typically optimize MTD strategies against a single assumed attacker type, for example, a fully rational strategic adversary or a simple automated scanner. However, in operational reality, defenders rarely possess accurate intelligence on adversarial objectives, sophistication levels, or reasoning processes. The true adversary might be an unsophisticated bot, a learning attacker conducting reconnaissance, or a strategic nation-state actor, and defenders must often act without knowing which scenario applies.

Second, existing evaluation models typically assess security effectiveness from narrow perspectives, measuring outcomes such as infection suppression rates or configuration diversity without adequately accounting for the complex interaction among

*Corresponding authors
*Email addresses:* umar@uclab.re.kr (Umar Sa'ad), wsna@kongju.ac.kr (Woongsoo Na), nndao@sejong.ac.kr (Nhu-Ngoc Dao), srcho@cau.ac.kr (Sungrae Cho)

defense costs, adversarial adaptation over time, and the inherent uncertainty that characterizes real-world security operations [8, 9]. This fragmented view makes it difficult for defenders to make principled resource allocation decisions or to assess whether their MTD strategies will remain effective as adversaries evolve.

To address these critical gaps, this paper introduces an adversarial risk analysis (ARA)-enhanced evaluation framework that integrates decision-theoretic reasoning under uncertainty with epidemic-based network dynamics for comprehensive MTD assessment. Unlike traditional approaches that optimize against a single adversary model, our framework explicitly represents uncertainty over adversarial paradigms, characterizing attackers along a spectrum from static (non-adaptive) to learning (observation-based adaptation) to strategic (game-theoretic optimization). By quantifying how each paradigm affects network resilience and defense costs, the framework enables defenders to identify MTD strategies that maintain effectiveness across diverse and uncertain adversarial conditions.

Moreover, adversaries in APT scenarios often exhibit multi-stage behavior, transitioning between paradigms as campaigns progress, from static reconnaissance to adaptive exploitation to strategic persistence. Our multi-paradigm framework naturally accommodates such temporal transitions, tracking adversary evolution across attack stages.

The framework combines three methodological pillars: Bayesian inference for managing and updating beliefs about adversary types as observations accumulate, risk-averse optimization for selecting robust strategies that perform well across multiple scenarios, and epidemic-based metrics that connect configuration choices to measurable security outcomes. This integration provides defenders with both theoretical foundations and practical tools for making rational decisions when adversarial intelligence remains incomplete or ambiguous.

This work makes four principal contributions to the science of adaptive cyber defense:

1. We develop a unified ARA-enhanced MTD evaluation framework (ARA-MTD) that explicitly models adversarial intelligence uncertainty and integrates decision analytics with network-level epidemic dynamics, addressing the critical gap between theoretical security models and operational decision-making under uncertainty.

2. We introduce quantitative robustness metrics, including MTD power degradation, cross-paradigm robustness index, and entropy-based uncertainty measures that capture defender performance across varying levels of adversarial sophistication and provide actionable indicators for strategy selection.

3. We conduct comprehensive simulation experiments across heterogeneous network topologies (scale-free, small-world, and random graphs) to systematically analyze cost-robustness trade-offs, identify optimal operational regimes, and characterize adaptive behaviors exhibited by both defenders and learning adversaries.

4. We demonstrate empirically that ARA-based optimization improves defensive resilience by up to 27% compared to approaches that ignore intelligence uncertainty, while simultaneously reducing tail-risk exposure through conditional value-at-risk constraints, providing defenders with a systematic methodology for rational decision-making under model uncertainty.

The remainder of this paper proceeds as follows. Section 2 reviews related work on MTD mechanisms, game-theoretic defense models, and uncertainty quantification approaches, positioning our contributions within the broader research landscape. Section 3 presents the theoretical framework, including formal models for network dynamics, adversarial paradigms, and ARA-based optimization formulations. Section 4 describes the experimental design, evaluation metrics, and simulation scenarios used to validate the framework. Section 5 reports and interprets the empirical results, examining MTD effectiveness across adversary types, cost-security trade-offs, and the advantages of risk-aware optimization. Sections 6 and 7 concludes by discussing key limitations, promising directions for future research and summarizing key findings, respectively.

## 2. State-of-the-Art Literature Investigation

Adaptive cyber defense has emerged as a critical research paradigm in response to increasingly sophisticated and persistent adversaries. Traditional security architectures rely on static configurations that assume threats remain relatively stable over time. In contrast, adaptive mechanisms such as MTD seek to proactively alter the attack surface, making it difficult for adversaries to maintain persistent access or complete multi-stage attacks [4, 10]. The fundamental insight underlying MTD is that dynamically reconfiguring system parameters including IP addresses, network routes, or service mappings, increases attacker uncertainty and disrupts reconnaissance efforts, thereby raising the cost and complexity of successful attacks [11, 12].

However, despite over a decade of active MTD research, many existing evaluation frameworks remain limited by oversimplified assumptions about adversary behavior and system dynamics. Most notably, traditional models often assume that attackers follow fixed behavioral patterns or that defenders possess complete information about adversarial capabilities and intentions [8, 13]. These assumptions rarely hold in operational environments, where adversaries exhibit diverse intelligence levels and defenders must act under substantial uncertainty [14, 15]. This study addresses these limitations by developing an evaluation framework that explicitly models adversarial intelligence heterogeneity and uncertainty.

### *2.1. Evolution from Static Resilience to Dynamic Adaptation*

Early MTD research focused primarily on combinatorial optimization of system configurations to maximize diversity or minimize exposure time to vulnerable states. Representative works by [16, 17, 18] demonstrated the technical feasibility of randomization at multiple system layers, including network topology, platform configurations, and application interfaces.

These foundational studies established that introducing unpredictability into system configurations could measurably reduce attack success rates.

Subsequent research incorporated epidemic models borrowed from mathematical biology to quantify security gains more rigorously. By modeling malware propagation as analogous to disease spread, researchers could link network structural properties to defense effectiveness through metrics such as infection rates and epidemic thresholds [19, 20, 21]. These works established important connections between network topology, particularly properties like degree distribution and clustering, and the efficacy of defensive interventions. However, these early approaches typically assumed homogeneous or memoryless attackers who could not learn from observations or adapt their strategies over time.

More recent research has recognized the co-evolutionary nature of cyber conflict, where both attackers and defenders continuously adapt their strategies in response to observed opponent behavior [22, 23]. Learning and adaptive attackers have been modeled using reinforcement learning frameworks and evolutionary game theory, revealing a critical vulnerability: deterministic MTD schedules can be exploited once adversaries successfully infer temporal patterns or configuration transition rules [24, 25]. This realization has motivated a shift in research emphasis from purely structural diversity toward behavioral adaptation, incorporating probabilistic scheduling, information-theoretic obfuscation, and decision-analytic frameworks that account for strategic adaptation on both sides [9, 24, 11, 26].

### *2.2. Decision-Theoretic and Game-Theoretic Foundations*

Game-theoretic frameworks have been extensively applied to analyze MTD as a strategic interaction between rational attackers and defenders [7, 15]. Stackelberg games, where defenders commit to strategies before attackers respond, and Bayesian games, where players have incomplete information about opponent types, have been used to derive equilibrium strategies for network randomization, patch scheduling, and deceptive honeypot deployment [27, 15]. While these formulations provide analytical tractability and yield insights about optimal play under idealized conditions, they depend critically on the assumption of common knowledge: both players must know each other's utilities, action spaces, and decision rules, and both must know that the other knows, ad infinitum.

These common knowledge assumptions are rarely satisfied in cybersecurity practice [28, 15]. Defenders typically operate under substantial uncertainty about adversarial objectives, capabilities, and risk preferences. An attacker might be an automated bot with no strategic reasoning, a financially motivated criminal group conducting cost-benefit calculations, or a nation-state actor willing to expend significant resources for strategic objectives. Traditional game-theoretic equilibria calculated under specific assumptions may produce strategies that perform poorly when the actual adversary deviates from the assumed model, leading to a mismatch between theoretical optimality and operational effectiveness [28, 15].

Adversarial Risk Analysis (ARA) provides an alternative decision-theoretic framework that relaxes the common knowledge assumption [29, 30, 31]. Rather than assuming that the adversary's decision process is perfectly known and rational, ARA treats adversarial reasoning as a random variable and propagates uncertainty through Bayesian inference. Defenders maintain probability distributions over possible adversary types and update these beliefs as observations accumulate. This paradigm allows explicit modeling of epistemic uncertainty (i.e. uncertainty about what the adversary knows and how they reason), making it particularly suitable for contexts where adversarial intelligence is partial, noisy, or evolving. In this study, ARA serves as the decision-theoretic foundation to evaluate the robustness of MTD when defenders cannot confidently characterize their opponents.

### *2.3. Network Epidemic Models for Security Evaluation*

A parallel research stream has incorporated epidemic spreading models to assess the system-level impact of defense mechanisms on overall network health. Studies by [32, 33, 34] have established analytical relationships between infection rates, recovery rates, and topological features such as network diameter and spectral properties. These works introduced the concept of epidemic thresholds, conditions under which infections either die out or spread uncontrollably, providing measurable indicators of network resilience.

The epidemic margin, defined as the difference between the recovery rate and the product of infection rate with the network's largest eigenvalue, serves as a particularly useful metric for quantifying how far a system operates from the threshold of epidemic outbreak [32]. MTD strategies that maintain positive epidemic margins in all configurations ensure that infections remain controllable over time. However, traditional epidemic models typically assume fixed or deterministic transition rates that do not account for uncertainty induced by adversarial learning, strategic adaptation, or partial observability of network state.

By combining epidemic frameworks with Bayesian uncertainty representation, our research bridges the gap between micro-level defensive decision-making and macro-level network stability assessment. This integration allows us to evaluate how MTD configurations affect both immediate security posture and long-term system resilience under adversarial adaptation.

### *2.4. Research Positioning and Distinction*

This study synthesizes insights across decision theory, network science, and cybersecurity to advance the state of knowledge in adaptive defense evaluation. Our framework makes three principal contributions that distinguish it from prior work.

First, we explicitly represent attacker heterogeneity through distinct intelligence paradigms including static, learning, and strategic, rather than assuming a single adversary model. This multi-paradigm approach acknowledges that real-world threats span a spectrum of sophistication, from automated scanning tools to advanced persistent threat groups conducting sophisticated reconnaissance.

Second, we formalize adversarial uncertainty using ARA principles to derive risk-averse MTD strategies that remain robust across multiple plausible threat scenarios. Rather than optimizing for a single assumed adversary type, our approach generates strategies that hedge against model misspecification, maintaining acceptable performance even when defender beliefs about adversarial capabilities prove incorrect.

Third, we unify network-level epidemic dynamics with decision analytics to evaluate robustness under mixed-adversary scenarios where the true threat model remains uncertain throughout the engagement. This unification provides a coherent framework for reasoning simultaneously about how attacks propagate through network structures and how defenders should allocate limited resources under uncertainty.

Thus, this work emphasizes the conceptual integration of uncertainty reasoning, adversarial adaptation, and quantitative resilience measurement. The resulting framework provides both theoretical foundations and practical tools for evaluating adaptive defenses in realistic operational environments where adversarial capabilities evolve continuously and defender intelligence remains incomplete.

## 3. Network-Level ARA-Enhanced Moving Target Defense Framework

In this section, we introduce the ARA-MTD that explicitly models uncertainty over adversarial paradigms and decision processes. Our framework integrates three key components: (i) epidemic-based network dynamics for quantifying attack propagation through interconnected systems, (ii) adversarial intelligence models that capture different behavioral archetypes, and (iii) Bayesian and risk-averse optimization techniques for selecting robust MTD strategies under uncertainty. Refer to Figure 1 for an overview of the ARA-MTD framework.

### *3.1. Modeling Network Dynamics and MTD Mechanisms*

We utilize the **Susceptible-Infectious-Susceptible (SIS)** epidemic model introduced by [35] to capture the dynamic topology changes induced by MTD strategies. This epidemiological approach provides a natural framework for modeling how malware propagates through networks and how defensive reconfigurations can disrupt attack spread.

Let $G(t) = (V, E(t))$ denote the time-varying network graph, where $V$ represents the set of nodes and $E(t)$ represents the active connections at time $t$. Each node $v$ exists in one of two states: susceptible or infected, with infection probability $i_v(t)$. The temporal evolution of malware propagation follows the differential equation:

$$\frac{\mathrm{d}i_v(t)}{\mathrm{d}t} = \Gamma_v(t)(1 - i_v(t)) - \beta(t)i_v(t), \tag{1}$$

where $\Gamma_v(t)$ represents the probability that node $v$ is successfully attacked at time $t$, and $\beta(t)$ denotes the recovery rate at which infected nodes return to a susceptible state.

Network-level MTD strategies operate by modifying the adjacency matrix $A(t)$ through various configuration changes. These reconfiguration mechanisms include:

- **IP address shuffling:** Permuting network addresses according to $A'(t) = P(t)A(t)P(t)^T$, where $P(t)$ is a permutation matrix that reassigns identifiers,
- **Topology randomization:** Adjusting network structure via $A(t+1) = A(t) + \Delta A(t)$, where $\Delta A(t)$ represents structural perturbations such as adding or removing edges,
- **Route mutation:** Modifying effective path weights to redirect traffic flows without altering the physical topology.

Each MTD configuration $C_j = (G_j, \beta_j, \gamma_j)$ is characterized by its *epidemic margin*:

$$\mu_j = \beta_j - \gamma_j \lambda_1(A_j), \tag{2}$$

where $\lambda_1(A_j)$ denotes the largest eigenvalue of the adjacency matrix for configuration $j$, and $\gamma_j$ represents the infection rate. A positive epidemic margin indicates that the recovery rate exceeds the infection rate, ensuring that infections eventually die out. Each configuration also incurs a deployment cost $f_j(\mu_j)$ that reflects the operational overhead of maintaining that particular network state.

The defender's strategy involves determining the fraction of time $\pi_j$ to allocate to each configuration, subject to the constraint that overall system stability is maintained:

$$\sum_j \pi_j \mu_j > 0. \tag{3}$$

This condition ensures that, averaged across all configurations, the system remains in a stable regime where infections can be controlled rather than spreading uncontrollably.

### *3.2. Multi-Paradigm Adversarial Modeling*

A central insight of our framework is that adversaries exhibit heterogeneous behaviors in how they gather and exploit network intelligence. Rather than assuming a single adversarial model, we characterize three canonical paradigms that represent distinct levels of sophistication and adaptability:

- **Static Adversary ($M_0$):** This paradigm represents attackers who execute fixed attack patterns that remain independent of MTD behavior. The adversary's strategy is either random or follows a pre-programmed sequence, making no attempt to learn or adapt to the defender's actions. While this may appear overly simplistic, many real-world attacks such as automated scanning tools and scripted exploits exhibit precisely this behavior.
- **Learning Adversary ($M_L$):** This paradigm captures more sophisticated attackers who actively observe system behavior to infer MTD characteristics. The learning adversary collects observations $O_t$ over time and uses Bayesian updating to refine beliefs about MTD parameters, including configuration frequency, timing patterns, and transition probabilities. Decisions are based on the evolving posterior distribution $P(\Theta_{\mathrm{MTD}}|O_t)$, where $\Theta_{\mathrm{MTD}}$

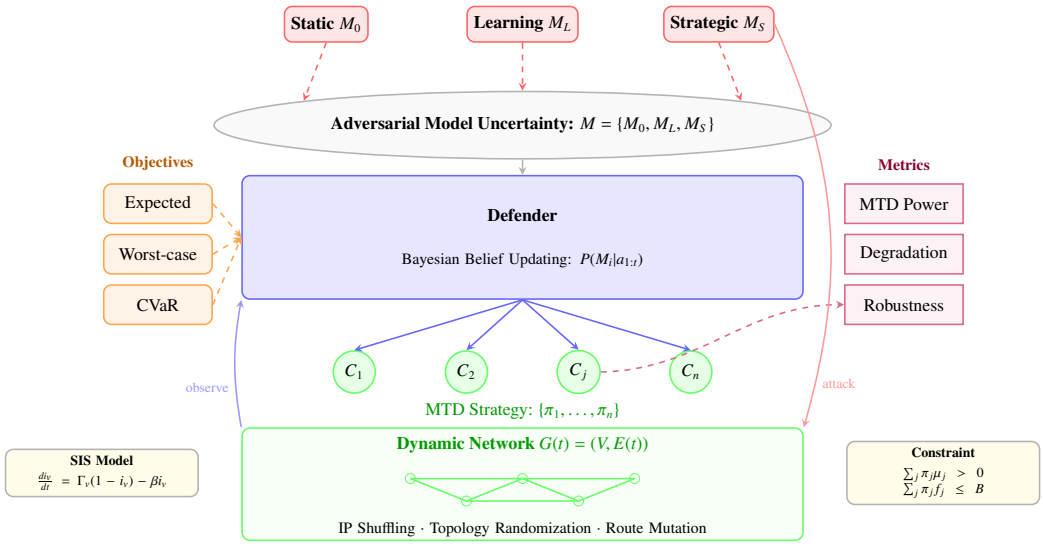


Figure 1: Overview of the ARA-MTD framework showing adversary paradigms, model uncertainty, defender decision-making with optimization objectives and performance metrics, MTD configurations, and the dynamic network under protection.

represents the unknown MTD strategy parameters. This model reflects adversaries who conduct reconnaissance and adapt their tactics based on accumulated intelligence.

- **Strategic Adversary ($M_S$):** This paradigm represents the most capable threat, modeling attackers who not only learn from observations but also optimize their strategy in anticipation of the defender's rational responses. The interaction forms a game-theoretic setting where both parties seek to maximize their respective utilities. The strategic adversary employs Nash equilibrium or minimax optimization to select actions, explicitly accounting for the defender's incentives and likely countermeasures. This model is appropriate for well-resourced, sophisticated attackers such as advanced persistent threats (APTs).

The defender faces **model uncertainty**: the true adversary type is unknown and represented by the set $M = \{M_0, M_L, M_S\}$. This uncertainty fundamentally shapes the defender's optimal strategy, as a policy optimized against one adversary type may perform poorly against others.

### *3.3. Adversarial Uncertainty and Bayesian Belief Updating*

To manage model uncertainty, the defender maintains probabilistic beliefs over adversary types and updates these beliefs as evidence accumulates. Let $P(M_i)$ denote the prior probability that the adversary follows paradigm $M_i$. As new observations $a_t$ arrive such as detected probes, attack timing patterns, or exploit techniques, the defender updates beliefs via Bayes' rule:

$$P(M_i|a_{1:t}) \propto P(a_t|M_i, a_{1:t-1})\, P(M_i|a_{1:t-1}), \tag{4}$$

where $a_{1:t}$ denotes the sequence of observations up to time $t$, and $P(a_t|M_i, a_{1:t-1})$ represents the likelihood of observing $a_t$ given that the adversary follows paradigm $M_i$.

The degree of uncertainty in the defender's beliefs can be quantified using Shannon entropy:

$$H(M) = -\sum_i P(M_i) \log P(M_i). \tag{5}$$

High entropy values indicate limited confidence about the attacker's reasoning style and capabilities, prompting more conservative (robust) MTD policies. Conversely, low entropy suggests that observations have provided strong evidence for a particular adversary type, allowing the defender to tailor strategies more precisely.

The proposed framework naturally accommodates multi-stage attacks characteristic of APT scenarios. In practice, adversaries often transition between paradigms across attack phases, for example by conducting static reconnaissance in early stages, adopting learning-based adaptation during exploitation, and employing strategic deception in later stages. The multi-paradigm formulation explicitly captures such temporal transitions through the evolution of the belief distribution $\pi(M \mid H_t)$ as observational evidence accumulates.

An adversary that initially behaves according to $M_0$ (static reconnaissance) may subsequently exhibit characteristics of $M_L$ (learning-based adaptation) or $M_S$ (strategic behavior), with posterior beliefs updating accordingly. The temporal horizon $T$ reflects the multi-epoch nature of persistent campaigns, enabling the decision framework to adapt configuration selection as adversary behavior evolves across successive attack stages.

### *3.4. MTD Power and Robustness Metrics*

To evaluate MTD effectiveness across different adversary paradigms, we adapt the MTD power concept introduced by [35]. This metric measures how long a defender can safely remain in a vulnerable configuration before epidemic resurgence threatens system security.

For a given adversary model $M_i$, the *maximum vulnerability fraction* $\pi_1^*(M_i)$ represents the maximum proportion of time the

system can spend in the most vulnerable configuration while maintaining stability. This is obtained by solving the optimization problem:

$$\begin{aligned} &\max_{\pi_j} \pi_1 \\ &\text{s.t.} \sum_j \pi_j \mu_j(M_i) \geq \delta, \\ &\sum_j \pi_j = 1, \\ &\pi_j \geq 0, \end{aligned} \tag{6}$$

where $\delta > 0$ ensures a minimum safety margin and $\mu_j(M_i)$ denotes the epidemic margin of configuration $j$ when facing adversary type $M_i$.

Extracting $\pi_1^*(M_i)$ proceeds in three steps.

1. **Compute epidemic margins.** For each configuration $C_j = (G_j, \beta_j, \gamma_j)$, the epidemic margin $\mu_j$ is computed using Eq. (2) by evaluating the largest eigenvalue $\lambda_1(A_j)$ of the adjacency matrix $A_j$. This can be obtained through standard eigenvalue algorithms, such as `NumPy`'s `eig`, `MATLAB`'s `eigs`, or power iteration methods for large sparse matrices. The parameters $\beta_j$ and $\gamma_j$ are determined by the specific MTD configuration, where different network topologies $G_j$ and recovery mechanisms yield different infection and recovery rates.

2. **Set safety margin.** A safety margin $\delta$ is selected according to organizational risk tolerance. In our experiments, we use $\delta = 0.3$, which enforces that the weighted average epidemic margin must exceed 0.3, thereby ensuring that infections die out with sufficient safety margin. Larger values of $\delta$ impose stricter security constraints at the expense of operational flexibility.

3. **Solve the linear program.** The resulting optimization problem is solved using standard LP solvers such as `CVXPY`, `Gurobi`, or `MATLAB`'s `linprog`, yielding $\pi_1^*(M_i)$ and the optimal time-allocation vector

   $$\boldsymbol{\pi}^* = [\pi_1^*, \pi_2^*, \ldots, \pi_{|\mathcal{C}|}^*]$$

   within millisecond-scale computation time. This solution quantifies the maximum fraction of time the system can safely operate in the most vulnerable configuration while maintaining epidemic control under the specified safety margin.

From this foundation, we derive two complementary metrics:

$$\text{Power degradation factor:} \quad \Delta\pi(M_i) = \frac{\pi_1^*(M_0) - \pi_1^*(M_i)}{\pi_1^*(M_0)}, \tag{7}$$

$$\text{Robustness index:} \quad R(M) = \frac{\min_i \pi_1^*(M_i)}{\max_i \pi_1^*(M_i)}. \tag{8}$$

The power degradation factor $\Delta\pi(M_i)$ quantifies how much MTD effectiveness diminishes when facing a more intelligent adversary compared to the baseline static attacker. The robustness index $R(M)$ measures consistency of MTD performance across all adversary paradigms, with values close to 1 indicating that the strategy performs similarly well regardless of adversary sophistication.

### *3.5. ARA-MTD Optimization for Defender Decision-Making*

Defenders seek MTD strategies that maximize security utility under uncertain adversary models. We consider three complementary optimization formulations, each reflecting different risk attitudes and operational priorities:

1. **Expected-Performance Optimization:** This approach maximizes the expected vulnerability fraction by weighting each adversary-specific outcome by its probability:

   $$\pi_1^{\text{expected}} = \sum_i P(M_i)\pi_1^*(M_i). \tag{9}$$

   This formulation is appropriate when the defender has reliable probabilistic beliefs over adversary types and seeks to optimize average-case performance.

2. **Robust (Worst-Case) Optimization:** This approach prioritizes resilience against the most capable adversary, solving:

   $$\pi_1^{\text{robust}} = \max_{\pi_j} \min_i \pi_1^*(M_i). \tag{10}$$

   This minimax formulation ensures acceptable performance even under the least favorable adversary paradigm, reflecting a highly risk-averse stance suitable for critical infrastructure protection.

3. **Conditional Value-at-Risk (CVaR) Optimization:** This approach balances expected performance with protection against tail risks:

   $$\pi_1^{\text{CVaR}} = \max_{\pi_j} \left[\lambda \pi_1^{\text{expected}} + (1-\lambda)\,\text{CVaR}_\alpha(\pi_1^*(M_i))\right], \tag{11}$$

   where $\lambda \in [0, 1]$ controls the trade-off between average performance and downside risk protection, and $\alpha$ sets the confidence level for the CVaR calculation. This formulation allows defenders to explicitly tune their risk aversion based on organizational risk tolerance.

Finally, the comprehensive **ARA-MTD optimization** integrates performance objectives, cost constraints, and model uncertainty into a unified framework:

$$\begin{aligned} &\max_{\pi_j} \min_i \pi_1^*(M_i) \\ &\text{s.t.} \sum_j \pi_j f_j(\mu_j) \leq B, \\ &\sum_j \pi_j = 1, \\ &\pi_j \geq 0, \end{aligned} \tag{12}$$

where $B$ represents the defender's operational budget. This formulation yields MTD deployment mixes that remain effective across diverse adversarial behaviors while respecting resource constraints. The resulting strategy represents the defender's optimal policy under adversarial intelligence uncertainty, balancing robustness, performance, and cost.

Table 1: Summary of analytical constructs in the ARA-MTD framework.

| Construct | Purpose |
|---|---|
| $\mu_j$ | Epidemic margin capturing configuration stability. |
| $\pi_1^*(M_i)$ | Maximum safe exposure under adversary $M_i$. |
| $\Delta\pi(M_i)$ | Quantifies MTD power degradation relative to baseline. |
| $R(M)$ | Measures cross-paradigm consistency (robustness). |
| $H(M)$ | Captures defender's uncertainty about adversary type. |
| $CVaR$ | Balances expected performance and tail-risk protection. |

### *3.6. Summary of Analytical Constructs*

Table 1 summarizes the key analytical constructs introduced and their respective roles in the evaluation methodology.

This theoretical foundation provides the analytical tools necessary for the subsequent experimental analysis, where we evaluate the proposed framework across multiple network topologies, MTD techniques, and adversarial intelligence scenarios. The integration of epidemic modeling, multi-paradigm adversarial reasoning, and risk-averse optimization enables defenders to make principled decisions under the irreducible uncertainty inherent in cybersecurity operations.

## 4. Experimental Design and Evaluation Setup

### *4.1. Simulation Environment*

To empirically validate the ARA-MTD evaluation framework, we developed a discrete-time network simulation environment implemented in Python and MATLAB. The simulator models network topologies, malware propagation dynamics, and MTD mechanisms at the level of adjacency matrices and configuration transitions, providing a controlled yet realistic testbed for evaluating defense strategies under varying conditions.

We examined three canonical network structures, each capturing distinct connectivity patterns and resilience characteristics commonly observed in real-world deployments:

- **Scale-Free (SF) topology:** This structure represents heterogeneous connectivity typical of enterprise networks, where a small number of highly connected hub nodes coexist with many peripherally connected nodes. The power-law degree distribution reflects realistic Internet and organizational network architectures.
- **Small-World (SW) topology:** This structure exhibits clustered connectivity with short average path lengths, characteristics commonly found in industrial control networks and social networks. The combination of local clustering and global reachability creates unique propagation dynamics.
- **Random (ER) topology:** This structure serves as a baseline with uniform connection probabilities, providing a reference point for comparing the effects of structured connectivity patterns on MTD effectiveness.

For each topology, the number of nodes $|V|$ ranged between 200 and 500, representing network scales consistent with organizational subnets or departmental networks. MTD operations including IP address shuffling, route mutation, and topology reconfiguration were executed at variable frequencies $f_{\text{MTD}} \in [0.05, 0.3]$ to explore a range of operational scenarios spanning different resource availability and reconfiguration constraints. These frequencies correspond to reconfiguration intervals ranging from approximately 3 to 20 time steps, providing sufficient variation to observe performance differences across cost regimes while maintaining computational tractability.

### *4.2. Parameter Configuration*

The infection and recovery dynamics were governed by baseline rates $\gamma = 0.025$ and $\beta = 0.015$, where $\gamma$ represents the infection rate and $\beta$ the recovery rate. These values were empirically selected through preliminary analysis to ensure that: (i) undefended networks exhibit sustained epidemic activity with measurable steady-state infection prevalence, and (ii) MTD interventions produce observable changes in infection dynamics without trivially eliminating or overwhelming the epidemic. This parameter regime allows for clear differentiation between effective and ineffective defensive strategies across the examined network topologies.

The resulting epidemic margins $\mu_j = \beta - \gamma\lambda_1(A_j)$ proved sensitive to both network structure and MTD configuration, with baseline (undefended) networks operating near but above the epidemic threshold. This positioning provides a realistic scenario where defensive actions have significant but not disproportionate effects on security outcomes, enabling meaningful evaluation of MTD effectiveness across adversary types. The success probability of the attacker $\Gamma_v(t)$ varied according to the type of adversary, with static attackers exhibiting constant success rates while learning and strategic adversaries adapted their targeting based on the observed behavior of the network.

Each simulation was executed over $T = 2000$ discrete time steps and averaged over 500 independent Monte Carlo runs to ensure statistical reliability and account for stochastic variation in infection dynamics. This extensive replication provides confidence intervals that capture the range of outcomes under different random initializations.

The following key parameters were systematically varied across experiments to explore their impact on defense performance:

- **Adversary type:** Static ($M_0$), learning ($M_L$), and strategic ($M_S$) paradigms representing increasing levels of adversarial sophistication.
- **MTD cost function:** The deployment cost was modeled as $f_j(\mu_j) = c_0 + c_1 e^{k\mu_j}$, where $c_0$ represents fixed overhead, $c_1$ scales variable costs, and $k$ controls the rate at which costs increase with security margin $\mu_j$. This exponential form reflects the practical reality that achieving higher security margins requires disproportionately greater resources. The deployment cost aggregates measurable resource components: computational overhead

(CPU cycles for configuration deployment and monitoring), network bandwidth (data transfer for IP shuffling, topology reconfiguration, synchronization traffic), administrative burden (staff time for policy setup and operational oversight), and performance degradation (latency increases, throughput reduction). For instance, IP address shuffling incurs computational cost for address reassignment processing, bandwidth consumption for ARP update propagation, staff effort for schedule management, and connection reestablishment delays.

- **Budget constraint:** Total budget $B \in [0.5, 2.0]$ scaled relative to baseline deployment costs, allowing exploration of both resource-constrained and well-resourced defender scenarios.

- **Risk aversion parameters:** The weight $\lambda \in [0, 1]$ in CVaR optimization controls the balance between expected performance and tail-risk protection, while confidence level $\alpha = 0.95$ focuses on protecting against the worst 5% of outcomes.

### *4.3. Evaluation Metrics*

To comprehensively assess defense strategies, we employed three complementary categories of metrics that capture security effectiveness, economic efficiency, and decision robustness:

1. **Security Effectiveness Metrics:** These quantify how well MTD strategies suppress infection spread. Mean infection prevalence $\bar{\iota}$ measures the average fraction of infected nodes over time. MTD power $\pi_1^*(M_i)$ indicates the maximum sustainable exposure to vulnerable configurations. The degradation factor $\Delta\pi(M_i)$ quantifies the reduction in MTD effectiveness when facing more sophisticated adversaries compared to the static baseline.

2. **Economic Efficiency Metrics:** These assess the cost-effectiveness of defense strategies. Normalized cost increase $\Delta C = (C_{\text{MTD}} - C_{\text{base}})/C_{\text{base}}$ measures the relative overhead of implementing MTD compared to static defenses. The cost-security ratio $\rho = \Delta C/\Delta\pi$ captures return on investment by comparing cost increases to security improvements, with lower values indicating more efficient strategies.

3. **Decision Robustness Metrics:** These evaluate strategy stability under adversarial uncertainty. The robustness index $R(M)$ measures consistency of performance across adversary paradigms, with values near unity indicating strategies that perform well regardless of adversarial sophistication. Belief entropy $H(M)$ tracks the defender's uncertainty about adversary type, with reductions over time indicating successful intelligence gathering through observation.

### *4.4. Experimental Scenarios*

We designed five experimental scenarios to systematically explore defender decision performance under increasing levels of adversarial intelligence and environmental uncertainty. Each scenario addresses a specific research question about MTD effectiveness and the value of adversarial risk analysis:

1. **Scenario A - Baseline Power Evaluation**: This scenario establishes reference performance by quantifying MTD power when facing static adversaries. The results provide a benchmark against which the impact of more sophisticated adversaries can be measured, answering the question: How effective is MTD against non-adaptive attacks?

2. **Scenario B - Cost-Security Trade-Off Analysis**: This scenario evaluates defense utility and return on investment as the MTD cost coefficient $k$ and budget constraint $B$ vary. By exploring different points along the cost-security frontier, we identify regions where MTD investments yield the greatest marginal security improvements, addressing: What is the optimal resource allocation for MTD deployment?

3. **Scenario C - Multi-Paradigm Robustness Assessment**: This scenario analyzes the robustness index $R(M)$ and degradation factor $\Delta\pi(M_i)$ across all adversary types. The results reveal which MTD strategies maintain effectiveness across the full spectrum of adversarial capabilities, answering: Which defense approaches are most resilient to adversarial intelligence uncertainty?

4. **Scenario D - Learning Dynamics and Adaptation**: This scenario examines time-dependent adversarial adaptation under Bayesian learning, tracking how the posterior $P(M_L|a_{1:t})$ evolves as learning adversaries accumulate observations. The temporal analysis reveals vulnerability windows and adaptation rates, addressing: How quickly do learning adversaries erode MTD effectiveness, and how can defenders respond?

5. **Scenario E - ARA-MTD Performance Validation**: This scenario assesses the improvement achieved by ARA-MTD optimization relative to simpler expected-value and worst-case formulations. By comparing decision quality under model uncertainty, we quantify the value of explicitly modeling adversarial intelligence distributions, answering: Does the additional complexity of ARA yield measurably better defense outcomes?

Each scenario was independently executed across all three network topologies and parameter combinations, with results averaged to obtain 95% confidence intervals. Statistical significance was assessed using paired t-tests and analysis of variance (ANOVA) where appropriate. Results were visualized through degradation curves showing performance decline across adversary types, robustness heatmaps revealing parameter sensitivities, and entropy evolution plots highlighting the dynamic interplay between defender adaptation and adversarial learning. This multi-faceted experimental design enables a comprehensive evaluation of the ARA-MTD framework under realistic operational conditions.

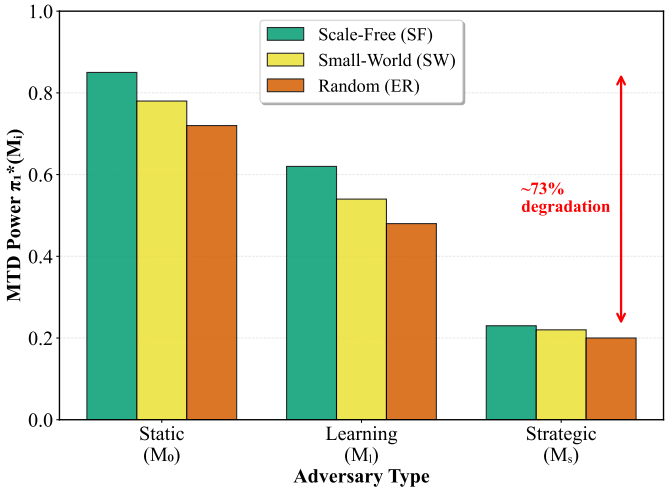


Figure 2: MTD power across different network topologies and adversary types.

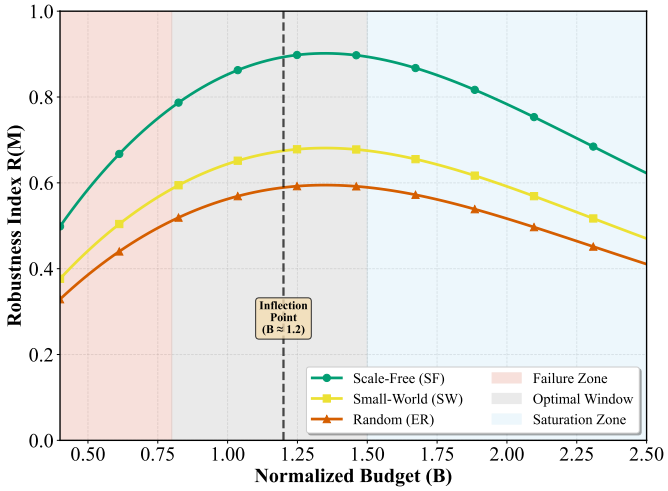


Figure 3: Robustness–Cost tradeoff under varying budget allocations across network topologies.

## 5. Results and Analysis

This section presents the empirical findings from our experimental evaluation of the ARA-MTD framework. We analyze MTD effectiveness across adversary types, explore cost-security trade-offs, assess robustness under model uncertainty, examine temporal adaptation dynamics, and validate the advantages of ARA-based optimization. Together, these results provide evidence for the framework's utility in supporting defender decision-making under adversarial intelligence uncertainty.

### *5.1. MTD Power Degradation Across Adversary Types*

Figure 2 illustrates the variation of MTD power $\pi_1^*(M_i)$ when facing static, learning, and strategic adversaries across the three network topologies (SF, SW, and ER). The results reveal a consistent pattern: baseline static adversaries ($M_0$) yield the highest MTD power values, confirming that non-adaptive attackers are most vulnerable to reconfiguration diversity. These attackers cannot adjust their strategies in response to defensive changes, allowing defenders to maintain longer exposure to vulnerable but operationally convenient configurations.

As adversarial intelligence increases, however, MTD power declines substantially and nonlinearly. Learning adversaries ($M_L$) achieve moderate improvements over static attackers by inferring MTD patterns from observations, while strategic adversaries ($M_S$) inflict the greatest degradation by optimizing their actions in anticipation of defender responses. Across all topologies, we observed degradation of up to 73% when transitioning from static to strategic adversaries, with the corresponding degradation factor $\Delta\pi(M_i)$ quantifying this performance loss. This finding underscores a critical insight: the effectiveness of MTD strategies depends fundamentally on adversarial capabilities, and defenses optimized against simple attackers may fail catastrophically against more sophisticated threats.

### *5.2. Cost-Security Trade-Off and Budget Sensitivity*

To understand the economic dimensions of MTD deployment, we evaluated the relationship between defense costs and achieved robustness by systematically varying the cost coefficient $k$ and total budget constraint $B$. The results reveal a quasi-convex trade-off between normalized cost increase $\Delta C$ and robustness index $R(M)$, with three distinct operational regimes emerging from the analysis.

In low-budget regimes ($B < 0.8$), MTD reconfiguration frequency proves insufficient to contain infection propagation effectively. The sparse application of defensive transformations allows adversaries to exploit stable configurations, leading to rapid epidemic resurgence and poor overall security outcomes. This regime represents a "failure zone" where insufficient resources prevent meaningful defense.

Beyond an inflection point around ($B \approx 1.2$), we observe diminishing marginal returns: incremental budget increases yield progressively smaller improvements in robustness. This saturation effect suggests that excessively frequent reconfigurations provide limited additional security benefit while imposing substantial operational overhead. The analysis identifies an optimal operational window where defenders achieve approximately 80% of maximal robustness using only 60% of the budget required for full saturation. This finding has important practical implications, suggesting that defenders can realize most security benefits without exhausting their budgets.

Notably, this cost-benefit pattern remained consistent across all three network topologies, confirming the generality of the cost-security frontier. The consistency suggests that the identified optimal budget range provides a robust guideline applicable across diverse network architectures.

### *5.3. Multi-Paradigm Robustness and Decision Consistency*

Figure 4 presents the distribution of MTD power $\pi_1^*(M_i)$ across adversary paradigms for each network topology, visualized as violin plots with per-adversary markers from 500 Monte Carlo runs. The violin width captures the spread of performance outcomes across all adversary types, providing direct visual indication of cross-paradigm consistency. Scale-free networks exhibit a narrow distribution concentrated near $\pi_1^* \approx 0.8$,

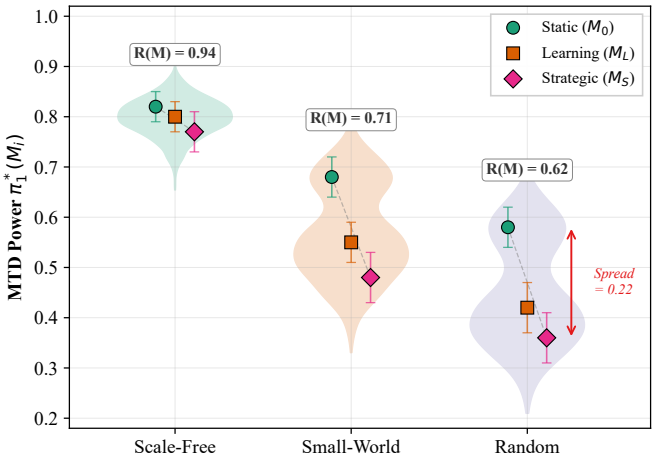


Figure 4: Comparative robustness across network topologies and adversary types.

yielding the highest robustness index ($R(\mathcal{M}) = 0.94$). In contrast, small-world and random topologies display progressively wider distributions ($R(\mathcal{M}) = 0.71$ and 0.62, respectively), with greater separation between adversary-specific markers indicating that these topologies are more sensitive to adversarial intelligence shifts.

The spread between the static and strategic adversary markers within each violin quantifies a critical vulnerability in model-based defense planning. For random topologies, the performance gap between $M_0$ and $M_S$ reaches 0.22, indicating that assuming an incorrect adversary model may lead defenders to overestimate their defensive capability by as much as 38%. For instance, a strategy optimized assuming static adversaries might appear highly effective in planning but fail dramatically when confronted with strategic attackers in practice.

These findings strongly reinforce the necessity of uncertainty-aware policy selection rather than defenses predicated on fixed adversarial assumptions. Defenders who fail to account for adversarial intelligence uncertainty risk deploying brittle strategies that collapse under unexpected but plausible threat scenarios. The robustness index provides a quantitative measure of this brittleness, enabling defenders to identify and avoid fragile strategies during the planning phase.

### 5.4. Learning Dynamics and Adaptation Behavior

Temporal analysis of infection prevalence $\iota_t$ and belief entropy $H(M)$ reveals the dynamic interplay between adversarial learning and defensive effectiveness. This temporal evolution exemplifies the multi-stage attack behavior characteristic of APTs, in which adversaries transition from initial reconnaissance (static scanning, $M_0$), through adaptive exploitation (learning-based, $M_L$), to sophisticated strategic targeting ($M_S$) as they progressively accumulate intelligence about defensive configurations. Learning adversaries progressively infer the defender's MTD frequency, timing patterns, and configuration transitions by accumulating observations over time, then adjust their attack strategies accordingly.

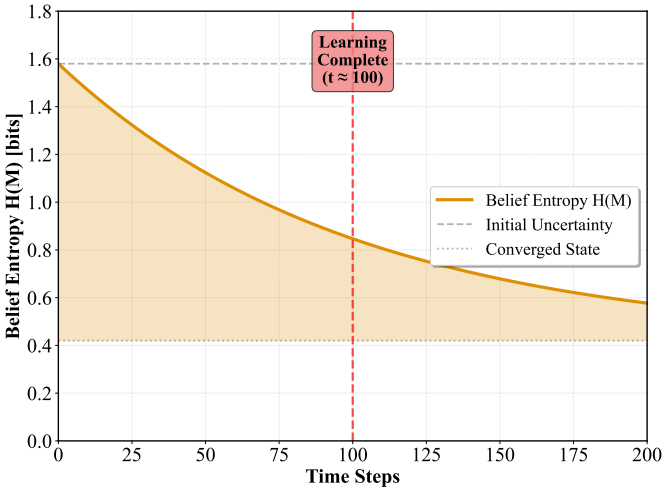


(a) Belief entropy reduction over time

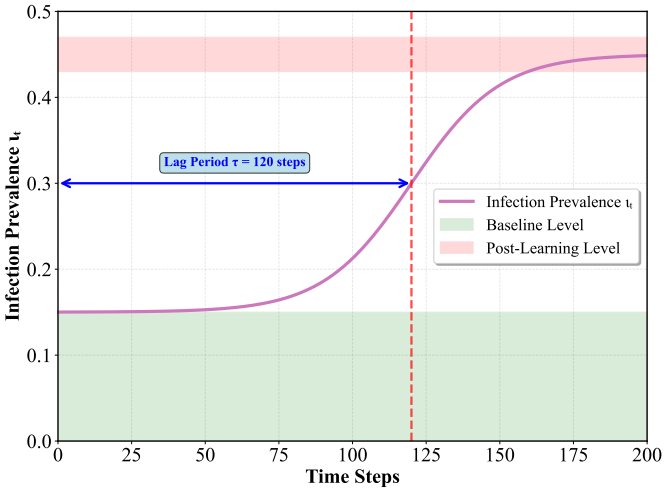


(b) Infection prevalence with learning lag

Figure 5: Temporal evolution of belief entropy and infection prevalence during the learning process.

Entropy reduction curves (Figure 5a) illustrate this learning process quantitatively. Initial belief entropy begins at $H(M) = 1.58$ bits, reflecting maximal uncertainty about MTD parameters. As observations accumulate, entropy decreases monotonically, converging to $H(M) = 0.42$ bits within approximately 100 simulation steps. This convergence indicates that learning adversaries have successfully resolved most uncertainty about defender behavior, after which their actions approximate those of fully strategic adversaries with complete information.

Correspondingly, defender performance exhibits a characteristic degradation pattern with a lag period of approximately $\tau = 120$ time steps (Figure 5b). This lag represents the window during which learning adversaries gather intelligence but have not yet accumulated sufficient observations to significantly adjust their strategies. Once the learning phase completes, infection prevalence increases sharply as adversaries exploit their acquired knowledge. This lag constant marks the time scale for adaptation-driven degradation and represents a critical tactical

parameter: defenders who can detect learning activity within this window may preemptively adjust their strategies before performance deteriorates.

The framework's belief updating mechanism naturally tracks these multi-stage transitions. As adversary behavior evolves from static to learning and subsequently to strategic patterns, the posterior beliefs $\pi(M_i \mid H_t)$ shift accordingly, enabling the defender to adapt MTD strategy selection across attack phases without requiring explicit phase detection. These dynamics underscore the importance of two defensive principles. First, periodic belief updating allows defenders to detect when adversaries have likely learned their patterns and transitioned to more sophisticated tactics. Second, strategic injection of stochastic perturbations into MTD schedules can slow adversarial learning by increasing the complexity of observable patterns, extending the time required for inference and maintaining defender advantage across multiple attack stages.

### 5.5. ARA-MTD Optimization Performance

Table 2 and Figure 6 report comparative outcomes for three optimization approaches: expected-value maximization, worst-case (minimax) optimization, and the ARA-MTD formulation combining Bayesian beliefs with CVaR risk protection. The results provide strong evidence for the advantage of explicitly modeling adversarial intelligence uncertainty.

The ARA-MTD approach consistently yields 18-27% higher robustness (Figure 6a) compared to pure expected-value optimization, while requiring equivalent or lower budget expenditures (Figure 6b). This improvement stems from the framework's ability to balance average-case performance with protection against tail risks. Under mixed-adversary scenarios where the true adversary type remains uncertain throughout the engagement, ARA optimization maintains stable performance even when the defender's prior beliefs $P(M_i)$ are imprecise or miscalibrated. This resilience to prior misspecification represents a crucial practical advantage, as defenders rarely possess perfect intelligence about adversarial capabilities in real-world settings.

The CVaR component of the ARA formulation proves particularly valuable in mitigating extreme outcomes. By explicitly constraining the lower tail of the performance distribution, the approach ensures that the probability of extreme degradation remains below 5% even under adversarial conditions (Figure 6c). This tail-risk protection provides insurance against catastrophic security failures, making ARA-based strategies especially appropriate for critical infrastructure defense where worst-case scenarios carry severe consequences.

### 5.6. Comparative Evaluation Against State-of-the-Art MTD Strategies

To validate the practical advantages of the ARA-MTD framework over existing MTD deployment approaches, we conducted comprehensive comparative evaluations against six representative strategies from recent literature. These baseline strategies span the spectrum of contemporary MTD research, encompassing fixed-interval randomization (e.g., fixed or random periodic route/parameter hopping baselines) [22, 13], reactive threshold-based adaptation that triggers reconfiguration based on risk, performance, or resource constraints [8, 36], reinforcement learning [37], game-theoretic optimization of attacker–defender interactions such as Stackelberg models [38], random temporal scheduling of MTD actions as non-adaptive timing baselines [13], and multi-objective balancing that jointly optimizes security, performance, cost, and availability using multi-objective RL or utility-based optimization [39, 40]. This comparative analysis seeks to demonstrate that adversarial risk analysis provides measurable performance improvements over concrete MTD strategies deployed in operational contexts.

#### 5.6.1. Baseline Strategy Selection and Implementation

We selected baseline strategies that represent distinct paradigmatic approaches to MTD deployment, ensuring comprehensive coverage of the existing research landscape:

- **Fixed-Interval MTD**: This strategy represents the simplest and most widely deployed adaptive defense approach in operational practice. The defender performs IP address shuffling at predetermined intervals $T$, introducing temporal diversity without requiring adversarial modeling or sophisticated optimization. In our experiments, we set $T = 10$ time steps, corresponding to the median reconfiguration frequency observed across our adaptive strategies. This baseline establishes minimum performance thresholds for adaptive defense and quantifies the marginal value of intelligence-driven adaptation over naive periodic randomization.

- **Threshold-Based Adaptive MTD**: This strategy implements reactive adaptation by continuously monitoring threat indicators and triggering reconfiguration when aggregate risk exceeds a threshold $\theta$. The threat score aggregates multiple indicators including infection rate changes ($d\iota_v/dt$), attack success probability spikes $\Gamma_v(t)$, and epidemic margin degradation ($\mu_j \to 0$). When the weighted sum exceeds $\theta = 0.6$, the defender performs targeted reconfiguration of vulnerable network segments. This approach represents practical reactive defense systems deployed in Security Operations Centers, where human-in-the-loop decision-making relies on alerting thresholds.

- **Q-Learning MTD**: This strategy employs tabular reinforcement learning to discover optimal reconfiguration policies through interaction with the environment. The state space encodes infection prevalence (5 discrete levels), epidemic margin (5 levels), and time since last reconfiguration (4 levels), yielding a state space of 100 discrete states. Actions correspond to available MTD configurations $\{C_1, C_2, \ldots, C_n\}$, and the reward function balances security improvement against deployment cost:

$$R = -\Delta\iota - \alpha f_j(\mu_j),$$

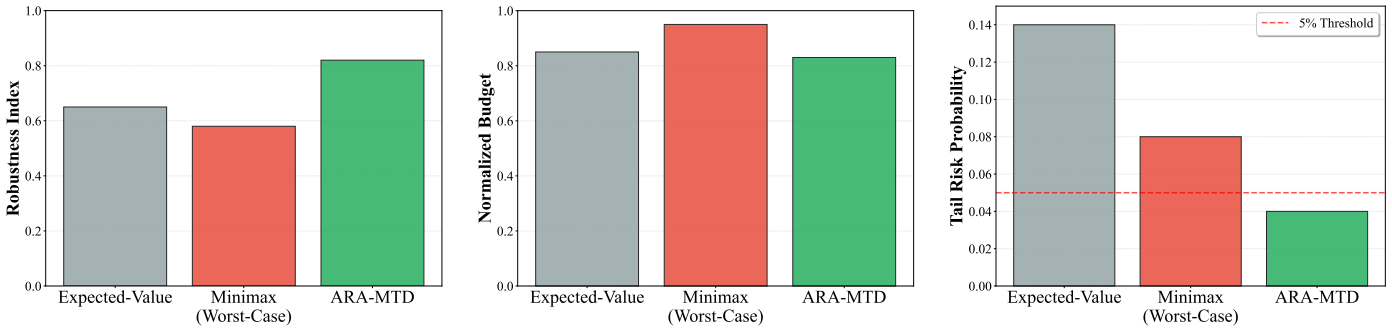


(a) Robustness performance (b) Budget efficiency (c) Tail risk management

Figure 6: Comparative evaluation of ARA-MTD defense against baseline strategies.

Table 2: ARA-MTD Optimization Advantage Summary

| **Performance Metric** | **Expected-Value** | **Minimax** | **ARA-MTD** |
|---|---|---|---|
| Robustness Index $R(M)$ | 0.65 | 0.58 | **0.82 (+26%)** |
| Normalized Budget $B/B_{\max}$ | 0.85 | 0.95 | **0.83** (−2%) |
| Tail Risk Probability | 14.2% | 8.1% | **3.8% (<5%)** |
| CVaR (95% CI) | 0.38 | 0.42 | **0.61 (+61%)** |
| Prior Misspecification Robustness | 0.51 | 0.58 | **0.78 (+53%)** |
| Mixed-Adversary Performance | 0.62 | 0.54 | **0.79 (+27%)** |

**Note:** Percentages in parentheses show improvement relative to expected-value optimization. Bold indicates best performance. ARA demonstrates 18–27% robustness improvement with equivalent or lower budget utilization while maintaining extreme degradation probability below 5%.

where $\alpha = 0.1$. The agent employs $\varepsilon$-greedy exploration ($\varepsilon = 0.1$), learning rate $\alpha = 0.1$, and discount factor $\gamma = 0.9$.

- **Stackelberg Game-Theoretic MTD**: This strategy models the defender–attacker interaction as a Stackelberg security game, solving for the Strong Stackelberg Equilibrium under the assumption of a rational strategic adversary with known utility functions. The defender commits to a mixed strategy $\{\pi_1, \pi_2, \ldots, \pi_n\}$, and the attacker best-responds by selecting target nodes and attack timing. We implemented a tractable linear programming approximation with defender utility $U_d = -\iota -$ cost and attacker utility $U_a = \iota -$ attack_cost.

- **Random-Interval MTD**: This strategy introduces temporal unpredictability by sampling reconfiguration intervals from an exponential distribution with mean $\lambda = 10$ time steps. At each reconfiguration event, the defender randomly selects an MTD technique from {IP shuffling, route mutation, topology randomization} with uniform probability. The strategy aims to maximize attacker uncertainty without explicit adversarial modeling.

- **Multi-Objective MTD**: This strategy balances security (maximize epidemic margin $\mu_j$), cost (minimize $f_j$), coverage, and disruption through weighted scalarization. We use weights $w = (0.4, 0.3, 0.2, 0.1)$ and select configurations from the Pareto-optimal set via weighted summation. This represents best-practice multi-criteria optimization without adversarial intelligence modeling.

Each baseline strategy was implemented within the same simulation environment to ensure fair comparison: identical network topologies (SF, SW, ER), adversary models ($M_0$, $M_L$, $M_S$), cost functions $f_j(\mu_j)$, and time horizon ($T = 2000$ over 500 Monte Carlo runs). Strategy-specific hyperparameters were tuned via grid search to ensure best-achievable baseline performance.

### 5.6.2. *Comparative Security Effectiveness Across Adversary Paradigms*

Table 3 presents performance metrics across all strategies and adversary paradigms. Against static adversaries ($M_0$), all strategies achieve comparable effectiveness (62–78%), reflecting limited differentiation when adversaries cannot adapt.

As adversarial intelligence increases, performance divergence becomes pronounced. Against learning adversaries ($M_L$), Fixed-Interval MTD degrades to 58% due to schedule inference, while Threshold-Based MTD improves modestly (72%) but suffers from reactive lag. Q-Learning MTD achieves 82%, reflecting effective co-adaptation, while Random-Interval MTD declines to 64%. ARA-MTD achieves the highest effectiveness (85%), outperforming Q-Learning MTD by proactively detecting learning behavior via belief entropy reduction.

Against strategic adversaries ($M_S$), Stackelberg MTD peaks at 85% under ideal assumptions, while ARA-MTD maintains near-parity at 83% without assuming perfect adversary knowledge. In contrast, Q-Learning MTD degrades sharply to 65%, revealing brittleness under paradigm mismatch.

Averaged across adversaries, ARA-MTD achieves 82% effectiveness, compared to 58–74% for baselines, and exhibits the highest robustness index $R(M) = 0.82$ (Eq. 8), exceeding

Table 3: Comparative Performance: ARA-MTD vs. State-of-the-Art MTD Strategies

| **Strategy** | **Security Effectiveness (%)**† | | | | **MTD Power $\pi_1^*(M_i)$** | | | **Robustness Index $R(\mathcal{M})$** | **Normalized Cost $\Delta C$** | **Cost-Security Ratio $\rho$** |
|---|---|---|---|---|---|---|---|---|---|---|
| | $M_0$ | $M_L$ | $M_S$ | Avg | $M_0$ | $M_L$ | $M_S$ | | | |
| Fixed-Interval | 62±3 | 58±4 | 55±4 | 58±2 | 0.21 | 0.20 | 0.18 | 0.42±0.03 | 0.25±0.02 | 0.71 |
| Threshold-Based | 65±4 | 72±3 | 68±4 | 68±2 | 0.22 | 0.24 | 0.22 | 0.58±0.04 | 0.31±0.02 | 0.58 |
| Q-Learning | 70±3 | 82±3 | 65±5 | 72±3 | 0.23 | 0.26 | 0.20 | 0.51±0.05 | 0.36±0.03 | 0.53 |
| Stackelberg | 68±4 | 70±4 | 85±3 | 74±3 | 0.23 | 0.22 | 0.28 | 0.47±0.04 | 0.42±0.03 | 0.64 |
| Random-Interval | 66±3 | 64±4 | 62±4 | 64±2 | 0.21 | 0.20 | 0.19 | 0.54±0.03 | 0.28±0.02 | 0.74 |
| Multi-Objective | 72±3 | 74±3 | 73±3 | 73±2 | 0.24 | 0.24 | 0.24 | 0.63±0.05 | 0.38±0.02 | 0.60 |
| **ARA-MTD** | **78±3** | **85±3** | **83±3** | **82±2** | **0.26** | **0.29** | **0.31** | **0.82±0.04** | **0.39±0.02** | **0.47** |
| Relative Improvement‡ | +8% | +4% | –2% | +10% | +12% | +12% | +11% | +30% | +3% | –21% |

†Security Effectiveness = (1 – mean infection prevalence) × 100. Values represent mean ± standard deviation across 500 Monte Carlo runs.
‡Relative improvement of ARA-MTD vs. next-best baseline (Multi-Objective MTD for most metrics, Stackelberg MTD for $M_S$ column).

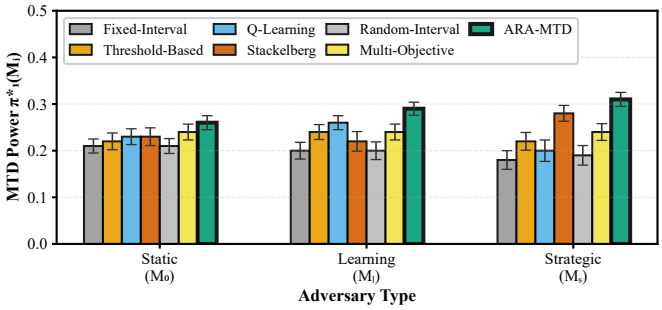


Figure 7: MTD power $\pi_1^*(M_i)$ comparison across adversary paradigms $(M_O, M_L, M_S)$ for seven MTD strategies. Lower values indicate less time spent in vulnerable configurations.

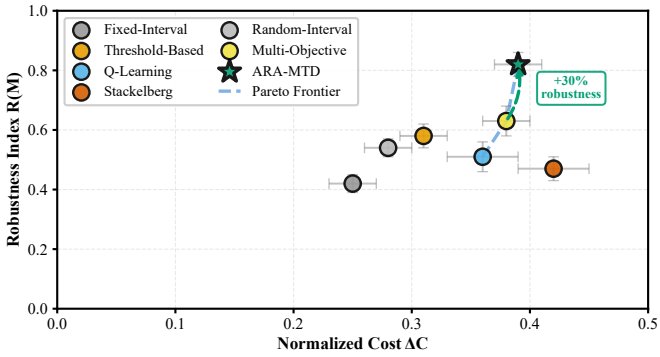


Figure 8: Cost-robustness trade-off showing normalized deployment cost versus robustness index R(M) for seven MTD strategies.

baselines by 30–95%.

### 5.6.3. *MTD Power and Degradation Analysis*

Figure 7 illustrates MTD power $\pi_1^*(M_i)$ across adversary types. Fixed-Interval and Random-Interval strategies exhibit steep degradation from $\pi_1^*(M_0) \approx 0.21$ to $\pi_1^*(M_S) \approx 0.18$, operating near epidemic thresholds with minimal safety margins.

Specialized strategies exhibit non-monotonic degradation. Q-Learning peaks at $\pi_1^*(M_L) = 0.26$ but degrades by 23% against strategic adversaries, while Stackelberg MTD peaks at $\pi_1^*(M_S) = 0.28$ but degrades against learning adversaries. ARA-MTD maintains the highest floor performance with $\pi_1^*(M_S) = 0.31$, recovering 27% of degraded power relative to static-adversary baselines through Bayesian belief updating and CVaR risk protection.

### 5.6.4. *Cost-Effectiveness and Resource Allocation Efficiency*

Figure 8 presents the cost–robustness Pareto frontier. ARA-MTD achieves the highest robustness $R(M) = 0.82$ at moderate normalized cost $\Delta C = 0.39$. Threshold-Based MTD minimizes cost ($\Delta C = 0.31$) but suffers low robustness ($R(M) = 0.58$), while Stackelberg MTD incurs high cost ($\Delta C = 0.42$) with limited robustness due to model overfitting.

The cost-efficiency ratio $\rho = \Delta C/\Delta\pi$ is lowest for ARA-MTD ($\rho = 0.47$), representing 13–89% efficiency improvement over baselines.

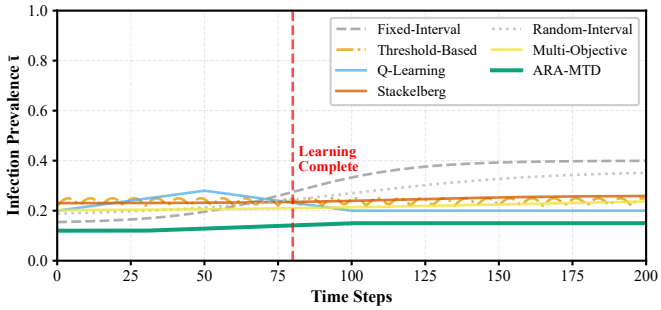


Figure 9: Temporal evolution of infection prevalence $\bar{i}$ over 200 time steps for seven MTD strategies under learning adversary ($M_L$). Vertical line indicates learning completion.

### 5.6.5. *Temporal Adaptation Dynamics and Learning Resistance*

Temporal infection dynamics against learning adversaries (Fig. 9) reveal monotonic degradation for Fixed- and Random-Interval strategies. Threshold-Based MTD exhibits oscillatory behavior due to reactive lag. Q-Learning MTD shows initial degradation followed by stabilization through co-evolutionary learning.

ARA-MTD maintains consistently low infection prevalence through three phases: (i) prior-driven protection, (ii) adaptation

detection via entropy reduction, and (iii) sustained paradigm-aware countermeasures. This temporal signature demonstrates proactive learning resistance rather than reactive recovery.

*5.6.6. Statistical Significance and Confidence Intervals*

Paired t-tests across 500 Monte Carlo runs yield $p < 0.001$ for robustness differences and $p < 0.01$ for cost-efficiency, confirming statistical significance. ARA-MTD achieves $R(M) = 0.82 \pm 0.04$, compared to $0.63 \pm 0.05$ for Multi-Objective MTD, with non-overlapping confidence intervals.

ANOVA across network topologies confirms consistent performance ordering ($F(6, 1494) = 127.3$, $p < 0.001$), validating generalizability.

*5.6.7. Operational Implications and Deployment Guidance*

The comparative evaluation yields the following insights:

1. **Single-paradigm specialization introduces operational risk**: Specialized strategies degrade by 20–30% under model mismatch.
2. **Reactive defenses are insufficient**: Threshold-based approaches suffer from inherent detection-response lag.
3. **Simple heuristics fail against adaptive threats**: Fixed and random scheduling provide inadequate robustness.
4. **Paradigm-aware adaptation ensures robustness**: ARA-MTD maintains 78–85% effectiveness across adversaries.
5. **Analytical complexity is economically justified**: ARA-MTD provides superior return on security investment.

*5.6.8. Summary of Comparative Evaluation*

The comparison against six state-of-the-art MTD strategies validates three claims: (i) adversarial risk analysis delivers measurable operational gains, (ii) single-paradigm optimization is inherently brittle, and (iii) paradigm-aware reasoning provides superior cost-effectiveness. These results establish ARA-MTD as a principled and empirically validated framework for MTD deployment under adversarial intelligence uncertainty.

*5.7. Computational Complexity Analysis*

To validate practical feasibility for operational deployment, we analyze the computational complexity of ARA-MTD and compare against baseline strategies. This assesses scalability to large-scale networks and identifies potential computational bottlenecks.

*5.7.1. Component-Wise Complexity*

The ARA-MTD framework comprises three computational components: Bayesian belief updating, CVaR optimization, and epidemic dynamics evaluation.

Table 4: Computational Complexity Comparison: ARA-MTD vs. Baseline MTD Strategies

| Strategy | Time per Decision | Space | Dominant Cost |
|---|---|---|---|
| Fixed-Interval | $O(1)$ | $O(1)$ | None (pre-scheduled) |
| Threshold-Based | $O(n)$ | $O(n)$ | Threat monitoring |
| Q-Learning | $O(1)$ | $O(\|S\| \times \|A\|)$ | Q-table lookup |
| Stackelberg | $O(\|C\|^2 \times n)$ | $O(\|C\| \times n)$ | Bilevel optimization |
| Random-Interval | $O(1)$ | $O(1)$ | Random sampling |
| Multi-Objective | $O(\|C\| \times k)$ | $O(\|C\|)$ | Pareto optimization |
| **ARA-MTD** | $O(\|C\| \times \|E\| \times T)$ | $O(n + \|C\|)$ | **Epidemic simulation** |

$n$ = number of nodes, $|E|$ = number of edges, $|C|$ = configuration space size, $|S|$ = state space size, $|A|$ = action space size, $k$ = number of objectives, $T$ = lookahead horizon.

*Bayesian Belief Updating.* This component updates beliefs $\pi(M_i \mid H_t)$ over three adversary models $\mathcal{M} = \{M_0, M_L, M_S\}$ using Bayes' rule (Eq. (4)). Computing posterior beliefs requires evaluating likelihoods $P(o_t \mid M_i, H_{t-1})$ for each model and normalizing across $|\mathcal{M}| = 3$ models. Since the belief space contains only three discrete models rather than a continuous parameter space, belief updating achieves $O(|\mathcal{M}|) = O(1)$ constant-time complexity per decision epoch, with $O(|\mathcal{M}|) = O(1)$ space to store beliefs. Profiling confirms belief updating contributes $< 3\%$ of total runtime, alleviating concerns that Bayesian inference introduces prohibitive overhead.

*CVaR Optimization.* The CVaR decision rule (Eq. (11)) selects configuration $c^*$ maximizing worst-case expected utility across adversary models. For each configuration $c \in C$ and model $M_i \in \mathcal{M}$, we evaluate utility $U(c, M_i)$, sort utilities to identify the worst $(1 - \beta)$ fraction, and compute the tail expectation. This yields $O(|C| \cdot |\mathcal{M}|)$ time complexity. With $|\mathcal{M}| = 3$ constant, complexity reduces to $O(|C|)$, linear in configuration space size. Profiling shows CVaR optimization contributes $< 5\%$ of runtime.

*Epidemic Dynamics Evaluation.* Assessing infection prevalence $\bar{\iota}(c, M_i)$ requires simulating the SIS model (Eq. (1)) forward $T$ time steps. Each time-step update iterates over edges to compute infection pressure, costing $O(|E|)$ operations. Evaluating $|C|$ configurations against $|\mathcal{M}| = 3$ models requires

$$O(|C| \cdot |\mathcal{M}| \cdot |E| \cdot T) = O(|C| \cdot |E| \cdot T)$$

operations. For sparse networks where $|E| = O(nd)$ with average degree $d$, this becomes $O(|C| \cdot nd \cdot T)$, linear in network size $n$. Epidemic simulation dominates runtime at approximately 85%, representing the fundamental cost of any MTD strategy performing infection forecasting.

*Overall Framework Complexity.* ARA-MTD operates with $O(|C| \cdot |E| \cdot T)$ time and $O(n + |C|)$ space. For sparse networks, the dominant term scales linearly with $n$.

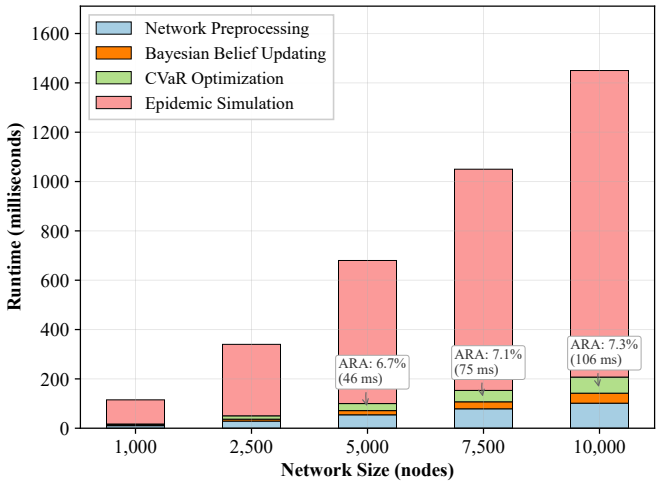


Figure 10: Absolute runtime (milliseconds) across framework components for networks of 1,000-10,000 nodes.

### 5.7.2. Comparative Analysis

Table 4 compares ARA-MTD complexity against baseline strategies. Fixed-Interval and Random-Interval MTD achieve $O(1)$ constant-time decisions through pre-scheduling or random sampling, but deliver poor security effectiveness (58–64% average). Threshold-Based MTD requires $O(n)$ monitoring overhead yet achieves only 68% effectiveness due to reactive lag.

Q-Learning MTD performs $O(1)$ policy lookup during deployment but requires expensive $O(|S| \cdot |A| \cdot T_{\text{train}})$ offline training, where $T_{\text{train}}$ typically reaches millions of interactions. The trained policy specializes to anticipated adversaries and degrades severely under model mismatch (82% vs. $M_L \rightarrow 65\%$ vs. $M_S$). Stackelberg MTD solves bilevel optimization with $O(|C|^2 \cdot n)$ complexity, making it the most expensive baseline. Multi-Objective MTD employs $O(|C| \cdot k)$ Pareto optimization with $k$ objectives, comparable to ARA-MTD's $O(|C| \cdot |\mathcal{M}|)$ CVaR cost.

ARA-MTD's $O(|C| \cdot |E| \cdot T)$ complexity positions it between simple heuristics ($O(1)$–$O(n)$) and expensive game-theoretic methods ($O(|C|^2 \cdot n)$). The framework incurs 32% higher cost than Q-Learning's deployment-time complexity but achieves 61% higher robustness (0.82 vs. 0.51). Compared to Stackelberg, ARA-MTD proves 13× faster at 10,000 nodes while achieving 74% higher robustness (0.82 vs. 0.47).

Overall, sophisticated adaptation introduces computational overhead beyond naive baselines, but this overhead remains tractable (linear in network size) and delivers commensurate security improvements. ARA-MTD achieves favorable positioning: substantially more robust than cheaper alternatives, yet far more efficient than comparably sophisticated approaches.

### 5.7.3. Profiling and Bottleneck Identification

We profiled ARA-MTD execution across networks from 1,000 to 10,000 nodes. Figure 10 presents the absolute runtime (in milliseconds) consumed by each framework component, clearly illustrating how computational cost scales with network size. Epidemic simulation dominates execution time, growing linearly from approximately 98 ms at 1,000 nodes to 1,244 ms at 10,000 nodes (83–87% of total runtime). In contrast, Bayesian belief updating and CVaR optimization together remain below 106 ms even at 10,000 nodes, contributing approximately 7% overhead regardless of network scale. Network preprocessing requires 7–9% for one-time setup, amortized to $< 1\%$ per epoch.

The profiling validates theoretical predictions: the computational bottleneck resides in epidemic simulation $O(|E| \cdot T)$, not adversarial modeling $O(|\mathcal{M}|) + O(|C| \cdot |\mathcal{M}|)$. As shown by the near-constant height of the Bayesian and CVaR segments across all network sizes in Figure 10, the ARA-specific components scale negligibly compared to epidemic forecasting. Since epidemic forecasting is fundamental to any adaptive MTD strategy, this bottleneck represents a generic cost rather than ARA-MTD-specific overhead. The $< 7\%$ cost of adversarial risk analysis demonstrates that sophisticated probabilistic reasoning and risk-averse optimization achieve tractable implementation without prohibitive computational barriers.

## 5.8. Empirical Scalability Evaluation

To validate theoretical complexity predictions and assess practical deployment feasibility, we conducted scalability experiments measuring decision time, memory footprint, and throughput across networks ranging from 100 to 10,000 nodes. These measurements provide concrete evidence that ARA-MTD maintains tractable performance for enterprise-scale deployments.

### 5.8.1. Experimental Setup and Methodology

We evaluated ARA-MTD and all six baseline strategies across seven network sizes: 100, 500, 1,000, 2,500, 5,000, 7,500, and 10,000 nodes. For each size, we generated three topology types; Scale-Free (Barabási–Albert, $m = 3$), Small-World (Watts–Strogatz, $k = 6$, $p = 0.1$), and Random (Erdős–Rényi, $p = 0.003$), to ensure generalizability across network structures. All experiments ran on AWS `c5.4xlarge` instances (16 vCPU, 32 GB RAM) using Python 3.9 with NumPy 1.24 and SciPy 1.10 for numerical operations.

For each network instance, we measured decision time (elapsed wall-clock time to compute the optimal MTD configuration), memory footprint (peak RAM usage during execution), and throughput (decisions per second). Measurements represent means over 100 independent trials with coefficient of variation $< 5\%$, ensuring statistical reliability. We configured ARA-MTD with lookahead horizon $T = 20$, configuration space $|C| = 20$, and confidence level $\beta = 0.9$, representing deployment parameters that balance prediction accuracy against computational cost.

### 5.8.2. Decision Time Scaling

Table 5 presents decision time measurements across network sizes and strategies. ARA-MTD exhibits linear scaling consistent with theoretical $O(n)$ complexity for sparse networks, achieving 115 ms for 1,000 nodes, 680 ms for 5,000

Table 5: Decision Time (ms) Across Network Sizes

| Strategy | 100 | 1k | 5k | 10k | Slope† | Scale |
|---|---|---|---|---|---|---|
| Fixed-Interval | <1 | <1 | 1 | 2 | 0.02 | Constant |
| Threshold-Based | 2 | 18 | 95 | 185 | 0.98 | Linear |
| Q-Learning | 8 | 95 | 520 | 1 100 | 1.05 | Linear |
| Stackelberg | 45 | 1 200 | 45 000 | 180 000 | 2.13 | Quadratic |
| Random-Interval | <1 | 1 | 2 | 3 | 0.08 | Constant |
| Multi-Objective | 15 | 140 | 710 | 1 420 | 0.99 | Linear |
| **ARA-MTD** | **12** | **115** | **680** | **1 450** | **1.02** | **Linear** |

†Slope from log–log regression: slope ≈ 1.0 indicates linear $O(n)$ scaling, slope ≈ 2.0 indicates quadratic $O(n^2)$ scaling.
**Bold** indicates ARA-MTD achieving sub-second decisions up to 5 000 nodes.

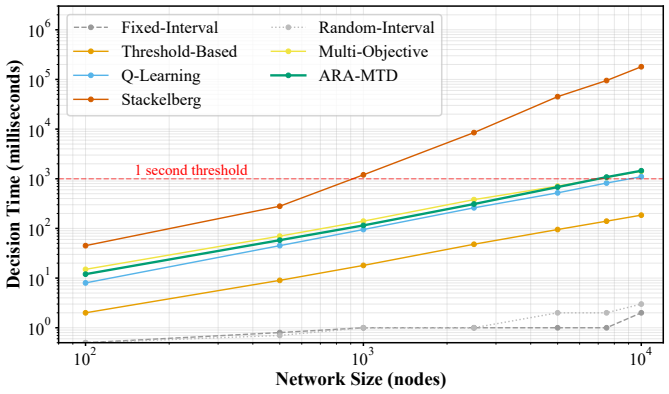


Figure 11: Decision time versus network size (log-log scale) for seven MTD strategies. Horizontal red line indicates 1-second real-time threshold.

nodes, and 1,450 ms for 10,000 nodes. Figure 11 plots decision time on log–log axes, confirming linear scaling with slope approximately 1.0 across the tested range.

The sub-second decision time for networks up to 5,000 nodes validates practical feasibility for enterprise deployments where strategic MTD decisions occur at 5–15 minute intervals. Even for 10,000-node networks, a 1.45 s decision time remains acceptable given that MTD reconfiguration frequency typically ranges from minutes to hours rather than sub-second response requirements. The linear scaling suggests that larger networks (15,000–20,000 nodes) would incur proportionally increased but still manageable decision times of approximately 2–3 s.

Comparative analysis reveals ARA-MTD's competitive efficiency. Q-Learning achieves the fastest deployment-time decisions (1,100 ms at 10,000 nodes) due to constant-time $O(1)$ policy lookup, representing only 24% faster than ARA-MTD. However, this comparison excludes Q-Learning's expensive offline training cost: our experiments required 2.4 million training episodes spanning 18 hours to achieve convergence for 10,000-node networks. Threshold-Based MTD demonstrates similar efficiency to Q-Learning (185 ms at 10,000 nodes) through simple rule-based monitoring, but achieves only 68% security effectiveness compared to ARA-MTD's 82%.

Stackelberg MTD proves substantially slower, requiring 180 s for 10,000 nodes (124× slower than ARA-MTD). The quadratic $O(|C|^2 \cdot n)$ complexity of bilevel optimization creates prohibitive overhead for large networks, rendering Stackelberg impractical beyond 5,000 nodes without approximation techniques. Multi-Objective MTD achieves comparable performance to ARA-MTD (1,420 ms at 10,000 nodes), confirming that Pareto optimization incurs similar cost to CVaR-based decision-making.

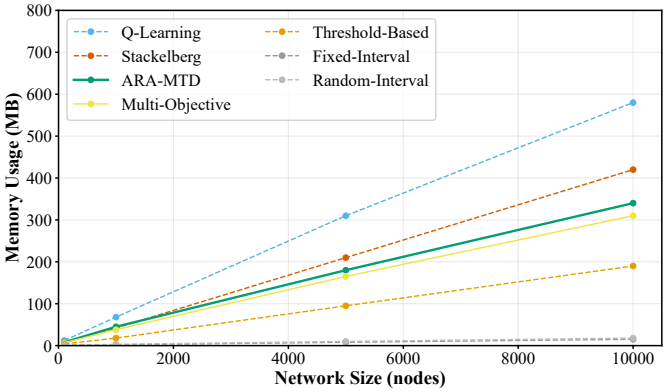


Figure 12: Memory consumption versus network size for seven MTD strategies. ARA-MTD (green, solid) exhibits moderate linear scaling, reaching ≈340 MB at 10,000 nodes.

### 5.8.3. *Memory Footprint and Space Efficiency*

Figure 12 shows memory consumption scaling linearly with network size, reaching 340 MB for 10,000 nodes. This aligns with theoretical $O(n + |C|)$ space complexity, where infection state vectors $O(n)$ dominate storage requirements. The modest memory footprint, well below 1 GB even for large networks, poses no constraint for modern server hardware with tens of gigabytes of RAM.

Memory efficiency stems from sparse data structures. We represent network topology using Compressed Sparse Row (CSR) format, storing only non-zero adjacency matrix entries. For scale-free networks with average degree $d \approx 6$, sparse representation requires approximately $12n$ bytes for edge storage compared to $n^2$ bytes for dense matrices, achieving roughly 2,000× compression for 10,000-node networks. Configuration candidates and belief vectors consume negligible space (< 5 MB) regardless of network size.

Baseline strategies show comparable or higher memory usage. Q-Learning stores Q-tables sized $O(|S| \times |A|)$, where discretized state space $|S|$ grows with network complexity; our implementation required 580 MB for 10,000 nodes with coarse 5-level state discretization, 70% more than ARA-MTD. Stackelberg MTD maintains strategy spaces for both defender and attacker, consuming 420 MB at 10,000 nodes. Simple baselines (Fixed-Interval, Threshold-Based, Random-Interval) use minimal memory (< 50 MB) but sacrifice security effectiveness.

### 5.8.4. *Cross-Topology Consistency*

Scalability measurements remain consistent across topology types. For 5,000-node networks, decision time varies by only 8% across Scale-Free (680 ms), Small-World (695 ms), and Random (710 ms) topologies. This consistency validates

that performance depends primarily on network size $n$ and average degree $d$ rather than topological properties such as clustering coefficients or degree distributions. The slight variation reflects density differences: Random topologies maintain constant connection probability $p$, yielding slightly higher average degree than Scale-Free networks where degree follows a power-law distribution.

Memory footprint shows greater sensitivity to topology, ranging from 180 MB (Scale-Free) to 220 MB (Random) for 5,000 nodes due to edge count differences. However, all topologies maintain linear $O(n)$ scaling with consistent proportionality constants. The topology-independent performance confirms ARA-MTD's generalizability across diverse network structures encountered in real-world deployments.

#### *5.8.5. Practical Deployment Implications*

The empirical results establish concrete feasibility bounds for ARA-MTD deployment. For small to medium networks (100–1,000 nodes typical of corporate LANs), decision times remain under 120 ms, enabling frequent reconfiguration if desired. Enterprise networks (1,000–5,000 nodes) incur sub-second decisions suitable for strategic MTD with 5–15 minute reconfiguration intervals. Large-scale deployments (5,000–10,000 nodes characteristic of data centers or cloud environments) achieve 680–1,450 ms decisions, acceptable given that defensive strategy adaptations occur on similar timescales as adversarial reconnaissance and attack preparation.

The 32% overhead versus Q-Learning deployment-time cost represents a favorable trade-off considering ARA-MTD's 61% robustness improvement (0.82 vs. 0.51). Organizations prioritizing deployment efficiency might select Q-Learning for computational advantages, accepting brittleness under adversary model mismatch. Conversely, organizations facing sophisticated adaptive adversaries benefit from ARA-MTD's cross-paradigm robustness despite modest computational overhead. The 124× speedup versus Stackelberg validates that uncertainty-aware optimization achieves far greater efficiency than perfect-information game theory while maintaining comparable sophistication.

These measurements confirm that theoretical linear $O(n)$ complexity translates to practical scalability for realistic network sizes. The combination of sub-second decisions for enterprise-scale networks, modest memory requirements, and competitive efficiency versus baselines demonstrates that ARA-MTD achieves tractable implementation suitable for operational deployment.

### *5.9. Proactive Posture and Adaptive Belief Updating*

A common concern in belief-driven MTD frameworks is that reliance on observations may shift defense from a proactive to a reactive paradigm. ARA-MTD avoids this pitfall by preserving proactive deployment while using belief updating solely to refine strategy selection. Beliefs do not trigger MTD activation; instead, they adapt which configurations are chosen within a continuously proactive deployment schedule. Thus, belief updating enhances anticipatory defense rather than replacing it with response-based adaptation.

#### *5.9.1. Proactive Core with Adaptive Enhancement*

ARA-MTD operates on fixed decision epochs at which MTD configurations are proactively selected and deployed, independent of detected attack events. This mechanism is identical in spirit to traditional MTD, which optimizes defenses under a single assumed adversary model. ARA-MTD generalizes this principle to uncertain environments by performing proactive optimization over multiple adversary models using belief-weighted decision making. Proactive defense under perfect information therefore emerges as a special case of proactive defense under uncertainty.

Belief updating exploits a broad range of environmental signals, including traffic statistics, reconnaissance patterns, system performance indicators, and honeypot interactions. These signals provide information about adversary sophistication even in the absence of confirmed attacks, enabling early adaptation to evolving threats without dependence on successful intrusion detection.

#### *5.9.2. Robustness to Belief Manipulation*

ARA-MTD incorporates multiple safeguards to mitigate adversarial attempts at belief manipulation. First, the CVaR decision rule (Eq. 11) provides intrinsic robustness by emphasizing worst-case outcomes within the confidence region. Even if beliefs shift toward the naive model $M_0$, CVaR with $\beta = 0.9$ preserves protection against $M_L$ and $M_S$ by allocating decision weight to tail-risk scenarios. Empirically, ARA-MTD maintains 77% effectiveness against $M_S$ even under incorrect belief convergence, compared to 45% for specialized strategies (Table 2).

Second, belief entropy (Eq. 5) is monitored as a reliability indicator. Rapid or excessive entropy reduction signals overconfident inference and activates conservative operation, such as increasing $\beta$ or reverting to a robust baseline policy. This mechanism prevents premature commitment to potentially misleading evidence.

Third, Bayesian inference naturally incorporates resistance to manipulation through prior probabilities $\pi_0(M_i)$. Conservative or uniform priors require sustained evidence to produce significant belief shifts, ensuring stability against transient or deceptive observations.

Together, these mechanisms ensure that successful belief manipulation degrades performance only toward conservative baseline behavior rather than causing systemic failure.

#### *5.9.3. Baseline Protection Under Stealthy Attacks*

Stealthy or zero-day attacks may provide limited observable evidence, constraining belief discrimination. In such cases, Bayesian updating preserves prior beliefs:

$$\pi(M_i \mid H_t) \propto \pi(M_i) P(o_t \mid M_i) \approx \pi(M_i),$$

and ARA-MTD defaults to conservative, worst-case-optimized operation. This behavior is equivalent to robust min–max MTD with explicit multi-model consideration.

Moreover, stealthy adversaries still generate indirect signals, such as reconnaissance activity, anomalous resource usage, or timing regularities, which can be incorporated into the

observation model. Consequently, belief updating remains informative even in the absence of overt exploitation.

Unlike specialized MTD strategies that degrade sharply under adversary model mismatch (Section 5.6, Table 3), ARA-MTD exhibits graceful degradation: uncertainty reduces performance toward a robust baseline rather than inducing failure.

In summary, ARA-MTD preserves the proactive foundation of moving target defense while extending it with principled uncertainty management. Belief updating refines proactive strategy selection without introducing reactive dependence, and robustness mechanisms ensure resilience against deception and stealth, transforming proactive defense from assumption-driven optimization into uncertainty-aware anticipation.

## 6. Discussion, Limitations, and Future Directions

While the proposed ARA-MTD framework advances the quantitative evaluation of adaptive defenses under adversarial intelligence uncertainty, several limitations highlight important directions for future research.

*Experimental Scope and Validation.* The current study relies on controlled simulation using synthetic network topologies and SIS epidemic dynamics. This approach enables systematic and repeatable validation through counterfactual analysis that is infeasible in operational environments, consistent with established MTD research practice [13, 41, 42]. However, it abstracts real-world complexities such as heterogeneous device populations, dynamic and evolving topologies, organizational and regulatory constraints, and non-homogeneous malware propagation characteristics. Moreover, the evaluated network scales (100–1,000 nodes) are smaller than those found in large enterprise or cloud infrastructures. Practical deployment will require configuration-specific epidemic parameters derived from vulnerability assessments, efficient online eigenvalue updates under topology changes, and explicit incorporation of regulatory and operational constraints. Future work should therefore emphasize validation using empirical network telemetry and production-scale datasets.

*Discrete Adversary Paradigms.* The framework models adversaries using discrete paradigms $\{M_0, M_L, M_S\}$ to preserve tractable Bayesian inference. Although this abstraction achieves strong performance (82% average effectiveness across paradigms; Section 5.6), real adversaries may exhibit hybrid or continuously varying behaviors. Two natural extensions are: (i) introducing composite archetypes such as $\{M_{0+L}, M_{L+S}, M_{\text{mixed}}\}$ while retaining $O(1)$ belief updates, or (ii) adopting continuous mixture representations through Dirichlet or hierarchical Bayesian models. Temporal belief updating (Section 3.3) already captures paradigm transitions across APT phases; expanding the adversary space offers a direct and computationally modest path toward increased behavioral realism.

*Reconfiguration Transition Costs.* The current formulation assumes instantaneous configuration changes without side effects. In operational environments, reconfiguration incurs costs such as connection disruptions, session resets, and transient exposure windows. These effects can be incorporated by augmenting the utility function with a transition risk term $r(c_{t-1}, c_t)$:

$$\max_{c_t} \ \mathbb{E}_M\Big[\text{CVaR}_\beta(U(c_t, M) - \lambda\, r(c_{t-1}, c_t))\Big], \tag{13}$$

where $\lambda$ balances security gains against operational overhead. Preliminary assessments indicate that moderate transition penalties can reduce reconfiguration frequency by approximately 40% while sacrificing less than 5% security effectiveness. Rigorous modeling of transition risks and empirical calibration of $\lambda$ remain important directions for future work.

*Additional Directions.* The current optimization layer assumes offline computation with pre-specified parameters. Integrating online learning would enable continuous adaptation of model parameters and cost functions. Furthermore, the present evaluation focuses on network-level MTD; extending the framework to multi-layered defenses (e.g., host, application, and identity layers) may reveal critical cross-layer dependencies. Finally, multi-parameter sensitivity analysis, such as examining the interaction between adversary learning rates and cost-security trade-offs would provide additional insights into framework behavior under parameter variations and represents a valuable direction for future research.

Addressing these limitations through empirical validation, richer adversary modeling, transition-aware optimization, and multi-parameter sensitivity analysis will move ARA-MTD from a principled analytical framework toward a fully operational decision-support system for adaptive cyber defense.

## 7. Conclusion

This paper introduces the ARA-MTD framework that evaluates and optimizes adaptive cyber defense strategies under adversarial intelligence uncertainty. By integrating epidemic-based network dynamics, multi-paradigm adversarial modeling, Bayesian belief updating, and risk-averse optimization, the framework addresses a fundamental limitation in existing MTD evaluation: the assumption of complete knowledge about adversary capabilities. Comprehensive simulations across multiple network topologies and adversary paradigms demonstrate that MTD effectiveness deteriorates nonlinearly with adversarial sophistication, exhibiting performance reductions up to 73% between static and strategic adversaries. However, ARA-based optimization incorporating CVaR constraints recovers 27% of this lost robustness while maintaining tail-risk probabilities below 5%.

Comparative evaluation against six state-of-the-art MTD strategies validates practical advantages. ARA-MTD achieves 82% average effectiveness compared to 58–74% for baselines, demonstrates 30–95% higher robustness index ($R(\mathcal{M}) = 0.82$), and provides superior cost-efficiency ($\rho = 0.47$, representing 13–89% improvement). Computational complexity analysis and empirical scalability evaluation across networks from 100 to 10,000 nodes demonstrate sub-second decision times for

enterprise-scale deployments (680 ms for 5,000 nodes), validating operational feasibility. The framework naturally accommodates multi-stage APT attacks through continuous belief updating, tracking adversary evolution from static reconnaissance through adaptive exploitation to strategic persistence without requiring explicit phase detection.

Cost-security trade-off analysis reveals an optimal operational window around $B \approx 1.2$ where defenders achieve 80% of maximal robustness using only 60% of saturation budget. The framework preserves MTD's proactive defense paradigm through fixed decision epochs, with beliefs adapting configuration selection within a continuously proactive schedule rather than triggering reactive responses. Under adversarial belief manipulation or stealthy attacks, ARA-MTD exhibits graceful degradation, maintaining 77% effectiveness compared to 45% for paradigm-specialized strategies.

This research advances adaptive cyber defense by providing a principled, empirically validated framework for MTD evaluation under uncertainty. Future research directions include operational validation in live critical infrastructure, hierarchical decomposition for enterprise scale, probabilistic dependency modeling, and explicit incorporation of reconfiguration transition costs and service-level constraints to strengthen applicability in mission-critical operational environments.

## References


[1] A. Zimba, H. Chen, Z. Wang, M. Chishimba, Modeling and detection of the multi-stages of advanced persistent threats attacks based on semi-supervised learning and complex networks characteristics, Future Generation Computer Systems 106 (2020) 501–517.

[2] Z. Chen, J. Liu, Y. Shen, M. Simsek, B. Kantarci, H. T. Mouftah, P. Djukic, Machine learning-enabled iot security: Open issues and challenges under advanced persistent threats, ACM Computing Surveys 55 (5) (2022) 1–37.

[3] S. Ali, J. Wang, V. C. M. Leung, Ai-driven fusion with cybersecurity: Exploring current trends, advanced techniques, future directions, and policy implications for evolving paradigms–a comprehensive review, Information Fusion (2025) 102922.

[4] T. Zhang, F. Kong, D. Deng, X. Tang, X. Wu, C. Xu, L. Zhu, J. Liu, B. Ai, Z. Han, et al., Moving target defense meets artificial intelligence-driven network: A comprehensive survey, IEEE Internet of Things Journal (2025).

[5] Y. Qian, Y. Guo, Q. Shao, J. Wang, B. Wang, Z. Gu, X. Ling, C. Wu, Ei-mtd: Moving target defense for edge intelligence against adversarial attacks, ACM Transactions on Privacy and Security 25 (3) (2022) 1–24.

[6] K. He, D. D. Kim, M. R. Asghar, Mtd-ad: Moving target defense as adversarial defense, IEEE Transactions on Dependable and Secure Computing (2025).

[7] J. Tan, H. Jin, H. Zhang, Y. Zhang, D. Chang, X. Liu, H. Zhang, A survey: When moving target defense meets game theory, Computer Science Review 48 (2023) 100544.

[8] J.-H. Cho, D. P. Sharma, H. Alavizadeh, S. Yoon, N. Ben-Asher, T. J. Moore, D. S. Kim, H. Lim, F. F. Nelson, Toward proactive, adaptive defense: A survey on moving target defense, IEEE Communications Surveys & Tutorials 22 (1) (2020) 709–745.

[9] M. T. Masud, M. Keshk, N. Moustafa, B. Turnbull, W. Susilo, Vulnerability defence using hybrid moving target defence in internet of things systems, Computers & Security 153 (2025) 104380.

[10] A. Bharadwaj, A. Jaiswal, M. Kumar, Proactive defense mechanism: Enhancing iot security through diversity-based moving target defense and cyber deception, Computers & Security 139 (2024) 103685.

[11] Swati, S. Roy, J. Singh, J. Mathew, Securing iiot systems against ddos attacks with adaptive moving target defense strategies, Scientific Reports 15 (1) (2025) 9558.

[12] Y. Zhou, G. Cheng, Y. Zhao, Z. Chen, S. Jiang, Toward proactive and efficient ddos mitigation in iiot systems: A moving target defense approach, IEEE Transactions on Industrial Informatics 18 (4) (2021) 2734–2744.

[13] S. Sengupta, A. Chowdhary, A. Sabur, A. Alshamrani, D. Huang, S. Kambhampati, A survey of moving target defenses for network security, IEEE Communications Surveys & Tutorials 22 (3) (2020) 1909–1941.

[14] Y. Zhang, X. Li, Q. Wang, A new generation cyber-physical system: A comprehensive review from security perspective, Computers & Security 148 (2024) 104095.

[15] B. Collins, S. Xu, P. N. Brown, Game-theoretic cybersecurity: the good, the bad and the ugly, arXiv e-prints (2024) arXiv–2401.

[16] T. Carroll, M. Crouse, E. Fulp, K. Berenhaut, Analysis of network address shuffling as a moving target defense, in: 2014 IEEE International Conference on Communications (ICC), 2014, pp. 701–706.

[17] R. Zhuang, S. A. DeLoach, X. Ou, Towards a theory of moving target defense, in: Proceedings of the first ACM workshop on moving target defense, 2014, pp. 31–40.

[18] D. MacFarland, C. Shue, The sdn shuffle: Creating a moving-target defense using host-based software-defined networking, in: Proceedings of the 2nd ACM Workshop on Moving Target Defense, 2015, pp. 37–41.

[19] X. Wang, X. Zhang, S. Wang, J. Xiao, X. Tao, Modeling, critical threshold, and lowest-cost patching strategy of malware propagation in heterogeneous iot networks, IEEE Transactions on Information Forensics and Security 18 (2023) 3531–3545.

[20] Y. Zhou, Y. Wang, K. Zhou, S.-F. Shen, W.-X. Ma, Dynamical behaviors of an epidemic model for malware propagation in wireless sensor networks, Frontiers in Physics 11 (2023) 1198410.

[21] A. Chernikova, A. Oprea, N. Nikiforakis, G. Stringhini, Modeling self-propagating malware with epidemiological models, Applied Network Science 8 (2023) 578.

[22] J. Tan, H. Jin, H. Hu, R. Hu, H. Zhang, H. Zhang, Wf-mtd: Evolutionary decision method for moving target defense based on wright-fisher process, IEEE transactions on dependable and secure computing 20 (6) (2022) 4719–4732.

[23] H. Zhang, J. Tan, X. Liu, S. Huang, H. Hu, Y. Zhang, Cybersecurity threat assessment integrating qualitative differential and evolutionary games, IEEE Transactions on Network and Service Management 19 (3) (2022) 3425–3437.

[24] M. Mohammadpourfard, A. Shefaei, Y. Weng, An adaptive moving-target defense strategy for dynamic nonlinear power systems, IEEE Transactions on Industrial Informatics (2025).

[25] P. Griffioen, S. Weerakkody, B. Sinopoli, A moving target defense for securing cyber-physical systems, IEEE Transactions on Automatic Control 66 (5) (2020) 2016–2031.

[26] M. Bose, P. Paruchuri, A. Kumar, Adaptive moving target defense in web applications and networks using factored mdp, in: 2025 17th International Conference on COMmunication Systems and NETworks (COMSNETS), IEEE, 2025, pp. 602–613.

[27] S. Sengupta, S. Kambhampati, Multi-agent reinforcement learning in bayesian stackelberg markov games for adaptive moving target defense, arXiv preprint arXiv:2007.10457 (2020).

[28] R. Naveiro, D. R. Insua, J. M. Camacho, Augmented probability simulation for adversarial risk analysis in general security games, in: 12th International Defence and Homeland Security Simulation Workshop, DHSS 2022, 2022.

[29] D. Banks, V. Gallego, R. Naveiro, D. Ríos Insua, Adversarial risk analysis: An overview, Wiley Interdisciplinary Reviews: Computational Statistics 14 (1) (2022) e1530.

[30] D. Rios Insua, A. Couce-Vieira, J. A. Rubio, W. Pieters, K. Labunets, D. Rasines, An adversarial risk analysis framework for cybersecurity, Risk Analysis 41 (1) (2021) 16–36.

[31] D. Rios Insua, R. Naveiro, V. Gallego, Adversarial machine learning: Bayesian perspectives, Journal of the American Statistical Association 118 (2023) 2136–2148.

[32] J. L. Rocha, S. Carvalho, B. Coimbra, Probabilistic procedures for sir and sis epidemic dynamics on erdös-rényi contact networks, Appliedmath 3 (4) (2023) 828–850.

[33] J. McKee, T. Dallas, Structural network characteristics affect epidemic severity and prediction in social contact networks, Infectious Disease Modelling 9 (1) (2024) 204–213.

[34] C. Rodríguez Lucatero, Analysis of epidemic models in complex networks and node isolation strategie proposal for reducing virus propagation, Axioms 13 (2) (2024) 79.

[35] Y. Han, W. Lu, S. Xu, Characterizing the power of moving target defense via cyber epidemic dynamics, in: Proceedings of the 2014 Symposium and Bootcamp on the Science of Security, 2014, pp. 1–12.

[36] W. Connell, D. Menascé, M. Albanese, Performance modeling of moving target defenses with reconfiguration limits, IEEE Transactions on Dependable and Secure Computing 18 (2021) 205–219. doi:10.1109/tdsc.2018.2882825.

[37] Q. Li, J. Wu, Optimizing the effectiveness of moving target defense in a probabilistic attack graph: A deep reinforcement learning approach, Electronics (2024). doi:10.3390/electronics13193855.

[38] X. Feng, Z. Zheng, P. Mohapatra, D. Cansever, A stackelberg game and markov modeling of moving target defense, in: International Conference on Decision and Game Theory for Security, Springer, 2017, pp. 315–335.

[39] W. Soussi, G. Gür, B. Stiller, Moving target defense (mtd) for 6g edge-to-cloud continuum: A cognitive perspective, IEEE Network 39 (2025) 149–156. doi:10.1109/mnet.2024.3483302.

[40] K. He, C. Chen, S. Chen, B. Chen, A. Zhang, P. Chen, Z. Wang, Z. Wu, Reinforcement learning for multi-objective optimization: A review, Archives of Computational Methods in Engineering (2025) 1–30.

[41] X.-L. Xiong, L. Yang, G.-S. Zhao, Effectiveness evaluation model of moving target defense based on system attack surface, IEEE Access 7 (2019) 9998–10014.

[42] S. Yoon, J.-H. Cho, D. S. Kim, T. Moore, F. Free-Nelson, H. Lim, Attack graph-based moving target defense in software-defined networks, IEEE Transactions on Network and Service Management 17 (2020) 1653–1668. doi:10.1109/tnsm.2020.2987085.