Calibrating Evolutionary Game Payoffs from Malware Execution Traces

A Signature-Conditional Equilibrium Study on WinMET 2025

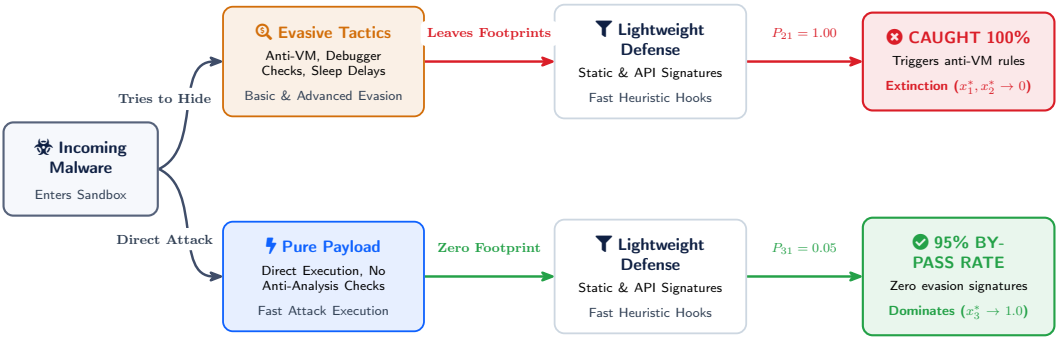
Tran Duc Le[1] Truong Duy Dinh[2] Igor Ushakov[3] Dinh Hai Nguyen[2]

[1]University of Wisconsin-Stout Polytechnic, USA [2]Posts and Telecommunications Institute of Technology, Vietnam
[3]The Bonch-Bruevich Saint Petersburg State University of Telecommunications, Russia

September 2026

Spacsec 2026 Conference — Research Presentation

The Evasion Paradox: Why "Clever" Malware Fails



The Core Paradox: Trying to evade dynamic analysis acts as a beacon that triggers heuristic signatures immediately (100% caught). Direct payload execution slips through lightweight filters 95% of the time, driving malware evolution toward pure payload!

Research Gap: The Hazard of Assumed Payoffs

Prior Game Theory (Manshaei et al., 2013)

- Traditional security game theory relies on **hypothesized payoffs**.
- Textbook assumption: Evasion bypasses lightweight checks ($P_{21} \approx 0.1$), while payload is easily caught ($P_{31} \approx 0.7$).
- **Textbook Advice:** "Invest everything in catching advanced evasion."

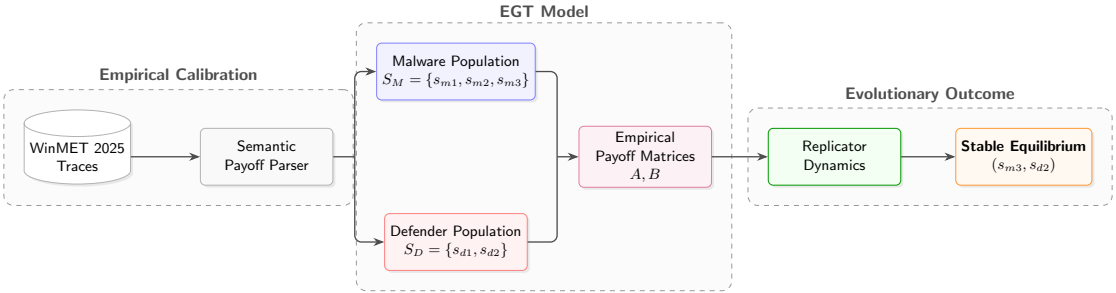
Empirical Grounding (Our Finding)

- Grounded in **198 real bare-metal traces** from WinMET 2025.
- Modern sandboxes catch evasion immediately via heuristic signatures ($P_{21} = 1.00$).
- **Operational Reality:** Pure payload slips through lightweight filters 95% of the time ($P_{31} = 0.05$)!

Our Objective

Calibrate an Evolutionary Game Theory (EGT) model against real execution data to uncover the **true empirical equilibrium** of malware evasion and sandbox defense.

Methodology: From Raw Traces to Evolutionary Equilibria



1. Ingestion

198 bare-metal JSON execution reports from WinMET 2025.

2. Classification

Signature semantics map tactics to s_{m1}, s_{m2}, or s_{m3}.

3. Payoffs

Extract C_i **&** P_{ij} from API overhead and proxy detections.

4. Equilibrium

Replicator dynamics and Jacobian invasion eigenvalues.

Game Formulation: What Drives Survival?

Malware Strategy Space (S_M)

- s_{m1} **Basic Evasion:** Sleep delays, debugger checks.
- s_{m2} **Advanced Evasion:** Anti-VM, process hollowing, timing checks.
- s_{m3} **Pure Payload:** Immediate attack, zero anti-analysis noise.

Defender Strategy Space (S_D)

- s_{d1} **Lightweight:** API hooks and static signatures ($K_1 = 1$).
- s_{d2} **Deep Bare-Metal:** Introspection and hardware traces ($K_2 = 5$).

Malware Fitness (a_{ij})

$$a_{ij} = V \cdot (1 - P_{ij}) - C_i$$

Payoff = (Reward $\times$ Escape Rate) $-$ Evasion Overhead

Trying to hide triggers alarms ($P_{ij} \to 1.0$), collapsing malware payoff!

Defender Fitness (b_{ji})

$$b_{ji} = R \cdot P_{ij} - L \cdot (1 - P_{ij}) - K_j$$

Payoff = Catch Reward $-$ Breach Loss $-$ Compute Cost

Balancing breach risk ($L = 10$) against deep analysis overhead ($K_2 = 5$).

Empirical Calibration: The Numbers That Flipped the Field

Ground Truth from 198 Real WinMET 2025 Traces (Zero Synthetic Data)

Calibrated across 32 malware families under the CAPE dynamic sandbox engine.

Malware Strategy (S_M)	**Traces (N)**	**API Overhead (C_i)**	P_{i1} **(Lightweight)**	P_{i2} **(Deep Bare-Metal)**
s_{m1} (Basic Evasion)	58	0.122	0.81 [0.71, 0.91]	1.00 [1.00, 1.00]
s_{m2} (Advanced Evasion)	63	0.004	**1.00 [1.00, 1.00]**	1.00 [1.00, 1.00]
s_{m3} (Pure Payload)	77	0.0002	**0.05 [0.01, 0.10]**	**0.53 [0.43, 0.65]**

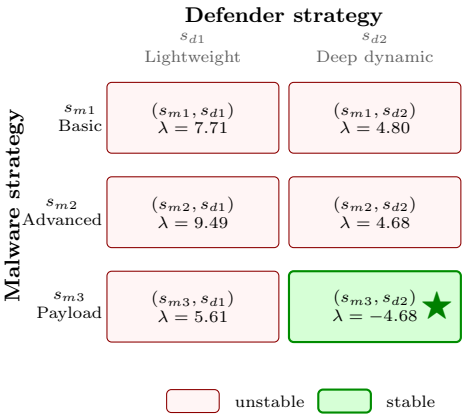
Evasion Is Caught (100%)

Anti-VM and timing checks leave behavioral footprints that trip lightweight heuristic hooks.

Pure Payload Slips Through (95%)

Omitting anti-analysis logic triggers zero evasion signatures, easily bypassing lightweight filters!

Jacobian Stability Analysis: Finding the Inescapable State



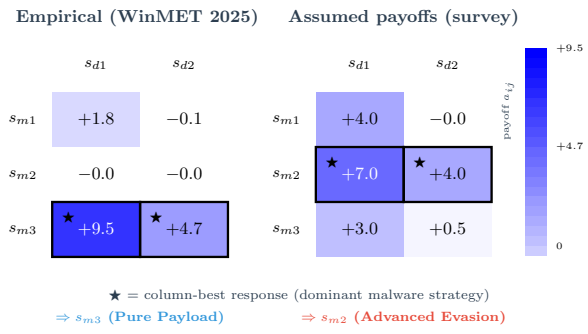
Five pure profiles are invadable ($\mathrm{Re}(\lambda) > 0$)*;* ★ *marks the unique ESS,* (s_{m3}, s_{d2}).

Hofbauer-Sigmund Stability Criterion

A profile is an **Evolutionarily Stable Strategy (ESS)** iff all invasion eigenvalues satisfy: $\max \mathrm{Re}(\lambda) < 0$.

- Five pure profiles have positive eigenvalues ($+4.68$ to $+9.49$): **unstable saddles**.
- **Profile** (s_{m3}, s_{d2}) **is the unique stable ESS** ($\max \lambda = -4.68$)**:**
  - $\lambda^{M}_{s_{m1}} = -4.80$: Basic evasion mutants repelled.
  - $\lambda^{M}_{s_{m2}} = -4.68$: Advanced evasion mutants repelled.
  - $\lambda^{D}_{s_{d1}} = -5.61$: Defender cannot downscale to light defense.
- **Strategic Reality:** Malware fixates on **Pure Payload**, forcing defenders to maintain **Deep Bare-Metal Inspection**.

Grounding Matters: Qualitative Equilibrium Inversion



Assumed Payoffs (Manshaei et al., 2013)

Assumes advanced evasion bypasses light analysis ($P_{21} = 0.1$).
→ **Predicts Advanced Evasion (s_{m2}) Dominance** ($x_2^* = 1.0$).

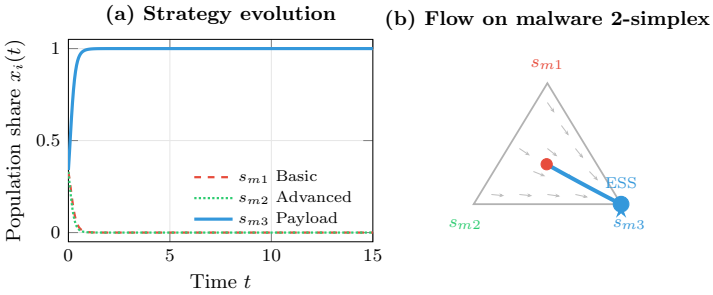
Empirical WinMET Calibration (Our Work)

Heuristic rules catch evasion ($P_{21} = 1.0$); payload escapes ($P_{31} = 0.05$).
→ **Predicts Pure Payload (s_{m3}) Dominance** ($x_3^* = 1.0$).

Strategic Implication

Payoff grounding does not just adjust numbers: it **qualitatively inverts** the predicted dominant strategy!

Evolutionary Dynamics: All Roads Lead to Pure Payload



Evolutionary Convergence

- **Time Plot (a):** Both basic and advanced evasion plummet to extinction.
- **Simplex Field (b):** Vector field funnels every initial malware mix straight to the Pure Payload vertex (s_{m3}).
- Defender population simultaneously shifts entirely to deep bare-metal analysis ($y_2 \to 1.0$).

Robustness Certificates: 3 Exhaustive Stress Tests

Sensitivity sweep certificate

Five one-at-a-time parameter sweeps; all outcomes select Pure Payload.

35/35
settings converge to s_{m3} dominance

$x_3^* = 1.00$

Parameter	Sweep range	Points	x_3^* range
V	2–30	7	1.00–1.00
R	2–30	7	1.00–1.00
L	2–30	7	1.00–1.00
K_1	0.5–5	7	1.00–1.00
K_2	2–15	7	1.00–1.00

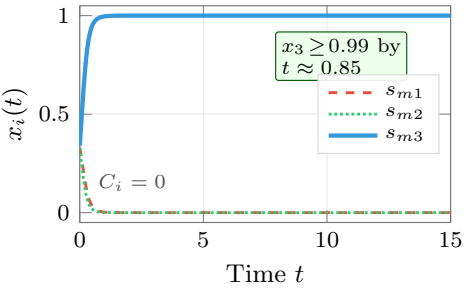
Because every sweep point gives the same equilibrium share, the useful result is invariance, not bar length: x_3^* remains exactly 1.00 throughout the audited grid.

Why This Result Is Invariant

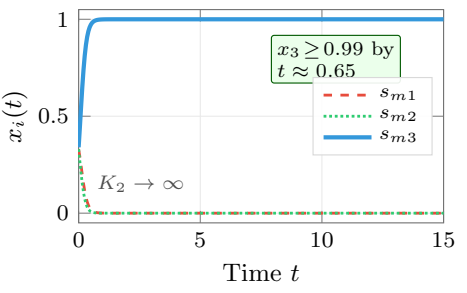
1. **35 Parameter Settings:** Sweeps across rewards, losses, and costs ($V, R, L \in [2, 30]$, $K_1 \in [0.5, 5]$, $K_2 \in [2, 15]$) yield $x_3^* = 1.00$ in **all 35 cases**.
2. **48 Random Basins:** 48 randomly generated Dirichlet starting states all converge to Pure Payload ($x_3 > 0.95$).
3. **Excluding Silent Traces:** Removing 35 empty traces still preserves (s_{m3}, s_{d2}) as unique ESS ($\max \lambda = -0.24$).

Ablation Studies: Testing Skeptical Counter-Hypotheses

Ablation 1: zero-cost evasion



Ablation 2: no deep inspection



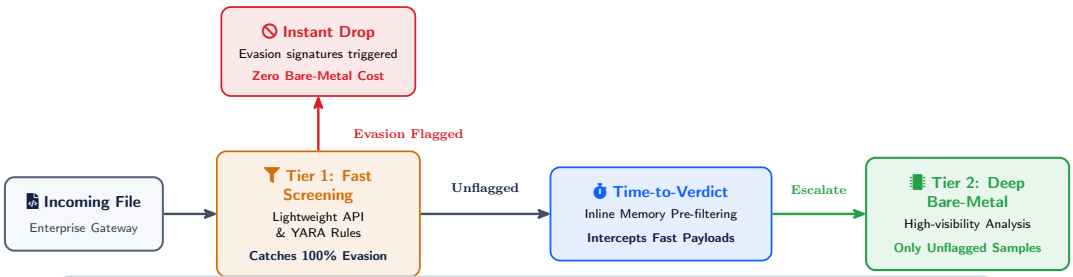
Ablation 1: What if Evasion Was Free? ($C_i = 0$)

Even if writing evasion code required zero effort, Pure Payload still wins by 7.61 points! Evasion fails due to high detection rates, not execution cost.

Ablation 2: What if Defenders Cannot Afford Deep Inspection? ($K_2 \to \infty$)

Restricting defenders to light inspection still drives malware to Pure Payload ($a_{31} = 9.48$ vs $a_{21} = -0.004$). Cheap filters let 95% of payloads through!

Strategic Implications for Sandbox Architecture



Pillar 3: Signature Hygiene as Strategic Anchor

Regularly updating anti-VM and anti-debug rules ensures evasion remains evolutionarily unviable ($P_{21} = 1.00$). Tiered escalation saves deep bare-metal compute while preserving equilibrium (s_{m3}, s_{d2}).

Scope Boundaries and Methodological Honesty

Signature-Conditional Scope

- The equilibrium reflects the sandbox-perceived landscape under CAPE's signature set.
- API hooking cannot see raw CPUID/RDTSC instructions unless flagged by signatures.
- Measured overhead C_i reflects execution-time API calls (lower bound for developer effort).

Corpus Context and Extensions

- **Sample Size:** Calibration study on $N = 198$ traces covering 32 families from WinMET 2025 (31,844 traces in full corpus).
- **Face Validity Confirmed:** Documented evasion-heavy families (Agent Tesla, Amadey, LokiBot, Redline) correctly categorized as s_{m2}.
- **Future Work:** Scaling to full 750 GB corpus and co-evolutionary dynamic modeling.

Conclusion: Grounding Game Theory for Real-World Defense

4 Key Takeaways to Remember

1. **Empirical Grounding Overturns Dogma:** EGT models must calibrate from real execution traces. Assumed payoffs lead to inverted strategic conclusions.
2. **The Evasion Trap:** Clever evasion triggers heuristic alarms (100% caught), while brute-force payloads slip through lightweight filters (95% bypass).
3. **Inescapable Equilibrium:** (s_{m3}, s_{d2}) (Pure Payload vs Deep Bare-Metal) is the unique ESS across 35 parameter sweeps and 48 random basins.
4. **Architectural Guidance:** Focus on **Time-to-Verdict** for fast payloads, maintain **Signature Hygiene**, and deploy **Tiered Escalation**.

Thank You! Questions & Discussion

Contact: let@uwstout.edu — University of Wisconsin-Stout Polytechnic

Backup: WinMET 2025 Full Corpus Family Distribution

Family Name	Trace Count	Percentage
Redline	1,811	5.7%
Agent Tesla	1,299	4.1%
Smokeloader	983	3.1%
LokiBot	695	2.2%
Amadey	436	1.4%
Other Identified Families (1,039 families)	19,364	60.8%
Unidentified / Consensus Pending	7,256	22.8%
Total Full Corpus	**31,844**	**100.0%**

Calibration Subset Representation

Calibration subset covers 32 of 1,044 families ($N = 198$ traces), correctly identifying documented evasion-heavy families (Agent Tesla, Amadey, LokiBot, Redline).

Backup: Proxy Detection Model Comparison

Severity Proxy (Primary)

- s_{d1} (Lightweight): Detects if any signature has severity ≤ 2.
- s_{d2} (Deep): Detects if any signature fires (full visibility).
- $P_{11} = 0.81$, $P_{21} = 1.00$, $P_{31} = 0.05$.
- $P_{12} = 1.00$, $P_{22} = 1.00$, $P_{32} = 0.53$.

Semantic Split (Cross-Check)

- s_{d1} (Lightweight): Detects only via static-capable signatures (YARA, hash, imphash, procmem).
- s_{d2} (Deep): Retains full visibility.
- Yields identical qualitative ordering and identical unique ESS: (s_{m3}, s_{d2}).