Published by

British Journal of Contemporary Research

BEXFORD PUBLISHING

Project DIANA: Digital Infrastructure and Adversarial Network Abuse. A Mathematical Analysis of AI-Driven Attacker Autonomy, DNS Security, and Network Openness

Original Research Article | Volume 1 | Issue 6 | 2026 | Article Number: 054

Accepted: 26 October 2026 | Published: 10 November 2026 | ISSN: 2979-8582 (Online)

Crossref : **https://doi.org/10.67693/BJCR-FKFLMAGC**

Ishma Deshmukh Britton

Pune University & British Computer Society (BCS), India

IDB 0009-0008-5068-4320

Correspondence: Ishma Deshmukh Britton, ishmawarrior@gmail.com

Abstract

The Domain Name System (DNS) is essential to Internet naming and service discovery, yet it is also implicated in malicious domain use, phishing, malware infrastructure, botnet coordination, and other forms of abuse. Detection and mitigation systems must reduce malicious activity without unnecessarily blocking legitimate queries or imposing unacceptable resolution delays. This paper develops PROJECT DIANA (Digital Infrastructure and Adversarial Network Abuse), a mathematical framework for analysing the relationship between attacker autonomy, DNS defensive effectiveness, and DNS-level preservation of legitimate activity. The framework represents the environment as a stochastic dynamical system, models attacker autonomy as a parameter of a state-dependent policy, defines security and openness measures, and formulates a constrained optimisation problem and security–openness Pareto set. It further specifies a feedback model for adaptive defence and a stochastic-game formulation for strategic interaction. A conditional monotonicity proposition is established for nested feasible sets. The manuscript identifies assumptions required to analyse equilibrium existence and local stability, while avoiding unsupported claims that greater autonomy must always reduce openness or that an equilibrium exists in every continuous-state game. DIANA is presented as a theoretical framework requiring further formal analysis and controlled empirical validation, not as an experimentally validated defence system.

Keywords: DNS Security, DNS Abuse, Attacker Autonomy, Adaptive Defence, Security–Openness Trade-Off, Constrained Optimisation, Pareto Frontier, Stochastic Games, Markov Perfect Equilibrium, Dynamical Stability

© 2026 The Authors. Published by Bexford Publishing Ltd

1. INTRODUCTION

1.1 Background and Motivation

The Domain Name System translates domain names into information used to locate Internet services. DNS is therefore a critical dependency for web access, email, software updates, cloud services, and other network activities. Its operational importance also makes it relevant to malicious activity involving phishing, malware, botnets, pharming, and spam used to deliver such harms. ICANN's DNS Abuse Mitigation Program identifies these categories as central forms of DNS abuse within its remit [1]. Research on DNS security covers malicious-domain detection, domain-generation algorithm (DGA) classification, DNS privacy, abuse mitigation, and infrastructure level threats. Surveys document a broad landscape of DNS threats and defensive techniques [2], [3]. A defender that increases filtering intensity may detect more malicious activity, but may also block legitimate requests, add processing delay, or disrupt service. These effects are not inevitable in every system: they depend on the detector, traffic distribution, policy, and environment. The central problem is therefore not to assume a universal trade-off, but to formulate a model in which the trade-off can be measured and its conditions investigated. The motivation for DIANA (Figure 1) arose from the NCUC quarterly session "AI and DNS Abuse: Keeping the Internet Safe and Open" in September 2026. The discussion raised questions about AI-assisted abuse, defensive false positives, and the preservation of legitimate Internet activity. The session motivates the research; it is not evidence that the proposed hypotheses are true or that the literature lacks relevant work.



Figure 1. PROJECT DIANA (Digital Infrastructure and Adversarial Network Abuse) project insignia.

1.2 Aim and objectives

The aim is to formulate a mathematical model of the interaction between attacker autonomy and defensive decisionmaking in DNS environments, with emphasis on security effectiveness, legitimate-service preservation, optimisation, and strategic stability.

- Represent DNS telemetry and attacker–defender interaction as a stochastic dynamical system.
- Define an operational autonomy parameter and a capability measure without assuming that autonomy always improves attack effectiveness.
- Specify security, DNS-level openness, and defensive cost measures
- Formulate constrained openness optimisation and a security–openness Pareto set.
- Identify conditions for monotonicity and sensitivity results.
- Specify an adaptive-defence feedback model and local-stability conditions.
- Formulate a stochastic game and distinguish established results from conditions still requiring

© 2026 The Authors. Published by Bexford Publishing Ltd

proof.
- Propose a reproducible, isolated testbed for future validation.

1.3 Research Questions

1. **RQ1** — Attacker autonomy: How can attacker autonomy be operationally represented, and under what conditions does it change attainable DNS-abuse capability?
2. **RQ2** — Security and openness: Under what conditions does increasing attacker autonomy reduce the maximum preservation of legitimate DNS activity when a minimum security level must be maintained?
3. **RQ3** — Adaptive defence: Under what conditions can adaptive defensive policies converge to a stable operating point while balancing security, openness, and operational cost?
4. **RQ4** — Strategic interaction: Under what assumptions does a dynamic attacker–defender game admit an equilibrium, and what can that equilibrium imply about defensive strategy?

1.4 Contributions & Status

The proposed contributions are: (i) a state-space formulation for DNS security under adaptive adversarial behaviour; (ii) a three-component DNS-level openness vector; (iii) a constrained optimisation formulation and Pareto outcome set; (iv) a conditional monotonicity proposition; (v) a feedback-system formulation for adaptive defence; and (vi) a stochastic-game model that makes equilibrium assumptions explicit. These are proposed analytical contributions. This manuscript does not report experiments, measured effect sizes, or a newly proved general Markov perfect equilibrium theorem.

2. RELATED WORK

2.1 Background & Motivation

DNS security research covers threats to resolvers, malicious domain use, DGA detection, DNS manipulation, denial-of-service, and privacy. Khormali et al. survey DNS security and privacy research across threat classes and methods [2]. Zhauniarovich et al. systematise malicious-domain detection through DNS data analysis and highlight the diversity of data sources, methods, and evaluation approaches [3]. Machine-learning models for DGA detection are one part of this broader field [4]. Such work supplies important detection methods, while DIANA focuses on how defensive decisions affect security metrics and legitimate-service performance under a changing adversarial policy.

2.2 Adaptive & Adversarial Learning

Adversarial learning studies interactions between learning systems and strategic attempts to manipulate their inputs or outcomes. For DNS, an abstract adaptive agent might change observable query features or domain-generation behaviour in response to defensive feedback. DIANA does not assume that an attacker necessarily uses reinforcement learning. It treats autonomy as a parameter that must be operationally mapped to adaptation capabilities. A future implementation could compare static, heuristic-adaptive, and learning-based agents in a controlled environment.

2.3 Game Theory & Moving-Target Defence

Game theory provides tools for representing strategic interaction between attackers and defenders. Moving-target defence research has examined the selection and timing of defensive changes under adversarial interaction; Tan et al. survey the intersection of moving-target defence and game theory [5]. Thus, DIANA cannot claim novelty merely for applying game theory to cybersecurity. Its potential distinctiveness must be established through a precise combination of autonomy modelling, DNS-level legitimate-service metrics, constrained openness optimisation, and derived analytical properties.

© 2026 The Authors. Published by Bexford Publishing Ltd

2.4 Stochastic-Game Equilibrium Theory

Existence of stationary Markov perfect equilibria is a non-trivial question for infinite-horizon stochastic games. He and Sun establish existence under a specific condition on transition kernels and discuss the limitations of general existence claims [6]. Research on the complexity of computing Markov perfect equilibria also demonstrates that equilibrium analysis in general-sum stochastic games requires care [7]. DIANA therefore does not claim that compactness, continuity, concavity, and discounting alone automatically establish an MPE for every continuous state formulation.

2.5 Research Positioning

The cited literature supports the relevance of DNS security, adaptive detection, game-theoretic defence, and stochastic equilibrium analysis. It does not yet establish that the complete DIANA formulation is novel. A formal literature review must compare the exact definitions, constraints, and theorems of closely related work before novelty can be claimed.

3. PROBLEM FORMULATION & SYSTEM MODEL

3.1 Time, State & Telemetry

Let time be indexed by $t \in N_0$. The DNS environment is represented by $x_t \in X \subseteq R^n$. The baseline model assumes χ is compact; convexity is imposed where required by a specific result. The state is partitioned into four conceptual telemetry groups:

- **Query telemetry $x_t(Q)$:** query volume, query-type distribution, and query-stream features.
- **Registration and domain telemetry $x_t(R)$:** domain-age summaries, TLD distribution, domain-name features, and observed address-mapping churn.
- **Infrastructure telemetry $x_t(T)$:** resolver cache statistics, response latency, resource utilisation, and packet-drop measures.
- **Legitimate-service telemetry x_t(O) :** successful legitimate resolutions, active service counts, and service availability indicators.

The groups are conceptual rather than a requirement to collect personally identifying registration data. A real testbed should minimise data collection, use synthetic or authorised data, and document provenance and privacy safeguards.

3.2 Action Spaces

At epoch t, the defender chooses $d_t \in D$ and the attacker chooses $a_t \in A$. A defender action may include detection threshold Θ_t , mitigation intensity Y_t , and resource allocation K_t . An attacker action is an abstract model input □ representing abuse load, infrastructure-change intensity, and evasion capability; it is not an operational attack recipe.

$$d_t = (\Theta_t, \gamma_t, \kappa_t) \in D; \quad a_t = \left(q_t, \mu_t, \rho_t\right) \in A$$

The action sets are assumed non-empty and compact. Convexity is imposed only where needed. If action variables are discrete, a continuous convex representation may be inappropriate and must be revised.

3.3 State Transition

Here ξ_t represents environmental uncertainty, such as benign traffic variation or transient service demand. The induced transition kernel is P(dx′ | x,a,d). The baseline model is time-homogeneous; explicitly

© 2026 The Authors. Published by Bexford Publishing Ltd

non-stationary traffic can be represented by augmenting the state with a regime variable. Continuity of f alone is not sufficient for every equilibrium result, so the relevant transition-kernel assumptions must be stated directly.

3.4 Attacker Autonomy

$$\alpha \in [0, 1], \quad a_t \sim \pi_A(\cdot \mid x_t; \alpha)$$

The parameter α indexes a family of attacker policies. At $\alpha = 0$, the baseline policy is specified as non-adaptive to defensive feedback. Intermediate values may represent increasing responsiveness to observed outcomes. Values near one represent the most adaptive policy class considered. These are modelling conventions, not universal measurements of autonomy. The application must define how α changes policy behaviour—for example, through an explicit adaptation rate, information set, or policy-update capability. α alone does not imply higher attack success, lower TPR, higher FPR, or faster convergence.

3.5 Abuse-Capabilities Index

$$B(x, a, d), = \Sigma_{i=1}^{4} w_i X_i(x, a, d), \quad w_i \geq 0, \quad \Sigma_i w_i = 1$$

The candidate components X_i represent malicious-query effectiveness, communication persistence, infrastructure adaptability, and resilience to defensive intervention. The weights w_i determine their relative contribution. Each component must be operationally defined and scaled before B can be interpreted quantitatively. Unless the weights are validated, B should be reported as a proposed index alongside its component values. Autonomy affects B through the action-policy distribution unless a direct dependence is explicitly introduced.

4. SECURITY, OPENNESS & OPTIMISATION

4.1 Security Effectiveness

$$S(d, \alpha; x) = TPR(d, \alpha; x) - \lambda FDR(d, \alpha; x), \quad \lambda \geq 0$$

TPR denotes the true-positive rate for malicious activity and FPR the false-positive rate for legitimate activity. λ encodes the relative operational cost of false alarms. Because this scalar can mask different detector behaviours, TPR and FPR should be reported separately. Ground-truth labels and denominators must be defined independently of the detector.

4.2 DNS-Level Openness Vector

$$\vec{O}(d, \alpha; x), = [A(d, \alpha; x), 1 - L(d, \alpha; x), 1 - D(d, \alpha; x)]^T$$

$$A = 1 - E\frac{E[N \text{ legitimate blocked}]}{E[N \text{ legitimate total}]}$$

$$L = min\left(1, \frac{E[\Delta T \text{ resolution}]}{T \text{ max}}\right)$$

$$D = E\frac{[N \text{ legitimate services disrupted}]}{N \text{ total active services}}$$

The vector represents legitimate DNS preservation through availability, normalised added latency, and service disruption. It is a DNS-level proxy, not a complete measure of Internet openness. Definitions

© 2026 The Authors. Published by Bexford Publishing Ltd

require non-zero denominators, a stated latency baseline, and a consistent definition of disrupted service.

The latency cap means differences above $T\,max$ are not distinguished by the aggregate score, so raw latency distributions should also be reported.

4.3 Aggregate Openness

$$O(d,\ \alpha; x) = \omega_1 A\ +\ \omega_2(1\ -\ D)$$
$$\omega_i \geq 0, \quad \Sigma_{i=1}^{\ 3}\ \omega_i = 1$$

The aggregate score supports scalar optimisation. Its weights should be justified by stakeholder requirements or sensitivity analysis. The three components must be reported alongside the aggregate to prevent metric masking.

4.4 Defense Cost

$$C\,D(d)\ = c_1||\gamma||_2^{\ 2}\ +\ c_2\,||k||_1, \quad c_1, c_2 >\ 0$$

The cost function approximates mitigation processing overhead and resource allocation. Coefficients must be calibrated or treated as symbolic parameters. Since O is dimensionless while $C\,D$ may use operational units, subtracting cost directly from O is meaningful only after normalisation or introduction of an explicit conversion weight. Let $\overline{C}\,D$ denote normalised cost.

4.5 Constrained Optimisation

$$D\,feas(\alpha; x)\ = \{d\ \in\ D\colon S(d,\ a; x)\ \geq\ S\,min\}$$
$$V\ *\ (\alpha;\ x)\ =\ max\,\{d\ =\ D\,feas(\alpha; x)\}\,[O(d,\ \alpha; x)\ -\ \beta\,C\ \ \overline{C}\,D(d)]$$

Here $\beta\,C\ \geq\ 0$ weights normalised defensive cost. The feasible set must be non-empty. If the set is compact and the objective continuous, a maximum is attained. If no defence meets $S\,min$ at a given autonomy level, the model should report infeasibility rather than an ordinary optimum.

4.6 Security Openness Pareto Set

$$F(\alpha; x) = \{\ (S(d,\ \alpha; x),\ O(d,\ \alpha; x))\ :\ d\ \in\ D\}$$

A strategy d_1 Pareto-dominates d_2 if it is no worse on both security and openness and strictly better on at least one. The Pareto frontier consists of non-dominated outcomes. The frontier describes available choices without presuming a universal trade-off. Defensive cost can be a third objective or handled separately through the constrained objective; the chosen interpretation must be stated consistently.

5. ANALYTICAL PROPOSITIONS & SENSITIVITY

5.1 Monotonicity Under Nested Feasible Sets

Proposition 1 (Conditional monotonicity). Let $\alpha_2 > \alpha_1$. Suppose

i. $D\,feas(\alpha_2; x)\ \subseteq\ D\,feas(\alpha_1; x)$
ii. the objective $h(d) = O(d; x)\ -\ \beta\,C\ \ \overline{C}\,D(d)$ is independent of α for each fixed d, and
iii. both feasible sets are non-empty and the maxima are attained. Then $V\ *\ (\alpha_2; x) \leq\ V\ *\ (\alpha_1; x)$.

Proof. Every strategy feasible at α_2 is also feasible at α_1. Maximising the same function h over the larger feasible set cannot produce a lower value than maximising it over the subset. Therefore,

© 2026 The Authors. Published by Bexford Publishing Ltd

$V^*(\alpha_2; x) \leq V^*(\alpha_1; x)$ ∎

This proposition does not prove that $D\,feas$ contracts as α increases. That inclusion is a separate modelling hypothesis requiring justification. If openness depends directly on α for a fixed defence, nested feasible sets alone are insufficient to compare the optimal values; additional assumptions on objective dependence are required.

5.2 Implications of Monotone Detection Rates

If, for a fixed defence, $TPR(d, \alpha_2; x) \geq TPR(d, \alpha_1; x)$ and $FPR(d, \alpha_2; x) \geq FPR(d, \alpha_1; x)$, with $\lambda \geq 0$ then $S(d, \alpha_2; x) \geq S(d, \alpha_1; x)$. This may imply nested security-feasible sets under consistent definitions. But these inequalities are substantive assumptions, not universal laws of autonomous attackers. They should be treated as a scenario to investigate or validate, not a universal fact.

5.3 Sensitivity of the Optimal Value

$$\mathcal{L}(d, \mu; \alpha) = O(d, \alpha; x) - \beta C \bar{C} D(d) + \mu[S(d, \alpha; x) - S\,min]$$

$$\frac{\partial V^*}{\partial \alpha} = \frac{\partial O}{\partial \alpha} + \mu^* \frac{\partial S}{\partial \alpha}$$

Under suitable differentiability, constraint qualification, and regularity conditions, an envelope result may be written using the Lagrangian. Derivatives are evaluated at an appropriate optimal primal–dual solution. The equality may fail at non-unique optima, changing active constraints, or non-differentiable value functions. Moreover, the sign cannot be inferred from monotonicity of TPR and FPR alone unless the signs of $\frac{\partial S}{\partial \alpha}$ and $\frac{\partial O}{\partial \alpha}$ are separately established. The non-positive derivative is therefore a conditional result, not a theorem established by the current definitions.

6. ADAPTIVE DEFENCE & DYNAMIC STABILITY

6.1 Feedback Update

$$\Theta_{t+1} = g(\Theta_t) = \theta_t + \eta[e(\Theta_t; \alpha) - e^*]$$

Here θ is a scalar threshold, $\eta > 0$ the adaptation rate, $e(\Theta; \alpha)$ a defined operational error measure, and e^* the target. A candidate is $e = FPR + \beta(1 - TPR)$, where $\beta \geq 0$ weights missed malicious activity. The update direction must match the meaning of Θ and the detector's response to threshold changes.

6.2 Fixed Point & Local Stability

$$g'(\Theta^*) = 1 + \eta \frac{\partial e(\Theta^*; \alpha)}{\partial \Theta}$$

A fixed point satisfies $g(\Theta^*) = \Theta^*$. For a differentiable scalar update, $|g'(\Theta^*)| < 1$ is a standard sufficient condition for local asymptotic stability. If, $e'(\Theta^*) < \Theta$ this gives the local condition:

$$0 < \eta < \frac{-2}{e'(\Theta^*)}$$

This bound applies to the stated scalar model and a fixed point with negative derivative. It is not a universal bound based on the largest absolute eigenvalue for arbitrary systems. For a vector update $\Theta_{t+1} = g(\Theta_t)$, the Jacobian $J\,g(\Theta^*)$ governs the linearisation; $\rho(J\,g(\Theta^*))$ is sufficient for local asymptotic stability under standard differentiability assumptions. Eigenvalues on the unit circle make linearisation inconclusive. The condition is local and does not establish global stability or imply that

© 2026 The Authors. Published by Bexford Publishing Ltd

instability necessarily produces chaos.

6.3 Non-Stationary Traffic

If traffic regimes change over time, the update is time-varying or stochastic. Stability then requires further analysis, such as bounded disturbances, stochastic stability, or robust stability over a specified family of regimes. **DIANA** *does not claim global asymptotic stability under arbitrary non-stationary traffic.*

7. DYNAMIC ATTACKER-DEFENCE GAME

7.1 Game Definition

$$\Gamma = \{X,\ A,\ D,\ P,\ U\,A,\ U\,D,\ \gamma\}, \quad \gamma \in (0,1)$$
$$U\,A\,(x,a,d;\ \alpha) = V\,A(B(x,a,d)) - C\,A\,(a) - R\,A\,(d,a)$$
$$U\,D(x,a,d;\ \alpha) = \lambda\,S\ S(d,a;x) + \lambda\,O\ O(d,\ a;x) - \beta\,C\ \overline{C}\,D(d)$$

The game has state space X, action spaces A and D, transition kernel P, stage utilities $U\,A$ and $U\,D$, and discount factor γ. V_A maps effective capability to attacker value; $C\,A$ represents attack cost; $R\,A$represents defensive risk. $\lambda\,S$ and $\lambda\,O$ are positive weights. If the defender must meet a security constraint at every state, the feasible action correspondence must be defined consistently and shown non-empty. A robust constraint takes an infimum over admissible attacker actions; an expected constraint depends on the attacker's policy. These are different models.

7.2 Stationary Markov Strategies & Values

Let $\sigma\,A(x)$ and $\sigma\,D(x)$ be stationary Markov strategies, mapping the current state to action distributions. Under a strategy profile, player *i*'s discounted payoff is:

$$Vi^{\{\sigma\,A,\,\sigma\,D\}}(x) = E\left[\Sigma_{t=0}^{\ \infty}\,\gamma^{t}\ U_i\left(x_t, a_t, d_t\right) \mid x_0 = x\right]$$

An MPE is a profile that is a best response for each player in every state-contingent subgame. In a general-sum game, the objectives are coupled through the transition law and the opponent's policy. The two Bellman expressions therefore cannot be treated as independent single-agent maximisation operators and assumed to define a contraction.

7.3 Equilibrium Existence: Scope of the Claim

A general stationary MPE existence theorem for the proposed continuous-state, continuous-action game is not established in this manuscript. Compactness, continuity, own-action concavity, and discounting do not by themselves complete the proof for every stochastic game. Existence depends on the transition kernel, feasible-action correspondence, payoff structure, and equilibrium concept. He and Sun [6] demonstrate the importance of additional transition-kernel conditions in general-state discounted stochastic games.

DIANA therefore treats equilibrium existence as a formal research objective: identify sufficient conditions for the selected game and prove that they apply. A tractable first step is a finite-state, finite-action discretisation with explicit transitions and payoffs, where established stochastic-game *existence results may apply.* Extending a result to continuous spaces requires a separate theorem *whose assumptions match the model.*

7.3 Equilibrium Implications

If an equilibrium exists, its policy profile describes mutually optimal stationary responses under the model's utilities and dynamics. It does not automatically provide a universal upper bound on abuse capability, guarantee operational safety, or show that high autonomy reduces openness. Those implications

© 2026 The Authors. Published by Bexford Publishing Ltd

must be derived from equilibrium conditions and additional comparative-statics assumptions.

8. DISCUSSION

8.1 Interpretation

DIANA's central premise is that DNS defensive performance should be assessed jointly across security and legitimate-service preservation. The framework does not assert that the trade-off is inevitable. It defines quantities that can reveal whether the attainable security–openness set changes under different attacker-policy regimes and defensive constraints. This separates three questions: what is measurable, what can be optimised, and what can be proved from explicit assumptions.

8.2 Operational Relevance

DNS operators should report false-positive rates, legitimate-query availability, resolution latency, service disruption, and resource cost alongside detection rates. Threshold adaptation should be tuned against a model of its feedback dynamics rather than a universal learning-rate formula. Game-theoretic analysis may help compare policies, but deployment decisions require empirical calibration, robustness analysis, and safeguards for legitimate users.

9. LIMITATIONS & FUTURE WORK

- The autonomy parameter α requires operational definition and calibration.
- The capability index and aggregate openness score depend on component definitions, scales, and weights.
- Monotonicity of detection rates, openness, or feasible sets is not universal.
- The continuous-state MPE existence theorem remains to be established for the selected model.
- Local stability does not imply global stability; non-stationary traffic requires additional analysis.
- Novelty must be evaluated through a deeper, reproducible literature review.

9.1 Future Analytical Work

Future analysis should specify an autonomy mechanism; formalise the dependence of transitions and metrics on α; establish feasibility and value-function regularity; derive a sensitivity result with complete constraint qualifications; prove stability for a specified update law; and choose a game class with a valid equilibrium-existence theorem. Counterexamples should be considered to identify which claims fail when assumptions are relaxed.

9.2 Future Controlled Validation

A later study may implement an isolated DNS testbed with synthetic legitimate and malicious-labelled traffic. Conditions may vary: autonomy regime, defensive strategy, and security threshold. Ground truth must be independent of the detector. The study should measure actual resolution latency and resource use, retain event-level logs, use repeated independent trials, report uncertainty intervals, and publish configuration and analysis code. Data must be used under its licence and access terms.

10. CONCLUSION

This paper presents **PROJECT DIANA**, a mathematical framework for analysing attacker autonomy, DNS security, and DNS-level preservation of legitimate activity. It defines a stochastic system model, candidate security and openness metrics, constrained optimisation, a Pareto outcome set, a feedback model for adaptive defence, and a dynamic attacker–defender game. A conditional proposition establishes non-increasing optimal value when feasible sets are nested and the objective is unchanged across autonomy levels. The framework identifies further conditions required for sensitivity, stability, and

© 2026 The Authors. Published by Bexford Publishing Ltd

equilibrium results. **DIANA's** *contribution will depend on deriving precise, non-trivial results that withstand comparison with existing theory and controlled empirical evidence.*

References

[1] Internet Corporation for Assigned Names and Numbers (ICANN), "DNS Abuse Mitigation Program," updated July 2026. https://www.icann.org/dnsabuse

[2] A. Khormali, J. Park, H. Alasmary, A. Anwar, M. Saad, and D. Mohaisen, "Domain name system security and privacy: A contemporary survey," Computer Networks, vol. 185, art. 107699, 2021. doi: 10.1016/j.comnet.2020.107699.

[3] Y. Zhauniarovich, I. Khalil, T. Yu, and M. Dacier, "A survey on malicious domains detection through DNS data analysis," ACM Computing Surveys, vol. 51, no. 4, 2018. doi: 10.1145/3191329.

[4] J. Mao, J. Zhang, Z. Tang, and Z. Gu, "DNS anti-attack machine learning model for DGA domain name detection," Physical Communication, vol. 40, art. 101069, 2020. doi: 10.1016/j.phycom.2020.101069.

[5] J. Tan, H. Jin, H. Zhang, Y. Zhang, D. Chang, X. Liu, and H. Zhang, "A survey: When moving target defense meets game theory," Computer Science Review, vol. 50, art. 100544, 2023. doi: 10.1016/j.cosrev.2023.100544.

[6] W. He and Y. Sun, "Stationary Markov perfect equilibria in discounted stochastic games," Journal of Economic Theory, vol. 169, pp. 35–61, 2017. doi: 10.1016/j.jet.2017.01.007.

[7] "On the complexity of computing Markov perfect equilibrium in general-sum stochastic games," National Science Review, vol. 10, no. 1, art. nwac256, 2023. doi: 10.1093/nsr/nwac256.

© 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the **Creative Commons Attribution (CC BY 4.0) licence** (https://creativecommons.org/licenses/by/4.0/)