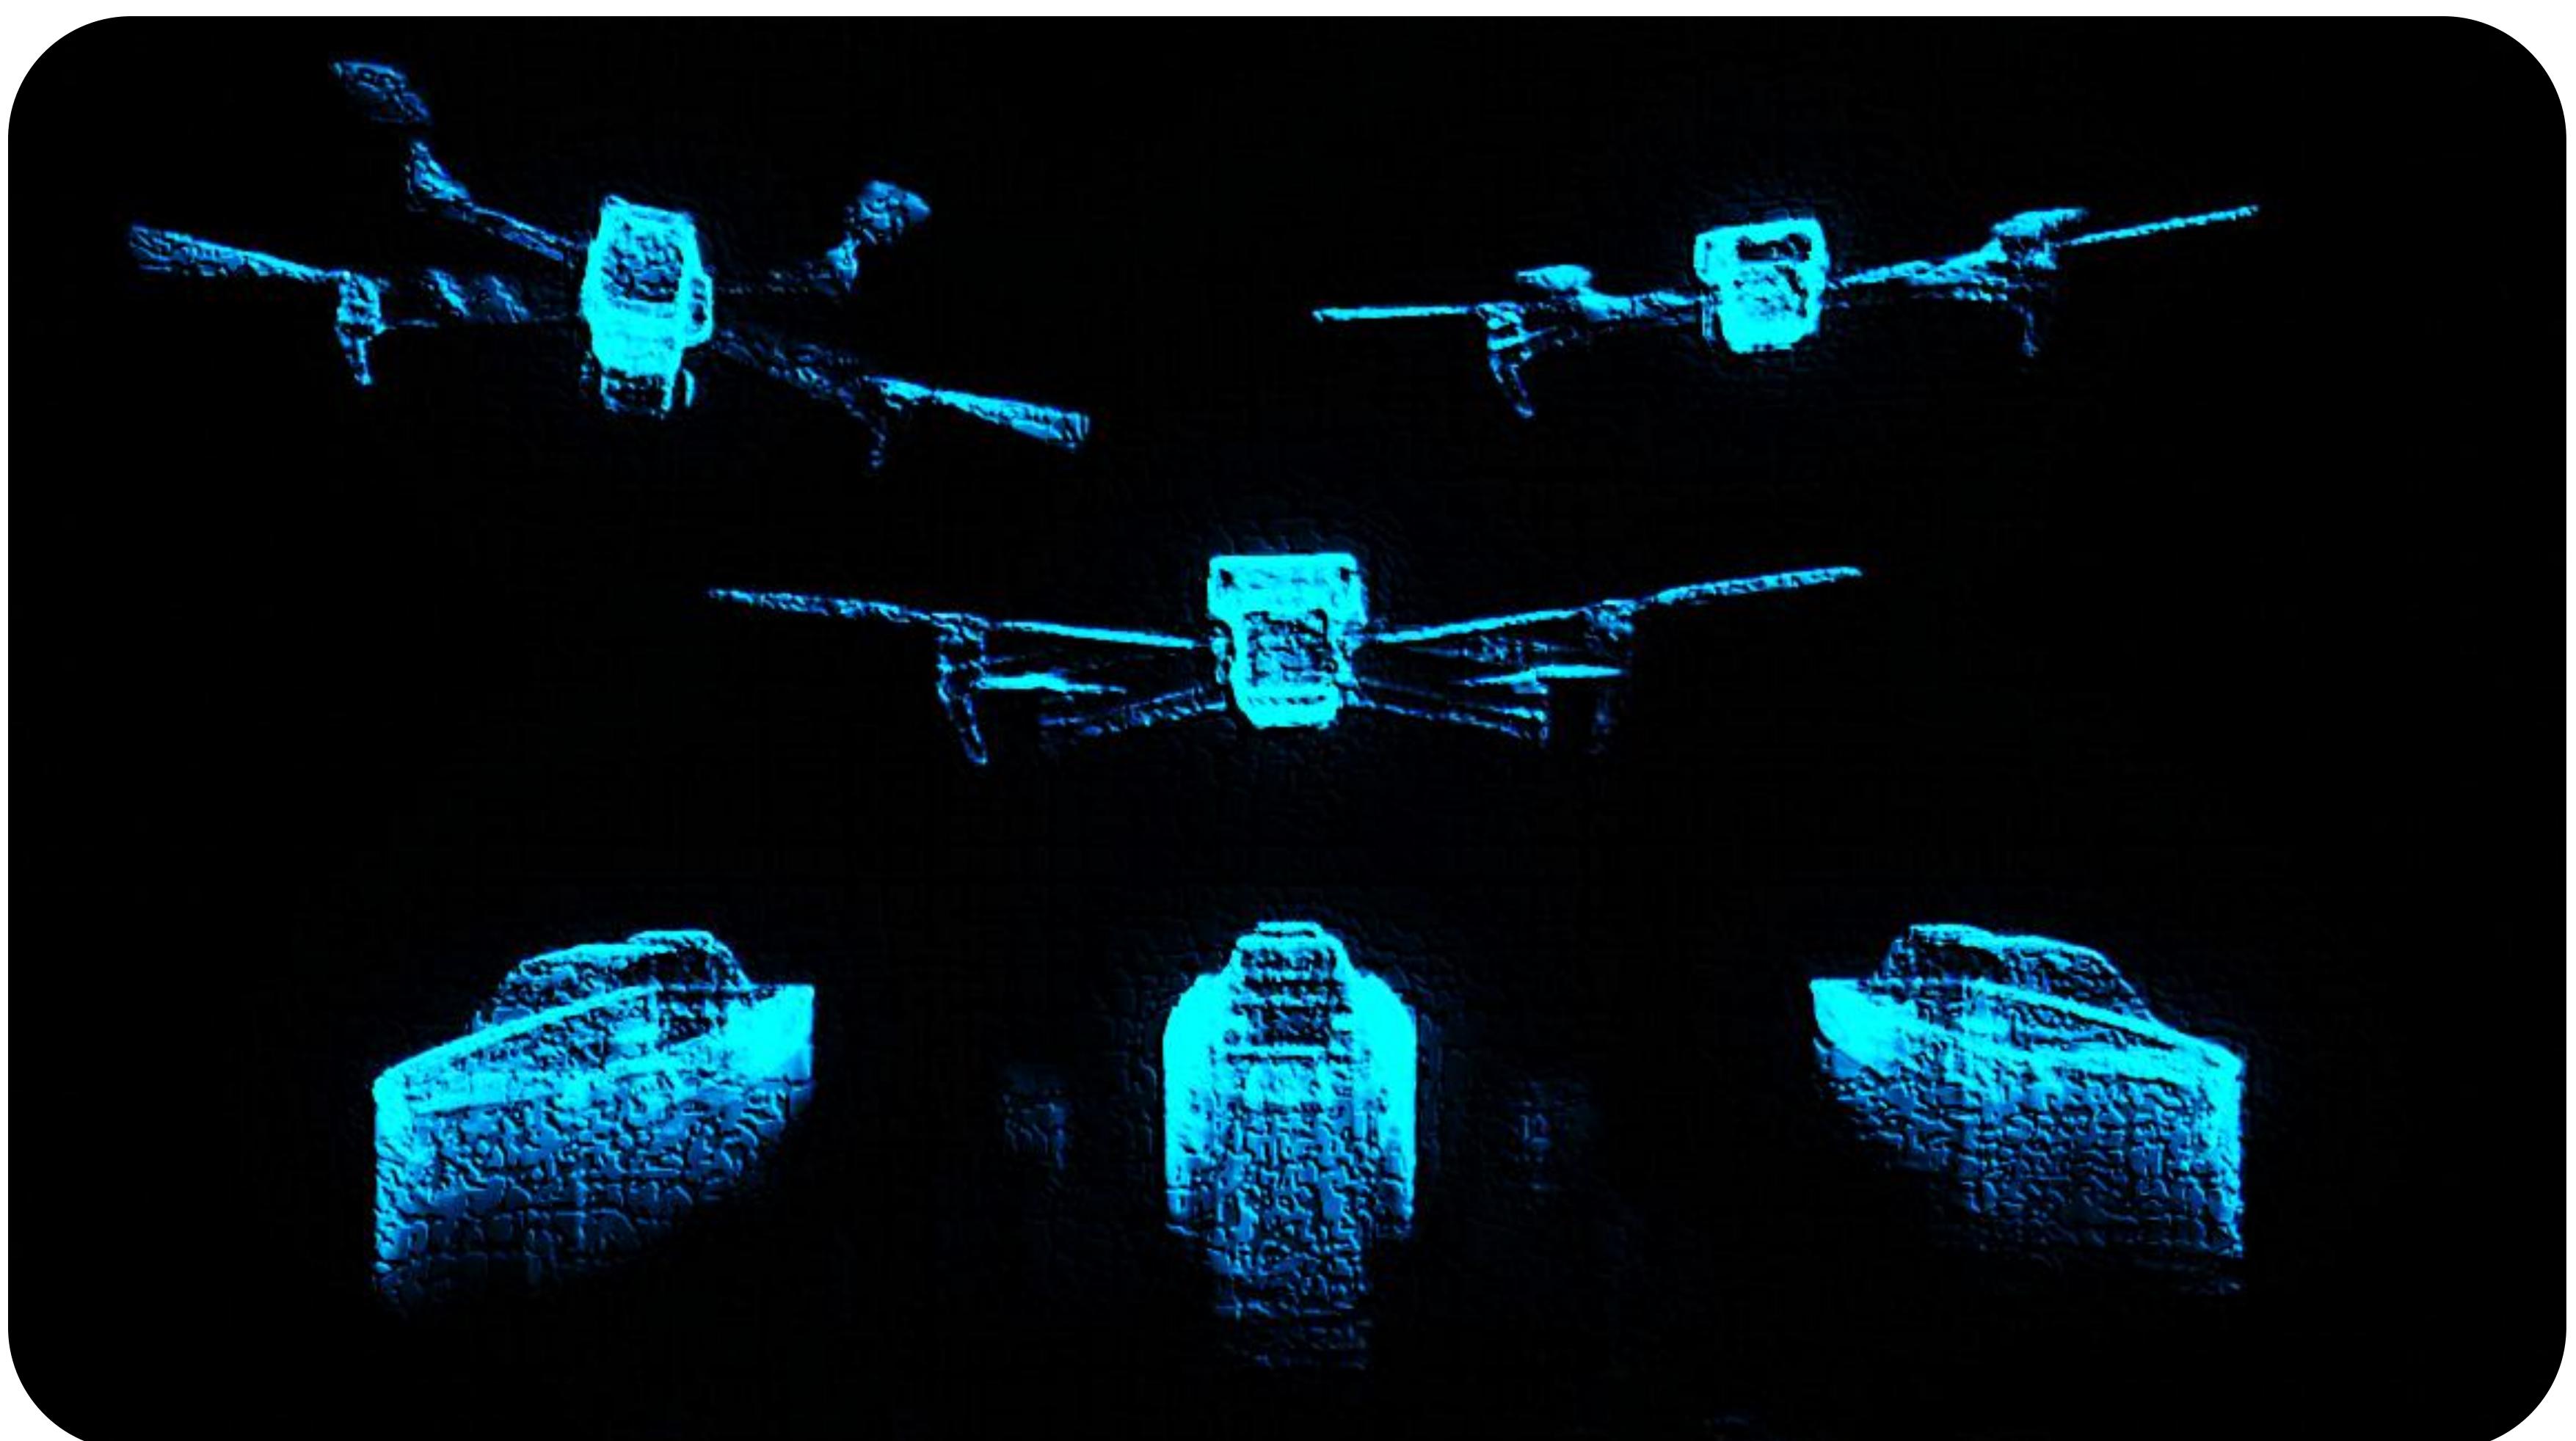


From Detection to Depletion: Cost-Exchange Limits in the Russia–Ukraine Drone War

Dr. Shaoyuan Wu

The Global AI Governance Research Center, EPINOVA



Research Report

Contents

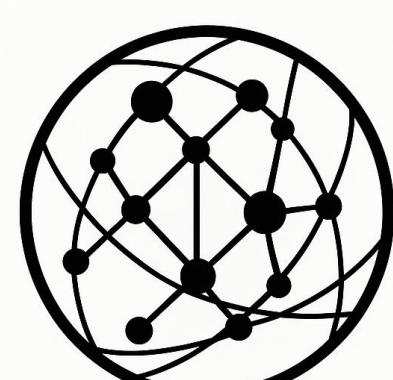
Series Information:
This research report is part of the EPINOVA Research Report Series on AI-Enabled Warfare, Sustainability, and Global Security Governance.

Recommended Citation:
Wu, S. (2025). *From Detection to Depletion: Cost-Exchange Limits in the Russia–Ukraine Drone War* (Research Report). EPINOVA–2025–01–RR. Global AI Governance Research Center, EPINOVA LLC.

Short Citation:
Wu (2025), *From Detection to Depletion*, EPINOVA RR–2025–01.

Disclaimer:
This report reflects the author's independent analysis based on publicly available information and does not represent the official views or positions of any government, organization, or institution.

DOI:
<https://doi.org/10.5281/zenodo.18036790>



GLOBAL AI
GOVERNANCE
RESEARCH CENTER

Executive Summary	03
Background	05
Methodology and Auditability	09
1. Observed Operational Frictions in the Drone Fight	11
1.1 Sensor Physics at Low Altitude	11
1.2 Temporal Mismatch	11
1.3 Infrared and Electro-Optical Constraints	12
1.4 Reduced Observability Is Not Stealth	12
1.5 Diminishing Communications Visibility	13
1.6 Spectrum Congestion and Frequency Agility	13
1.7 Cost–Magazine Frictions	13
1.8 Structural Conclusion: System-on-System Friction	14
2. Analytic Framework: From Detection to Depletion	15
2.1 Unit of Analysis and Minimum Viable, Auditable Standard	15
2.2 Attack Intensity: Definition and Proxies	15
2.3 Cost Structure	16
2.4 Reference Loss Construction and Observational Loss Suppression	17
2.5 Decision-Usable Marginal Efficiency: CPLA	18
2.6 Cost Sustainability and Composite Indicators	18
2.7 Counterfactual Stress-Test Module	19
2.8 Mission Outcome Anchoring: Key Asset Preservation	19
2.9 Interpretation Logic	20
3. Assessment Based on Current Public Data of the Russia–Ukraine War	21
3.1 Public Data Basis and Primary Proxies	21
3.2 National-Level Damage Context	21
3.3 Critical-Asset Loss Proxies	21
3.4 Attack Intensity Proxies	22
3.5 Interpretation of KAPS	23
3.6 Scope and Caution	23

Research Report

Contents

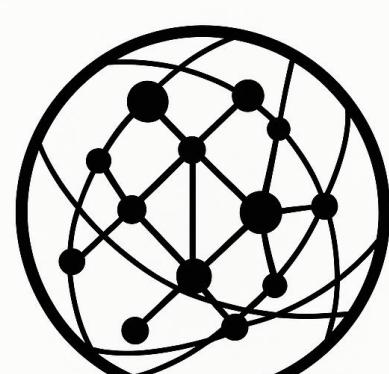
Series Information:
This research report is part of the EPINOVA Research Report Series on AI-Enabled Warfare, Sustainability, and Global Security Governance.

Recommended Citation:
Wu, S. (2025). *From Detection to Depletion: Cost-Exchange Limits in the Russia–Ukraine Drone War* (Research Report). EPINOVA–2025–01–RR. Global AI Governance Research Center, EPINOVA LLC.

Short Citation:
Wu (2025), *From Detection to Depletion*, EPINOVA RR–2025–01.

Disclaimer:
This report reflects the author's independent analysis based on publicly available information and does not represent the official views or positions of any government, organization, or institution.

DOI:
<https://doi.org/10.5281/zenodo.18036790>

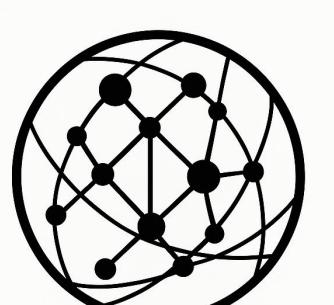


GLOBAL AI
GOVERNANCE
RESEARCH CENTER

4. Baseline Results (2023–2025)	24
4.1 Core Unified Assumptions	24
4.2 Observational Sustainability and Mission Outcomes	25
4.3 Counterfactual Stress-Test Results (2025)	27
4.4 Findings	28
5. Interpretation and Limitations	29
5.1 Interpretation of Baseline Results (2023–2025)	29
5.2 Practical Use Guidance	31
5.3 Key Caveats: What the Baseline Can and Cannot Claim	31
5.4 Roadmap for Refinement	32
6. Policy Options	33
6.1 Mission-Based Defense: Prioritize Critical Assets and Resilience	33
6.2 Rebalance the Intercept Mix Toward Low-Cost Terminal Layers	33
6.3 Invest in Detection and Track Continuity	34
6.4 Magazine Depth and Sustainment Strategy	34
6.5 Recommended Decision Criterion	35
Conclusion	36
References	38
Appendix A: Abbreviations and Glossary	41
Appendix B: Tables	43
Appendix C: Decision Rules for Interpreting Counter-Drone Performance under Saturation	45



Executive Summary



Executive Summary

By late 2025, the Russia–Ukraine war demonstrates that unmanned aerial systems (UAS) have shifted from a tactical supplement to a system-level driver of military sustainability. Low-cost, mass-producible, expendable, and networked drones—combined with AI-enabled sensing, targeting, and coordination—have altered both the economics and operational logic of air defense. Under sustained saturation, defensive systems face simultaneous constraints across three dimensions: **physical limits** (low-altitude clutter and weak signatures), **temporal limits** (mismatches across sensing, decision, and interception timelines), and **resource limits** (magazine depth, operational tempo, and fiscal sustainability). In this environment, performance metrics centered solely on intercept or shoot-down rates become increasingly misleading.

This report addresses a core policy question emerging from the conflict: not whether defenses can intercept more drones, but whether they can **sustain an acceptable cost–loss exchange while preserving critical national functions over time**. Across the 2023–2025 period, indicators such as Cost per Loss Avoided (CPLA)—the marginal cost of preserving one unit of critical function—and the composite Cost Exchange Ratio (CER^{*}) consistently deteriorate before intercept rates visibly decline. This pattern highlights their value as early-warning signals of structural exhaustion. Using a Minimum Viable, Auditable (MVA) analytical baseline grounded in publicly available data and conservative proxies, the study reframes counter-drone effectiveness away from intercept-centric measures toward sustainability- and mission-outcome-oriented evaluation.

The analytical framework integrates **five** complementary indicators. **CER** (Cost Exchange Ratio) captures relative economic efficiency between attack and defense; **OLER** (Observational Loss Suppression Ratio) measures asset-level functional loss under observed attack pressure; **CER^{*} _{obs}** tracks the joint evolution of cost pressure and residual loss; **CPLA_{obs}** quantifies the marginal cost of preserving one unit of critical function; and **KAPR/KAPS** (Key Asset Preservation Ratio/Score) anchors assessment in mission outcomes such as power continuity and fuel supply. Together, these metrics translate the intuitive dilemma of high-cost interceptors versus low-cost threats into a structured, comparable, and decision-relevant framework.

Applied to the 2023–2025 period, the baseline assessment yields three central findings.

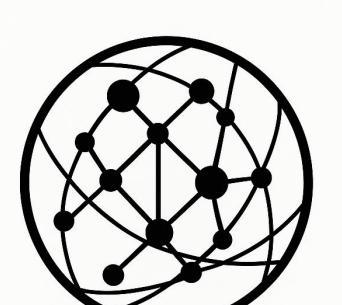
First, **high intercept performance does not imply sustainability**: defenses can maintain shoot-down rates by relying on increasingly expensive interceptors, while CPLA and CER^{*} deteriorate well in advance of visible declines in interception performance. Second, **scale and saturation make unit cost curves decisive**: as attack frequency rises, long-term outcomes depend less on peak technical capability than on the availability of low-cost terminal interception layers capable of absorbing volume. Third, **mission outcomes matter more than counts**: in both Ukrainian air defense and Ukrainian deep strikes against Russian energy infrastructure, the preservation—or degradation—of critical asset functionality, as captured by KAPS, provides a more accurate measure of strategic effect than interception or strike tallies alone.

Executive Summary

Beyond the immediate conflict, these findings point to a broader shift in the logic of warfare. Tactically, effective counter-drone defense increasingly requires layered architectures centered on low-cost terminal interception, multi-sensor fusion, and high-frequency command-and-control adaptation. Strategically, counter-UAS capability must be treated as an endurance capability, the ability to sustain function under continuous pressure, rather than as a measure of peak interception performance. At the governance level, the expanding operational role of AI-assisted detection and automated engagement intensifies unresolved questions of responsibility, proportionality, and meaningful human control, pushing these debates from abstract principles toward institutional stress testing.

Finally, this report emphasizes methodological restraint. Because the analysis relies on open-source data and proxy variables, especially for 2025, all absolute values should be interpreted as illustrative and provisional. The framework is designed for trend comparison, sensitivity analysis, and policy exploration, rather than precise attribution. Priority areas for future research include auditable data on interceptor mix and unit costs, asset-level recovery and resilience curves, and attack data disaggregated by wave intensity and decoy composition. Progress along these lines would enable the present MVA framework to support higher-confidence budgetary, force-structure, and strategic planning decisions under sustained drone saturation.

Background



Background

As of 14 December 2025, approximately three years and nine months (≈ 3.8 years) have elapsed since Russia launched its full-scale invasion of Ukraine on 24 February 2022. Over the course of this prolonged, high-intensity war of attrition, one of the most consequential structural transformations has been the large-scale integration of artificial intelligence (AI) into battlefield systems, alongside the rise of unmanned platforms, including unmanned aerial vehicles (UAVs), loitering munitions, and unmanned surface vehicles (USVs), as dominant instruments **of low-cost, high-frequency, and expendable strike**.

Together, these developments have accelerated a transition in the character of warfare from platform-centric force employment toward a **sensor–algorithm–munition network–centric paradigm**, in which detection, targeting, and engagement are increasingly governed by data flows and automated or semi-automated decision processes.

1. Drones as a Paradigm-Shaping Technology

The paradigm-level significance of drones in the Russia–Ukraine war does not arise from superior performance at the level of individual platforms. Rather, drones have functioned as a scaling mechanism that has brought several latent tensions of modern warfare into sustained empirical exposure at unprecedented volume.

First, drones have contributed to a **compression of the kill chain and increased battlefield transparency**. Across reconnaissance, targeting, guidance, and terminal strike phases, a growing share of functions is assisted, or partially automated, by algorithms. The interval between target detection and engagement has been reduced to minutes or even seconds, rendering the battlespace increasingly observable, traceable, and contestable in near real time.

Second, drones have driven a **reconstruction of air-defense economics**, often summarized as “shooting missiles at mosquitoes.” The persistent employment of inexpensive unmanned systems has imposed continuous pressure on defenders’ stocks of costly interceptors and associated readiness resources. This asymmetry has forced a reorientation toward cheaper countermeasures and denser, layered defensive architectures. Ukraine, for example, has reportedly explored **AI-assisted and partially automated counter-drone firing concepts** as a means of restoring cost balance at scale.

Third, the routine integration of **mass-expendable unmanned systems, electronic warfare, and algorithm-assisted detection and cueing** has produced spillover effects on doctrine, organizational practice, and legal–ethical frameworks. Established approaches to air defense, depth protection, and critical infrastructure security are increasingly strained. At the same time, foundational legal and ethical principles—distinction, proportionality, attribution of responsibility, and the feasibility of meaningful human control—are being tested in practice rather than debated in abstraction.

Beyond their immediate tactical utility, the large-scale employment of drones in the Russia–Ukraine war has exposed **structural mismatches within legacy military systems**, particularly on the defensive side, revealing the difficulty of adapting inherited force postures and operational doctrines to a drone-dominated battlespace.

Background

2. Structural Mismatches in Legacy Defensive Systems

One such mismatch lies in **legacy force structures and sensor architectures optimized for high-signature threats**. A significant portion of Russia's air-defense inventory was designed under earlier technological assumptions, prioritizing the detection and interception of large, fast-moving, high-altitude platforms with strong radar and infrared signatures, such as aircraft, cruise missiles, and ballistic threats. Contemporary battlefield drones, by contrast, are typically small, low-flying, slow-moving, highly maneuverable, and characterized by weak or intermittent electromagnetic and thermal signatures. This divergence reduces the effectiveness of traditional radar cueing, missile interceptors, and engagement envelopes when confronted with massed, low-cost unmanned systems. The result is not merely degraded performance, but a **systemic misalignment between sensor design assumptions and the dominant threat profile** of modern conflict.

A second mismatch appears at the **terminal defense layer**, where doctrinal and procedural adaptation has lagged behind operational reality. Many existing rules of engagement, fire-control procedures, and target-classification thresholds remain calibrated for manned or high-value platforms. Such standards are poorly suited to environments characterized by swarms of low-speed, expendable, and semi-autonomous systems. Effective defense under these conditions requires updated engagement doctrines, revised identification and discrimination criteria, accelerated decision cycles, and greater tolerance for automated or semi-automated responses. Absent such adaptation, defenders face chronic delays, over-reliance on high-cost interceptors, and an inability to scale responses proportionally to the volume and persistence of drone threats.

3. From Tactical Adaptation to Paradigm Exhaustion

Taken together, these dynamics suggest that the strategic significance of drones in the Russia–Ukraine war lies not in any single technological breakthrough, but in the **exhaustion of an inherited defensive paradigm** under conditions of sustained saturation and cost asymmetry. Air-defense architectures, sensor systems, and engagement doctrines developed to counter scarce, high-value, high-signature threats are being persistently stressed by the opposite logic: ubiquitous, low-cost, low-signature, and expendable systems deployed at scale.

This mismatch manifests most acutely at the terminal layer of defense, where **human-centered decision cycles**, legacy rules of engagement, and interceptor-centric response models prove increasingly unable to cope with the speed, volume, and ambiguity of drone threats. Under sustained saturation, the requirement for real-time discrimination and proportional response exceeds what manual or purely human-led processes can reliably deliver. As a result, operational pressure mounts to delegate detection, classification, cueing, and, in some cases, firing decisions to algorithm-assisted or partially automated systems—not as a matter of preference, but of necessity.

Background

Crucially, this transition marks a **qualitative shift in the governance of force**. When defensive effectiveness depends on machine-speed responses and automated engagement thresholds, the traditional assumption that meaningful human control can be exercised at each step of the kill chain becomes increasingly fragile. Legal and ethical frameworks grounded in human deliberation are thus strained in practice by the temporal and economic realities of contemporary warfare.

In this sense, drone warfare functions as a **paradigm-shaping stress test**, revealing not only the limits of existing military technologies, but also the boundaries of doctrinal, institutional, and legal systems designed for a different tempo, scale, and ontology of conflict.

4. Scale and Trends

Because the classification and reporting of “drones” vary widely, encompassing attack platforms, decoys, reconnaissance UAVs, FPV systems, and maritime unmanned assets, and because much information is shaped by wartime information operations, the figures below are used **solely to illustrate scale and trend**, with original reporting categories preserved.

Russian deep strikes against Ukraine (Shahed/Geran family; proxy indicators):

- **13 September 2022 – 30 August 2023:** Airwars records approximately **1,956 Shahed launches**, capturing the early phase of scale formation.
- **2024 (to 1 November):** Ukraine’s General Staff reports **6,987 attack UAVs** launched since the start of the year.

2025 (high-saturation phase):

- Sky News, citing CSIS, reports **over 44,000 Shahed and variants** launched in the first ten months of 2025, alongside declining interception ratios relative to 2024.
- CEPA estimates **over 38,000 launches** (including decoys) between January and November 2025, with pronounced monthly peaks.

Ukrainian deep strikes against Russia (energy and refining infrastructure; proxy indicators):

- **2024:** BBC Russian reporting, relayed by Meduza, documents **at least 81 drone strikes** on Russian oil, gas, and refining facilities.
- **2025:** Reuters reports **more than 60 strikes** on Russian energy facilities since early August, with a narrower visual investigation identifying **at least 58 incidents**.

5. From Scale to Sustainability

Taken together, these data indicate that drones have evolved from a supplementary battlefield

Background

tool into a **central variable shaping deep-strike dynamics, air-defense depletion, energy and industrial resilience, and cost distribution**. Drones function simultaneously as weapons, instruments of economic attrition, and vectors of information-psychological warfare, tightly coupled with AI-enabled sensing and decision chains.

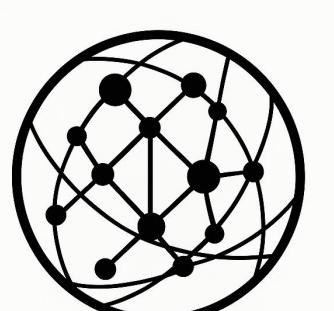
Operational experience between **2023 and 2025** has exposed a persistent set of frictions, including:

- detection and tracking failures in ultra-low-altitude and cluttered environments;
- temporal mismatches between sensor scan cycles and target maneuvering;
- limitations of infrared and passive sensing against low-thermal-signature platforms;
- reduced radar cross-sections driven by materials and airframe design;
- spectrum congestion and escalation dynamics in electronic warfare contests; and
- most prominently, **extreme cost asymmetries**, in which per-intercept costs vastly exceed per-target costs.

These frictions manifest not as isolated technical shortcomings, but as **volatile cost-exchange outcomes**, observable through indicators such as **CER, CPLA, and loss-exchange ratios**.



Methodology and Auditability



Methodology and Auditability

1. Objective and Analytic Standard

This report adopts a **Minimum Viable, Auditable (MVA)** standard to evaluate counter-drone sustainability under real-world conflict conditions. The objective is to establish a **baseline analytical framework** that can be replicated, cross-checked, and stress-tested using **publicly available information**, without reliance on classified sources or unverifiable assumptions.

Under the MVA standard, all key variables and indicators are required to satisfy three conditions:

- **traceable sources and definitional clarity**;
- **explicit inclusion and exclusion rules**; and
- **bounded uncertainty ranges**, explored through sensitivity analysis rather than point assertions.

Where direct measurement is infeasible, the analysis relies on **conservative proxy variables**, prioritizing systematic underestimation over overstatement. A data dictionary and calculation workflow are provided in the Appendix to support independent replication and auditability.

2. Unit of Analysis and Scope Discipline

The basic unit of analysis is defined as **year × operational side (attacker/defender) × critical asset set**. This structure enables longitudinal comparison while maintaining alignment with **mission-relevant outcomes**, rather than aggregate or symbolic measures of wartime damage.

Attack activity, defensive effort, and realized losses are treated as **analytically distinct dimensions**. Measures of attack intensity are used strictly as **pressure proxies** and are not interpreted as direct counts of effective strike platforms. Outcome variables are confined to **functional impacts on critical assets**, rather than national-level damage totals, in order to preserve analytic focus and cross-period comparability.

3. Inclusion and Exclusion Principles

To prevent indicator inflation and category drift, the following principles govern variable construction throughout the report.

Included outcomes are limited to **functional degradation or restoration costs** of clearly defined critical assets, such as electricity supply interruptions, refinery downtime, or infrastructure repair proxies.

Excluded outcomes include **aggregate national damage estimates**, political signaling effects, and indirect macroeconomic losses that cannot be directly attributed to drone-related strikes. Such figures may be referenced only for **contextual background**, not as inputs to analytic indicators.

Methodology and Auditability

Ambiguous observations, such as mixed-use platforms, suspected decoys, or partially reported incidents, are handled through **conservative reconciliation rules** and explicitly parameterized uncertainty bands, rather than discretionary inclusion. These principles ensure that reported indicators reflect **operationally meaningful loss suppression and preservation dynamics**, rather than rhetorical or symbolic measures.

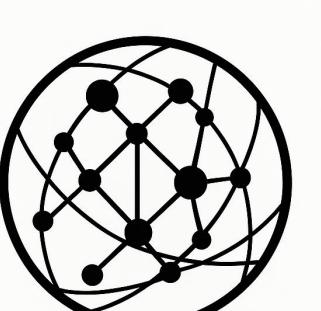
4. Replicability and Cross-Verification

Each core variable is cross-verified against **at least two independent source categories**, including official assessments, established independent trackers, and corroborated open-source intelligence (OSINT) datasets. Where discrepancies arise, the baseline adopts the **more conservative interpretation**, with alternative bounds explored through sensitivity analysis.

Counterfactual assumptions are **strictly separated from observational metrics**. Scenario-based parameters are introduced exclusively for stress testing and are never permitted to mechanically determine observed results. This separation preserves the interpretive integrity of trend analysis while enabling controlled exploration of extreme or stress conditions.



Observed Operational Frictions in the Drone Fight



1

Observed Operational Frictions in the Drone Fight

This section documents the principal operational frictions observed in the drone contest during the Russia–Ukraine war. Rather than treating these frictions as isolated technical shortcomings, the analysis organizes them along the **detection–identification–tracking–engagement–assessment–cost-exchange chain**, highlighting how localized failures propagate into **system-level sustainability constraints** under conditions of sustained saturation.

1.1 Sensor Physics at Low Altitude: When “Detectable” Does Not Mean “Engageable”

Operational experience shows that **ultra-low-altitude flight profiles**, dense ground clutter, and strong surface reflections consistently degrade radar and optical performance. Even when targets are intermittently detected, defenders often struggle to achieve **fire-control–quality tracking**, producing gaps between initial detection and actionable engagement.

This effect reflects a classic **signal-to-clutter ratio (SCR) degradation problem**. Low-altitude targets operate within environments dominated by terrain, vegetation, and urban structures, where weak target returns are easily masked by background noise. Detection thresholds must therefore be raised, increasing miss rates. Track stability further degrades as frequent clutter-induced drops prevent reliable prediction. RAND analyses similarly emphasize that small UAVs are significantly harder to detect, classify, and track in high-clutter environments, particularly in urban settings.

Low-altitude geometry compounds these challenges. As targets approach the radar horizon, elevation-angle resolution compresses, while multipath effects and ground reflections undermine the generation of reliable fire-control solutions. In such conditions, **“seeing” a target does not reliably translate into an ability to engage it**.

Operationally, this friction shifts the primary bottleneck in counter-drone defense from interceptor availability to **early detection and track continuity**. Without stable tracks, defenders cannot confidently allocate fires or exploit low-cost terminal interception layers. This dynamic helps explain Ukraine’s emphasis on combining mobile fire units, electro-optical systems, electronic warfare, and interceptor drones, rather than relying exclusively on surface-to-air missiles.

1.2 Temporal Mismatch: Scan–Fusion–Decision Latency Versus Drone Maneuver Cycles

A second friction arises from **temporal mismatches** between drone maneuver cycles and air-defense system update rates. Many low-speed drones, often operating around 180 km/h, execute random or evasive maneuvers on timescales shorter than some radar scan intervals. Delays accumulate across sensor refresh, data fusion, decision-making, and engagement authorization.

Air-defense kill chains consist of sequential stages—sensor refresh, track filtering and fusion, identification and prioritization, authorization, weapon cueing, and terminal guidance. Latency at

1

Observed Operational Frictions in the Drone Fight

any stage compounds downstream. When a target's maneuver cycle is shorter than system update intervals, **track divergence** emerges: by the time a new update arrives, the target may fall outside predicted gates, degrading track confidence or causing track loss altogether.

Importantly, low speed does not necessarily make drones easier to defeat. In cluttered low-altitude environments, slow-moving targets may resemble ground traffic or environmental noise, increasing classification ambiguity. Moreover, slower drones extend engagement timelines, imposing sustained strain on defender readiness and ammunition expenditure. CSIS analyses of Russian saturation tactics explicitly frame this temporal pressure as a **deliberate operational advantage**, rather than a limitation.

From an evaluation perspective, indicators such as **sensor-to-shooter latency, track continuity, and engagement opportunity density** often explain penetration outcomes more effectively than nominal detection range or interceptor performance.

1.3 Infrared and Electro-Optical Constraints: The Thermal Contrast Problem

Infrared and electro-optical sensing faces persistent limitations against electric motor–driven drones, which frequently exhibit surface temperatures close to ambient conditions. Under many environmental conditions, this sharply reduces acquisition ranges and lock-on reliability.

The underlying constraint is **thermal contrast (ΔT)** rather than absolute temperature. Battery-powered drones often present weak or fragmented heat signatures that approach sensor noise levels, especially against warm backgrounds. The result is shorter detection ranges, higher false-alarm rates, and longer lock-on times, each of which compresses viable engagement windows.

Operationally, these constraints reinforce the **economic and logistical unsustainability** of using expensive surface-to-air missiles against low-signature drones, accelerating efforts to deploy cheaper terminal defenses such as guns, interceptor UAVs, and layered point-defense systems.

1.4 Reduced Observability Is Not Stealth: Multi-Sensor Fusion as a Necessity

Small airframes, composite materials, and low radar cross-sections, combined with visual blending against ground backgrounds, significantly complicate detection. However, reduced observability does not constitute true stealth.

By lowering radar and optical signatures, drones force defenders toward two costly alternatives: deploying denser, higher-sensitivity sensor networks, or accepting partial leakage while prioritizing protection of critical assets over universal interception. RAND and JAPCC frameworks consistently emphasize that effective counter-UAS operations require **multi-layered, multi-sensor fusion**, integrating radar, electro-optical, acoustic, RF, and ground-based intelligence rather than reliance on any single modality.

1

Observed Operational Frictions in the Drone Fight

1.5 Diminishing Communications Visibility: From Link Suppression to Platform Defeat

As drone autonomy increases, the visibility of command-and-control links has diminished. Emissions control, pre-programmed routes, inertial navigation, and terrain matching reduce reliance on continuous external communications, limiting the effectiveness of traditional signal interception and jamming.

This shift does not eliminate the relevance of electronic warfare, but it **changes its function**. Rather than serving as a decisive single-point intervention, EW increasingly operates as a **system-level suppression tool**, whose effectiveness depends on integration with sensing, timing, and engagement processes. RAND research highlights that EW effects must now be synchronized with detection and terminal defense rather than treated as a standalone solution.

1.6 Spectrum Congestion and Frequency Agility: Entering the System-Overload Regime

High rates of frequency hopping and mutual interference have produced congested and dynamically unstable electromagnetic environments. As both sides employ EW as a default capability, increased jamming power yields diminishing returns.

Instead of linear gains, congestion degrades friendly systems, increases error rates, and destabilizes coordination. These challenges are rarely solvable through individual platform upgrades; they require coordinated spectrum management, task allocation, and rapid reconfiguration cycles—a point emphasized in RAND research on small-UAS and EW integration.

1.7 Cost–Magazine Frictions: From “Can Intercept” to “Can Sustain”

Perhaps the most consequential friction is economic. Persistent asymmetries between the unit cost of interceptors and that of incoming drones generate the classic **“missile-versus-drone” dilemma**.

The core issue is not solely the price of individual interceptors, but the interaction between **magazine depth, replenishment rates, and sustained attack volume**. CSIS analyses of the Shahed campaign characterize low-cost, mass-deployable drones as a strategic tool designed to push defenders onto unfavorable cost curves.

Recent Ukrainian efforts to deploy interceptor drones as a **low-cost aerial shield** reflect an attempt to correct this imbalance by substituting cheaper engagement layers for portions of missile-based defense, thereby preventing uncontrolled escalation of cost-exchange ratios.

Analytically, this friction underscores why counter-drone effectiveness under saturation must be evaluated using indicators such as **Cost per Loss Avoided (CPLA)**, **composite cost–loss ratios (CER*)**, and **key asset preservation metrics**, rather than intercept rates alone.

1

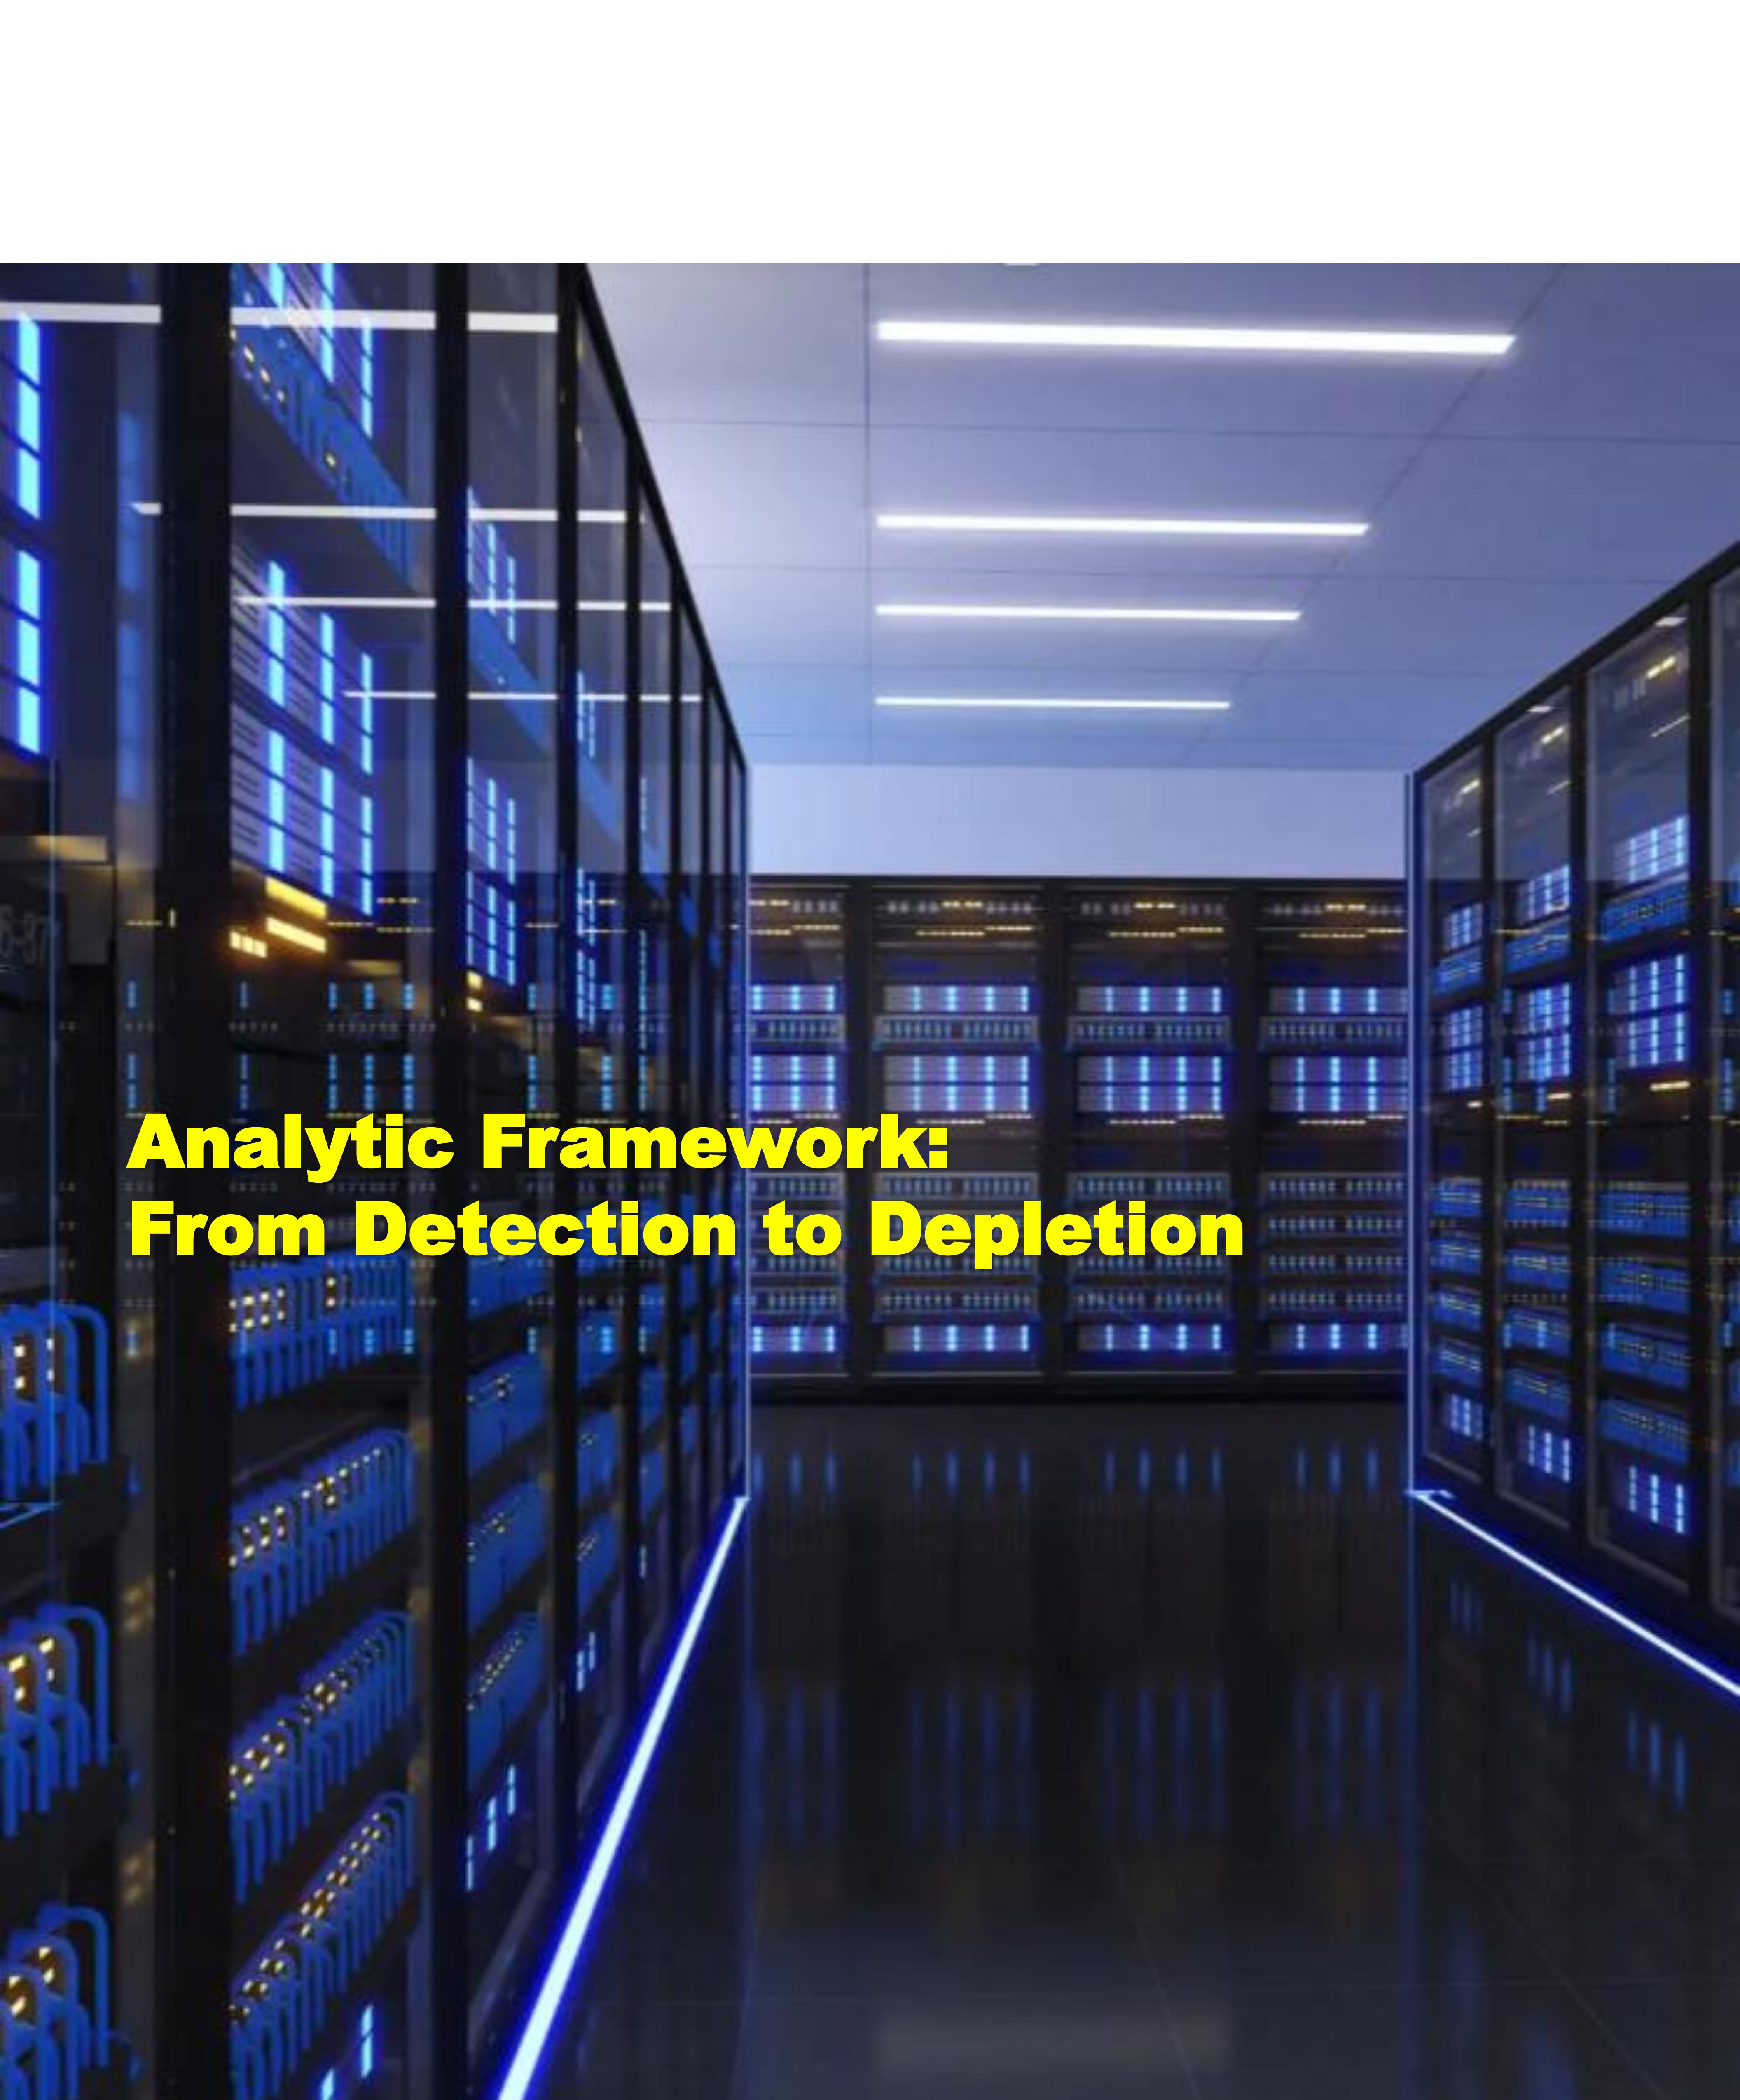
Observed Operational Frictions in the Drone Fight

1.8 Structural Conclusion: System-on-System Friction, Not Platform Competition

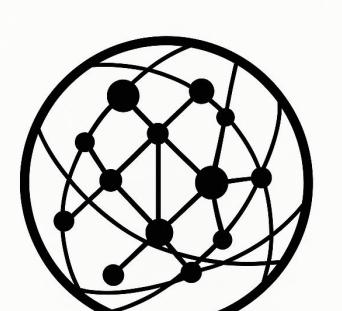
Taken together, these frictions indicate that the central challenge of drone warfare lies not in defeating individual platforms, but in operating within a regime constrained simultaneously by **physical limits** (low-altitude clutter and weak signatures), **temporal limits** (refresh, fusion, and authorization delays), **spectrum limits** (congestion and mutual interference), and **economic limits** (magazine depth and replenishment capacity).

Russian employment of scale, decoys, and iterative adaptation has consistently pushed defenders toward saturation thresholds, even as interception performance improves at the margin. Defensive responses, in turn, increasingly emphasize cheaper terminal layers and drone-on-drone interception, reserving high-end missiles for fast, high-value threats.

This **system-on-system interaction** provides the analytical bridge to the next section, which formalizes observed frictions into **auditable sustainability indicators** capable of comparing defensive performance across time, force structures, and cost regimes.



Analytic Framework: From Detection to Depletion



2

Analytic Framework: From Detection to Depletion

This section introduces a sustainability-oriented evaluation framework designed to answer a single question: whether counter-drone defenses can preserve critical functions at acceptable cost over time. The framework proceeds in three steps:

- define **comparable attack pressure**;
- measure **observed loss suppression and marginal cost**; and
- anchor results in **mission outcomes** rather than intercept counts.

Under sustained drone saturation, counter-UAS performance cannot be assessed through intercept rates alone. A defense may preserve high nominal shoot-down ratios by shifting toward expensive or scarce interceptors, while drifting toward strategic unsustainability through rising marginal costs, magazine depletion, and operational fatigue. Accordingly, this framework evaluates counter-drone effectiveness along three linked dimensions: **cost sustainability, loss suppression, and mission outcome preservation**.

2.1 Unit of Analysis and Minimum Viable, Auditable Standard

The basic unit of analysis is **year × defending side × target set**. To ensure cross-year comparability and auditability, the report adopts a **MVA** standard:

- all variables are derived from **publicly available and cross-verifiable sources**;
- where direct observation is unavailable, **conservative proxies** are used and explicitly labeled;
- **counterfactual parameters** are strictly separated from observational metrics and reported as **scenario bands**.

2.2 Attack Intensity: Dual-Track Definition of N

Because public reporting frequently aggregates attack drones, decoys, reconnaissance UAVs, and mixed strike waves, attack intensity is represented in two forms:

- **N_{total}** : Total reported aerial objects associated with drone attack activity (including decoys);
- **N_{attack}** : Estimated number of effective attack platforms.

The relationship between the two is modeled as:

$$N_{attack} = N_{total} \cdot (1 - d) \quad (\text{Eq. 1})$$

where **d** represents the decoy proportion. Because **d** is not directly observable, it is treated as a **sensitivity parameter** with Low / Base / High values in scenario analysis. All core indicators are computed using **N_{attack}** to avoid inflating effective attack pressure.

Attack intensity variables defined in this section enter the cost and loss modules that follow; outcome indicators are formally introduced in Sections 2.4–2.8.

2

Analytic Framework: From Detection to Depletion

2.2.1 Wave Intensity Proxy (Robustness Add-on)

Annual totals conceal the operational reality that counter-UAS failures often occur during peak salvos. To minimally represent saturation without requiring classified telemetry, define a wave intensity proxy.

The wave-intensity adjustment does not introduce new outcome variables. When $\beta > 0$, the adjusted effective pressure $N_{\text{attack},y}^{\text{eff}}$ replaces $N_{\text{attack},y}$ **only as an input** in subsequent modules—specifically: (1) the engagement count \mathbf{E} used in defensive cost construction (Section 2.3.1), and (2) the reference-loss construction $\mathbf{D}_{y,\text{ref}}$ (Section 2.4). All outcome indicators remain defined and interpreted exclusively in their respective sections below.

$$w_y = \frac{N_{\text{peak},y}}{\bar{N}_{\text{day},y}} \quad (\text{Eq. 2})$$

- $N_{\text{peak},y}$: maximum reported daily (or single-wave) inbound objects in year y from tracker/official daily reports.
- $\bar{N}_{\text{day},y}$: average daily inbound objects in the same reporting series.

Optionally adjust effective pressure for robustness testing:

$$N_{\text{attack},y}^{\text{eff}} = N_{\text{attack},y} \cdot [1 + \beta \cdot (w_y - 1)] \quad , \beta \in [0, 0.3] \quad (\text{Eq. 3})$$

MVA baseline uses $\beta = 0$ (no adjustment), and reports $\beta > 0$ as robustness only.

Baseline: $\beta = 0$.

Robustness: when $\beta > 0$, replace N_{attack} with $N_{\text{attack},y}^{\text{eff}}$ consistently in (1) engagement count \mathbf{E} ; and (2) reference-loss construction $\mathbf{D}_{y,\text{ref}}$ (See below).

2.3 Cost Structure

2.3.1 Defense Cost

Defense cost is decomposed as:

$$C_{\text{def}} = C_k + C_{\text{ops}} + C_{\text{attr}} \quad (\text{Eq. 4})$$

where:

- C_k : kinetic and non-kinetic interception costs;
- C_{ops} : operations, readiness, maintenance, and personnel costs;
- C_{attr} : attrition and depreciation of defensive assets (optional in MVA baseline).

2

Analytic Framework: From Detection to Depletion

The core component \mathbf{C}_k is modeled as a weighted interception mix:

$$\mathbf{C}_k = \sum_{m \in \{SAM, AAA, EW, ID\}} s_m \cdot E \cdot u_m \quad (\text{Eq. 5})$$

where:

- s_m : the engagement share of method m ;
- E : number of defensive engagements (proxied by N_{attack} or $N_{\text{attack}}^{\text{eff}}$);
- u_m : unit cost of method m (public cost ranges).

To preserve auditability, three interception-mix scenarios are defined: **SAM-heavy**, **Balanced**, and **Low-cost-heavy**, allowing CER and CPLA to be evaluated as functions of force-structure choice rather than fixed constants.

2.3.2 Attack Cost (C_{atk})

$$C_{\text{atk}} = N_{\text{attack}} \cdot u_{\text{atk}} \quad (\text{Eq. 6})$$

where u_{atk} represents the estimated unit cost of attack drones, modeled as a public cost range reflecting scale, substitution, and decoy inclusion.

2.4 Reference Loss Construction and Observational Loss Suppression

To ensure replicability under a MVA standard, the reference loss $D_{y,\text{ref}}$ is generated by a fixed rule rather than analyst discretion.

2.4.1 Baseline-year anchoring rule (MVA default).

Select a baseline year y_0 with the most stable reporting coverage (default: **2023**). Define:

$$\lambda = \frac{D_{y_0}}{N_{\text{attack},y_0}} \Rightarrow D_{y,\text{ref}} = \lambda \cdot N_{\text{attack},y}$$

a) Interpretation boundary:

$D_{y,\text{ref}}$ is a **comparability anchor**, not a “no-defense counterfactual.” It represents the expected functional loss if the **baseline-year loss-per-effective-attack relationship** remained unchanged under observed effective attack pressure.

b) Optional saturation refinement:

To partially capture wave/saturation dynamics without requiring classified data, a bounded nonlinearity can be introduced:

$$D_{y,\text{ref}}^{\text{sat}} = \lambda \cdot N_{\text{attack},y}^{\alpha}, \alpha \in [1.0, 1.3] \quad (\text{Eq. 7})$$

2

Analytic Framework: From Detection to Depletion

The report retains $\alpha = 1.0$ as the MVA baseline and reports $\alpha > 1.0$ only as a robustness band.

2.4.2 Observational Loss Suppression Ratio (OLER)

$$OLER_y = \frac{D_y}{D_{y,ref}} \quad (\text{Eq. 8})$$

- D_y : observed critical-asset functional loss proxy in year y (e.g., blackout-hours, refinery downtime-hours, or normalized functional loss index).
- $D_{y,ref}$: reference loss generated by the fixed rule above.

a) Interpretation.

Lower $OLER_y$ value indicates stronger loss suppression under comparable effective attack pressure. OLER is an **observational** metric and does not embed counterfactual multipliers.

2.5 Decision-Usable Marginal Efficiency: CPLA

Define loss avoided:

$$\Delta D_y = D_{y,ref} - D_y$$

Then:

$$CPLA_{obs,y} = \frac{c_{def,y}}{\Delta D_y} \quad (\text{Eq. 9})$$

a) Unit discipline:

- If D is measured in **blackout-hours**, then CPLA is **USD per blackout-hour avoided**.
- If D is a **normalized loss index** (0–1), then CPLA is **USD per 0.01 loss-index avoided** (recommended to scale by 0.01 so values are interpretable).

b) Interpretation:

CPLA is the primary resource-allocation indicator. It shows the marginal dollars required to preserve one unit of mission-relevant function.

2.6 Cost Sustainability and Composite Indicators

Cost Exchange Ratio (CER) is defined as:

$$CER = \frac{C_{def}}{C_{atk}} \quad (\text{Eq. 10})$$

where **CER** captures the economic asymmetry of the contest but does not, by itself, indicate effectiveness. C_{def} denotes total defensive cost and C_{atk} denotes total attack cost.

2

Analytic Framework: From Detection to Depletion

a) Observational composite indicator:

$$CER_{obs}^* = CER \cdot OLER \quad (\text{Eq. 11})$$

CER_{obs}^* captures the **joint evolution of cost pressure and residual functional loss** under observed conditions. It increases when defense becomes more expensive **and/or** when loss suppression weakens, even if intercept rates remain high.

2.7 Counterfactual Stress Test Module

To explore degraded-defense or saturation scenarios, introduce a loss amplification multiplier **M**:

$$CLER = \frac{1}{M}$$

The counterfactual composite indicator is defined as:

$$CER_{cf}^* = CER \cdot CLER$$

a) Interpretation:

- CER_{cf}^* is not an observed metric.
- It is used exclusively for scenario comparison and stress testing under explicitly stated assumptions about defensive degradation.
- All counterfactual outputs are reported as scenario bands, not point estimates, and should not be interpreted as empirical reality.

2.8 Mission Outcome Anchoring: Key Asset Preservation

For each critical asset class **i** in year **y**, define the **Key Asset Preservation Ratio (KAPR)**:

$$KAPR_{i,y} = 1 - \frac{D_{i,y}}{D_{i,ymax}} \quad (\text{Eq. 12})$$

where the upper-bound degradation scenario is defined as:

$$D_{i,y}^{max} = D_{i,y} \cdot (1 + \rho_i)$$

Here, ρ_i represents an asset-specific stress factor capturing plausible worst-case escalation or cascading failure.

The aggregate outcome measure is:

$$KAPSy = \sum_i (w_i \cdot KAPR_{i,y}), \quad \sum_i w_i = 1 \quad (\text{Eq. 13})$$

In the baseline implementation, the asset set is limited to electricity and energy/refining infrastructure to preserve auditability.

2

Analytic Framework: From Detection to Depletion

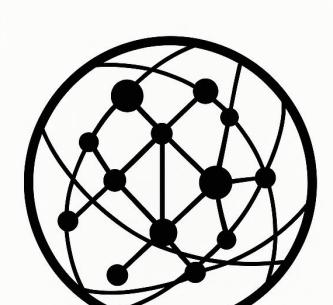
2.9 Interpretation Logic

These rules are **heuristic classification criteria** rather than deterministic thresholds.

- Low CER + Low OLER + High KAPS → Sustainable and effective defense;
- High CER + Low OLER + High KAPS → Effective but economically stressed;
- Low CER + High OLER + Low KAPS → Inexpensive but ineffective;
- High CER + High OLER + Low KAPS → System approaching saturation or failure.



Assessment Based on Current Public Data of the Russia–Ukraine War



3

Assessment Based on Current Public Data of the Russia–Ukraine War

3.1 Public Data Basis and Primary Proxies (Illustrative)

This assessment employs a **MVA** standard to enable cross-year comparison while maintaining transparency and replicability under conditions of limited access to classified military data. The objective is not to produce a definitive accounting of drone-only losses or precise defense expenditures, but to establish a consistent and conservative baseline suitable for trend analysis, sensitivity testing, and policy-relevant inference.

Under the MVA approach, the analysis prioritizes authoritative, publicly verifiable sources and relies on explicit proxy variables and scenario parameters where direct observation is unavailable. All proxies are selected according to three criteria: (1) consistency across years, (2) traceability to primary sources, and (3) conservative bias to avoid overstating effects.

Crucially, the assessment distinguishes between **national-level damage context** and **critical-asset–level operational effects**, using each strictly for analytically appropriate purposes.

3.2 National-Level Damage Context

For Ukraine, national-level damage estimates are drawn from the Rapid Damage and Needs Assessment (RDNA) series jointly produced by the Government of Ukraine, the World Bank, the European Union, and the United Nations. These figures are used solely as contextual background indicators of overall war damage and are not interpreted as drone-attributable losses.

Key reference points include:

- **RDNA3:** Estimated direct damage of approximately **USD 152 billion** as of **31 December 2023**.
- **RDNA4:** Estimated direct damage of approximately **USD 176 billion** as of **31 December 2024**.

These estimates aggregate damage from multiple sources—including missile strikes, artillery, ground combat, and drone attacks—and therefore do not isolate the marginal contribution of unmanned systems. Their function in this report is limited to providing a stable macro-level loss context against which the scale of counter-drone resource allocation and strategic pressure can be understood.

3.3 Critical-Asset Loss Proxies

This section establishes the macro damage context to prevent misattribution of drone-specific effects.

For operational analysis—specifically the computation of **OLER**, **CPLA**, and **KAPS**—the report relies on **critical-asset–level functional loss proxies**, rather than national damage totals, as specified in the analytic framework, rather than national damage aggregates..

3

Assessment Based on Current Public Data of the Russia–Ukraine War

On the Russian side, deep-strike effects are proxied using energy-sector damage estimates derived from satellite imagery and OSINT-based assessments combining RFE/RL reporting with Frontelligence Insight analyses. These estimates capture **order-of-magnitude impacts** on refining and fuel infrastructure (e.g., facility damage and operational downtime) and are conservatively aggregated into annual proxy values. For illustrative purposes, recent assessments suggest losses on the order of **tens of billions of rubles** over multi-month periods, recognizing that these figures represent partial sectoral coverage rather than comprehensive national loss. **These figures should be interpreted strictly as order-of-magnitude sectoral impact proxies**, reflecting partial operational disruption and repair costs rather than comprehensive asset valuation or national economic loss.

On the Ukrainian side, critical-asset impacts are proxied using indicators such as electrical power disruptions (e.g., large-scale blackout duration affecting major population centers) and energy-system functionality, as reported by Ukrainian authorities, international organizations, and corroborated media investigations. These measures are selected because they directly reflect **mission-relevant outcomes**, continuity of power and fuel supply, rather than aggregate monetary damage.

3.4 Attack Intensity Proxies

Attack intensity proxies are used solely to approximate pressure, not lethality or realized damage.

Attack intensity is proxied using publicly reported counts of drone activity, with explicit recognition of definitional heterogeneity across sources.

Key illustrative proxies include:

- **Russian Shahed/Geran campaign intensity:** Reporting by the Institute for Science and International Security (ISIS), including figures such as **15,011 launches between 1 August 2024 and 1 March 2025**, which capture the scale and tempo of sustained saturation attacks.
- **Ukrainian deep strikes on Russian energy assets:** Compilations by **Reuters** and other OSINT organizations indicating **at least 81 strikes in 2024 and no fewer than 58 attacks since early August 2025**, depending on the temporal and definitional scope applied.

Because public counts may include decoys, reconnaissance platforms, and mixed strike waves, the analysis distinguishes between **total reported aerial objects** (N_{total}) and **effective attack platforms** (N_{attack}). The latter is estimated through sensitivity parameters reflecting plausible decoy proportions. This prevents overstatement of effective attack pressure and preserves cross-year comparability.

Differences in reporting definitions (e.g., launches, strikes, inbound objects) are **explicitly handled through the dual-track $N_{\text{total}}/N_{\text{attack}}$ framework and sensitivity analysis on decoy proportions**, rather than assumed away.

3

Assessment Based on Current Public Data of the Russia–Ukraine War

3.5 Interpretation of KAPS

The **KAPS** reframes counter-drone effectiveness in terms of **mission outcomes** rather than interception counts. Instead of asking how many drones were shot down, KAPS evaluates whether critical national functions—specifically electrical power continuity and fuel supply—were preserved under sustained attack.

Under this framework, a defense posture may remain strategically effective even with incomplete interception, provided that power grids continue to operate and fuel systems avoid prolonged disruption. Conversely, high nominal intercept rates that nonetheless coincide with extended blackouts or refinery downtime yield low KAPS values, signaling strategic vulnerability irrespective of tactical performance.

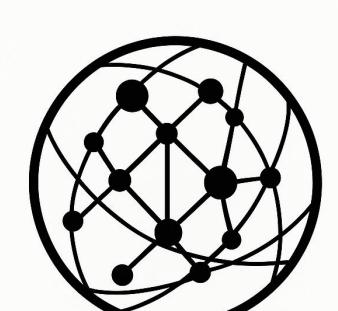
This distinction is central to understanding modern drone warfare. Saturation campaigns are designed not merely to penetrate defenses, but to impose cumulative functional degradation. KAPS therefore serves as a bridge between cost-based indicators (CER, CPLA) and strategic resilience, ensuring that economic efficiency is evaluated alongside the preservation of core societal and military capabilities.

3.6 Scope and Caution

All figures presented in this section are **illustrative and conservative by design**. They are intended to support comparative analysis, sensitivity testing, and policy discussion, not to establish precise causal attribution or definitive loss accounting. Where 2025 data remain provisional or incomplete, values are explicitly marked and treated as scenario placeholders to preserve methodological continuity without overstating confidence.

While this section does not yet present indicator results, the structure of the proxies defined here is designed to detect **structural stress and sustainability erosion** before degradation becomes visible in interception statistics or tactical success metrics alone.

Baseline Results (2023–2025)



4

Baseline Results (2023–2025)

This section reports baseline observational and counterfactual results for the 2023–2025 period under a unified MVA framework. Results are presented to illustrate structural trends in cost sustainability, loss suppression, and mission outcomes under sustained drone saturation, rather than to provide precise accounting of drone-only losses or defense expenditures.

Unless otherwise stated, all results are derived under a consistent set of illustrative assumptions applied symmetrically across years and sides to enable trend comparison.

4.1 Core Unified Assumptions

Baseline results are generated under a common set of assumptions regarding attack intensity, decoy composition, force structure, and unit costs.

a) Attack intensity (N)

Illustrative baseline ranges for inbound UAV activity are constructed from multiple public trackers and official statements and are intended to capture order-of-magnitude pressure rather than exact counts.

Ukraine vs. Shahed-family UAVs

- 2023: 3,500–4,000 inbound UAVs
- 2024: 10,000–11,000 inbound UAVs
- 2025: 38,000–45,000 inbound UAVs

The baseline decoy share is set at $d = 0.30$, yielding:

$$N_{attack} = (1 - d) \cdot N_{total} = 0.7 \cdot N_{total}$$

b) Defensive force structure.

The baseline configuration adopts **S2 (Balanced)** force structure:

Method	Share
SAM	30%
AAA	30%
EW	25%
Interceptor drones	15%

c) Unit cost assumptions (public mid-range estimates).

- Attack UAV unit cost u_{atk} : USD 25,000
- Defensive weighted-average unit cost u_{def} : USD 35,000–45,000 (implied by S2 mix)

These assumptions are held constant across years to isolate structural effects driven by scale, saturation, and force-structure choice.

4

Baseline Results (2023–2025)

4.2 Observational Sustainability and Mission Outcomes

Table 1A (Appendix B) reports baseline observational indicators of counter-drone sustainability and mission outcomes for 2023–2025. Results are expressed as bounded ranges reflecting limited sensitivity over decoy share, attack unit cost, and defensive unit-cost mix.

Several patterns emerge. Figures 1a–1c visualize the year-by-year evolution of the five observational indicators (CER, OLER, $\text{CER}^*_{\text{obs}}$, CPLA_{obs}, and KAPS) for Ukraine and Russia in 2023–2025.

First, sustainability deterioration precedes visible intercept degradation. Across both sides, indicators of marginal efficiency and cost pressure—most notably CPLA and the composite $\text{CER}^*_{\text{obs}}$ —deteriorate consistently before any collapse in loss suppression or mission outcomes becomes visible. This confirms that intercept or shoot-down rates alone are lagging indicators under saturation conditions.

Second, rising attack scale dominates marginal cost dynamics. Between 2023 and 2025, effective attack volume increases by an order of magnitude. Even where OLER remains broadly stable, the marginal cost of preventing additional functional loss rises sharply. By 2025, CPLA values increase nonlinearly, reflecting the interaction between saturation pressure and growing reliance on higher-cost interception layers to preserve acceptable loss levels.

Third, mission outcomes remain intact despite mounting sustainability stress. Throughout the period, KAPS remains above collapse thresholds on both sides. For Ukraine, this indicates continued preservation of core functions—most notably electricity and energy supply—even under sustained Shahed saturation. For Russia, deep-strike effects on energy infrastructure impose measurable but bounded degradation. These results demonstrate that operational effectiveness can persist alongside accelerating economic strain.

Taken together, Table 1A shows that defenses can remain tactically effective and mission-functional while drifting toward strategic unsustainability, a divergence that intercept-centric metrics fail to capture.

Clarification on attack-intensity asymmetry

The markedly lower values of N_{total} and N_{attack} on the Russian side reflect fundamental differences in attack mode rather than data inconsistency. Ukrainian deep strikes against Russian energy and refining infrastructure are characterized by low-frequency, high-value, node-targeted operations, for which publicly auditable proxies capture confirmed strike events rather than total drone sorties. Accordingly, N is not numerically symmetric across sides and is not intended to be compared on an absolute scale. Cross-side comparison in Table 1A is therefore mechanism- and trend-focused, with outcome indicators (OLER, CPLA, KAPS) absorbing differences in value density and functional impact.

4

Baseline Results (2023–2025)

Figure 1a. Observational Indicators (2023)

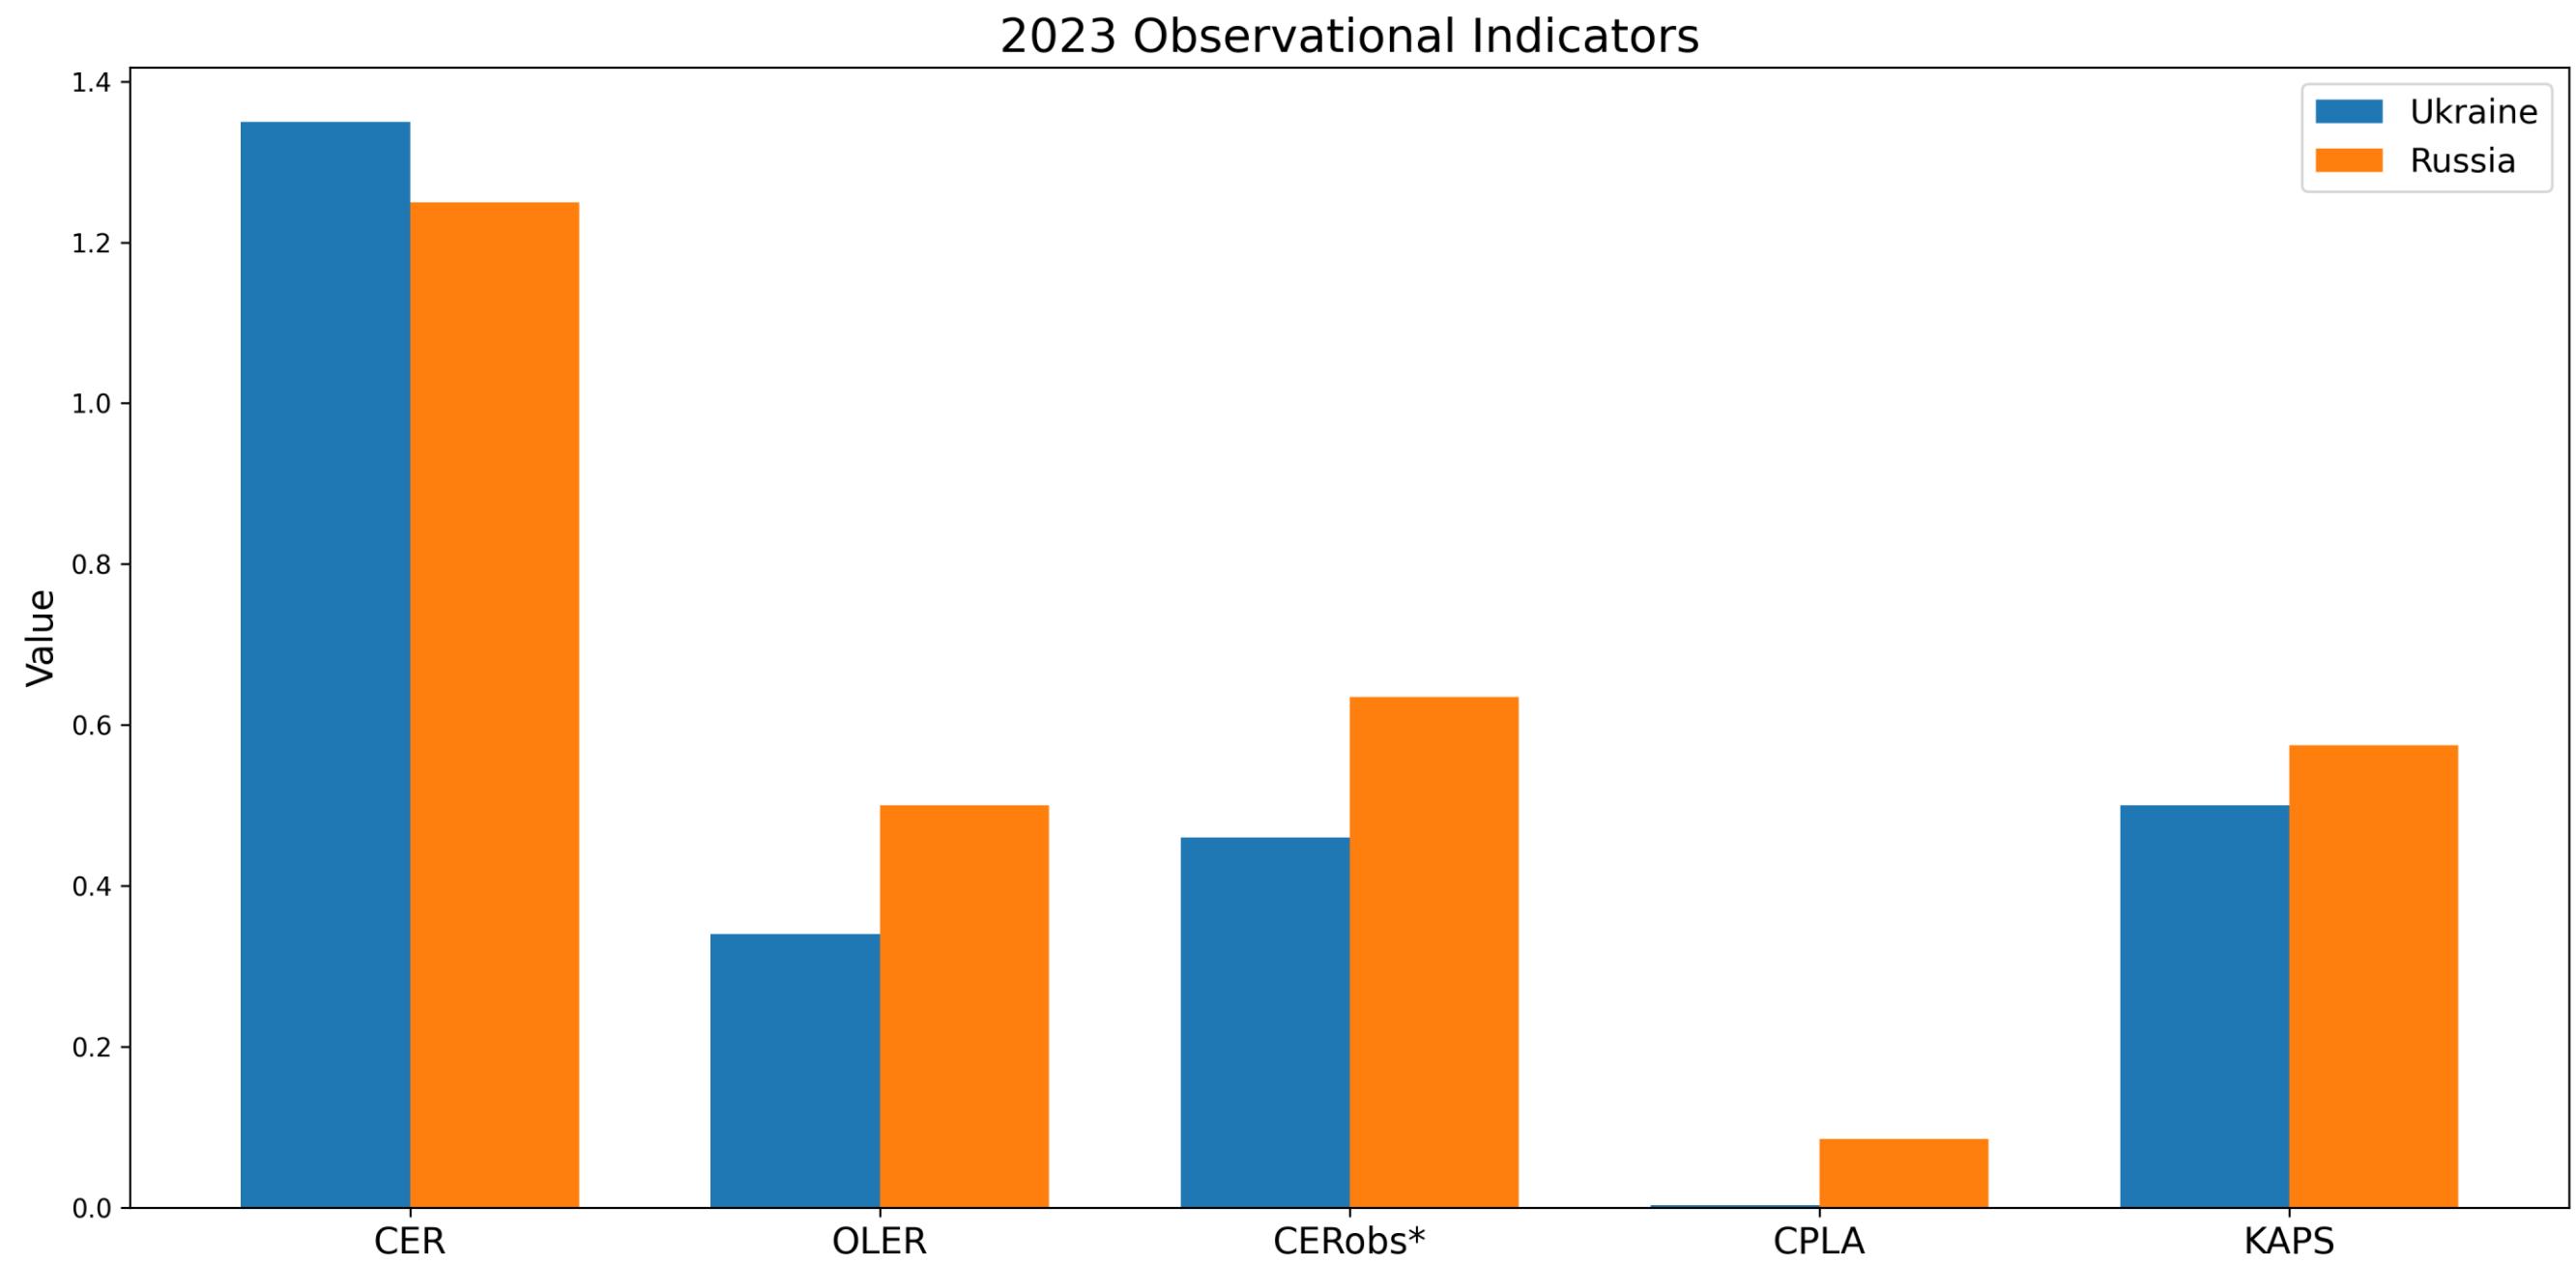


Figure 1b. Observational Indicators (2024)

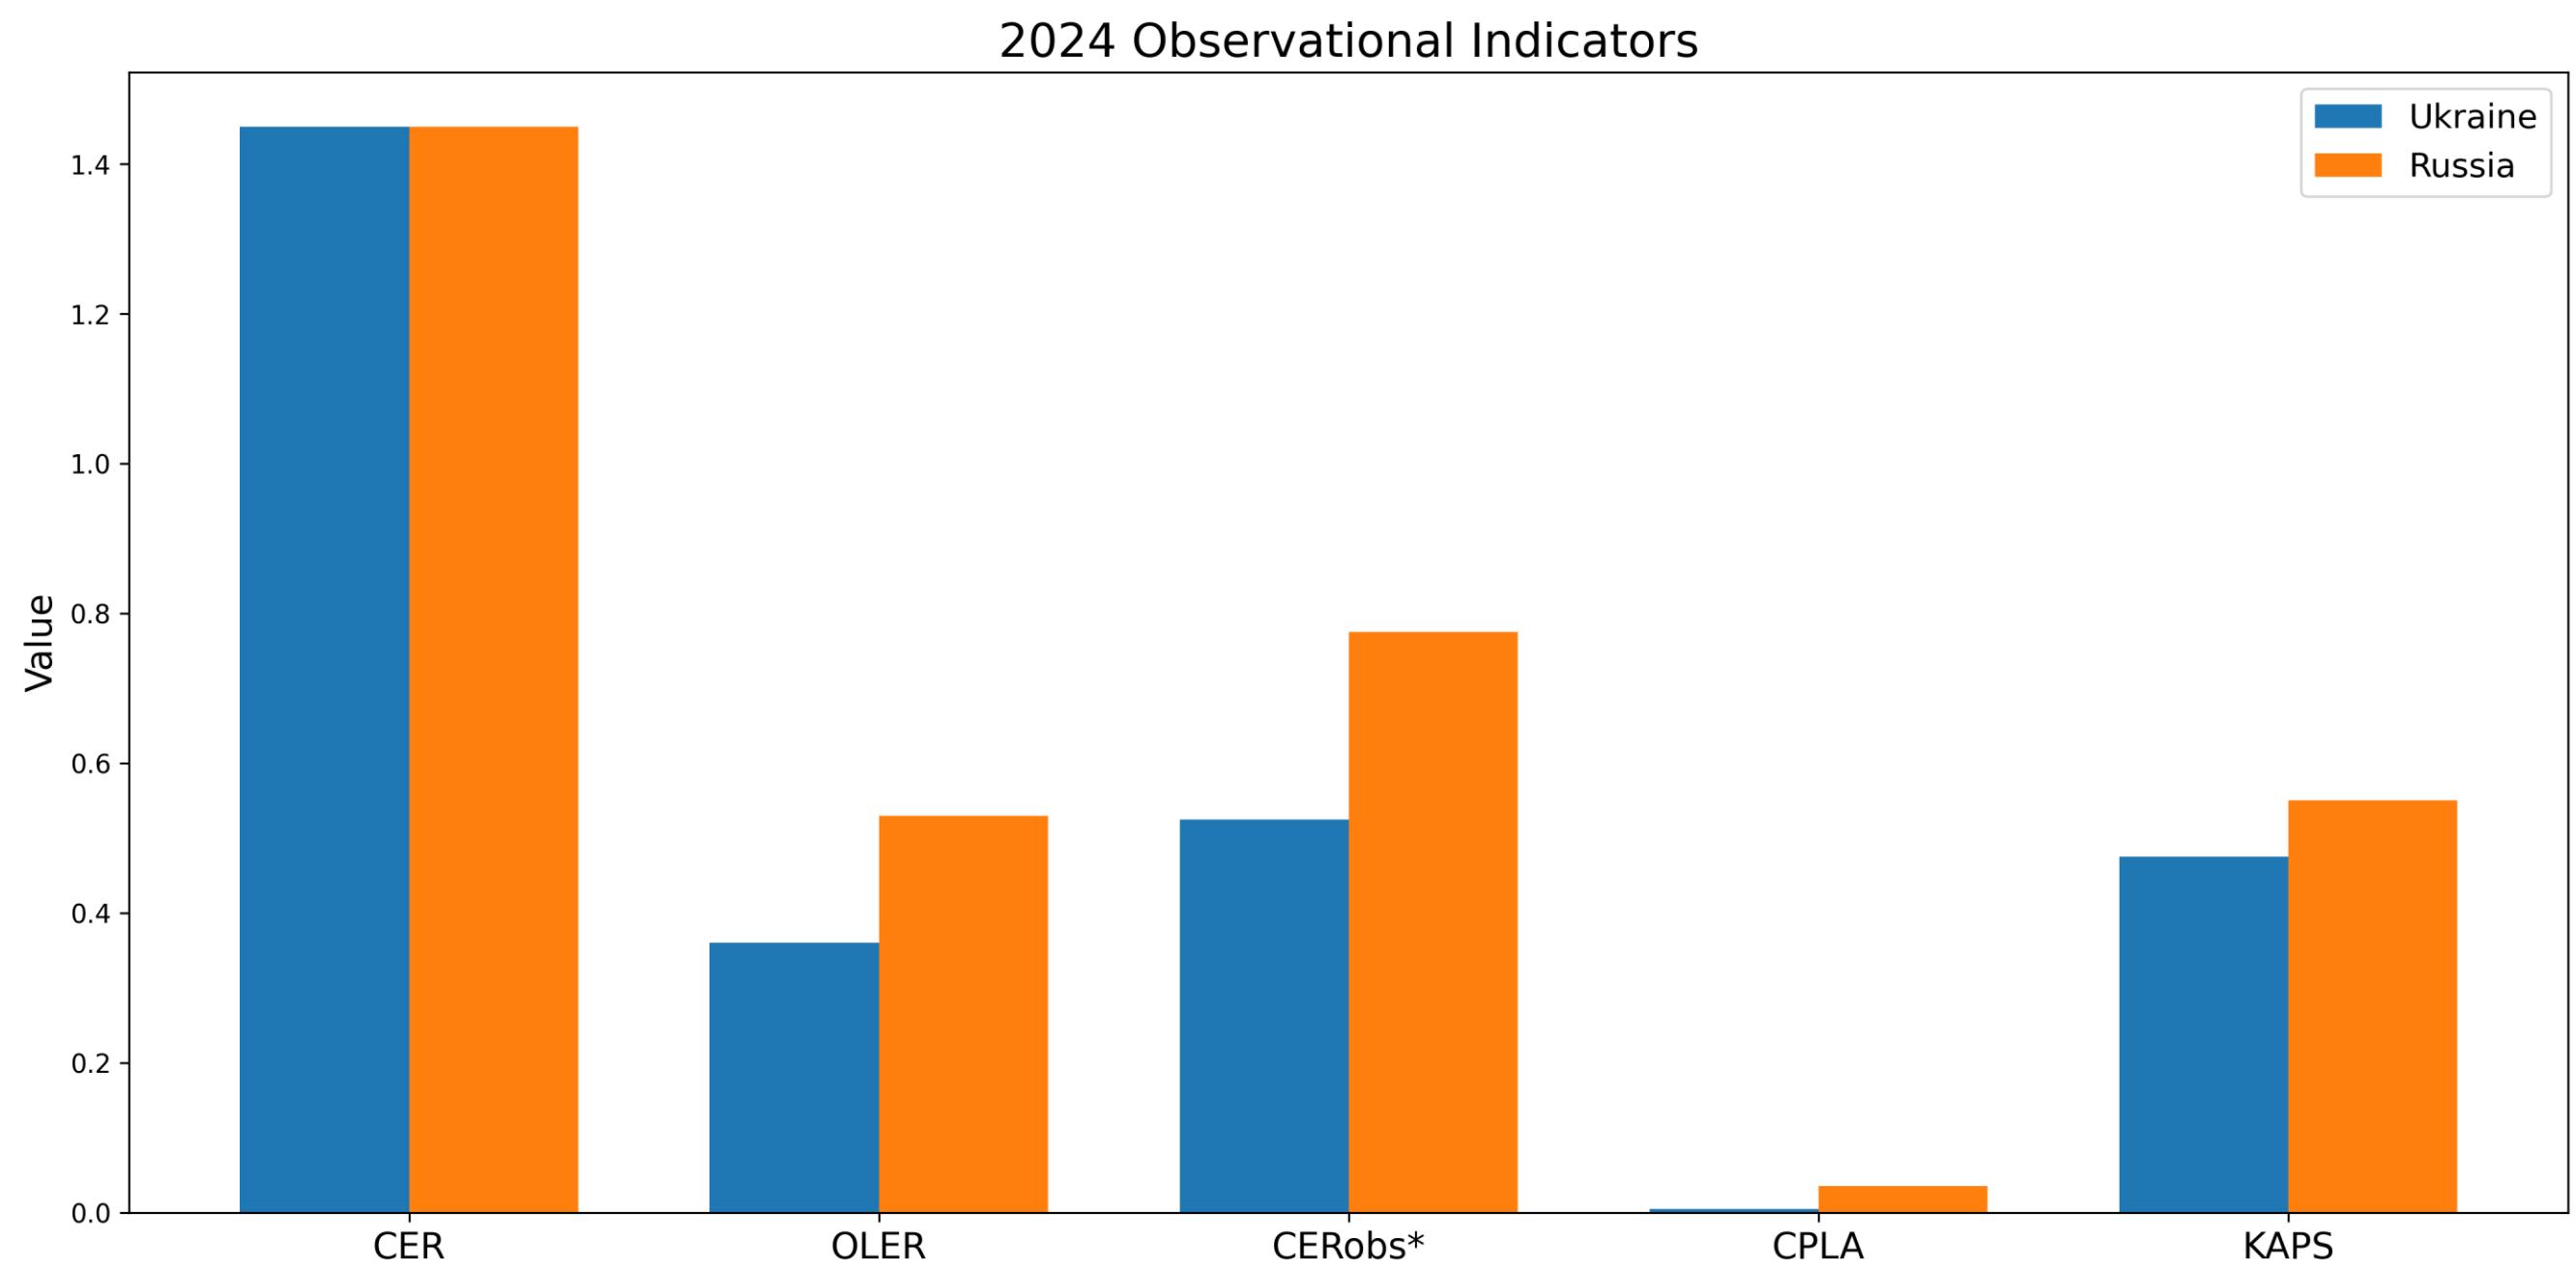
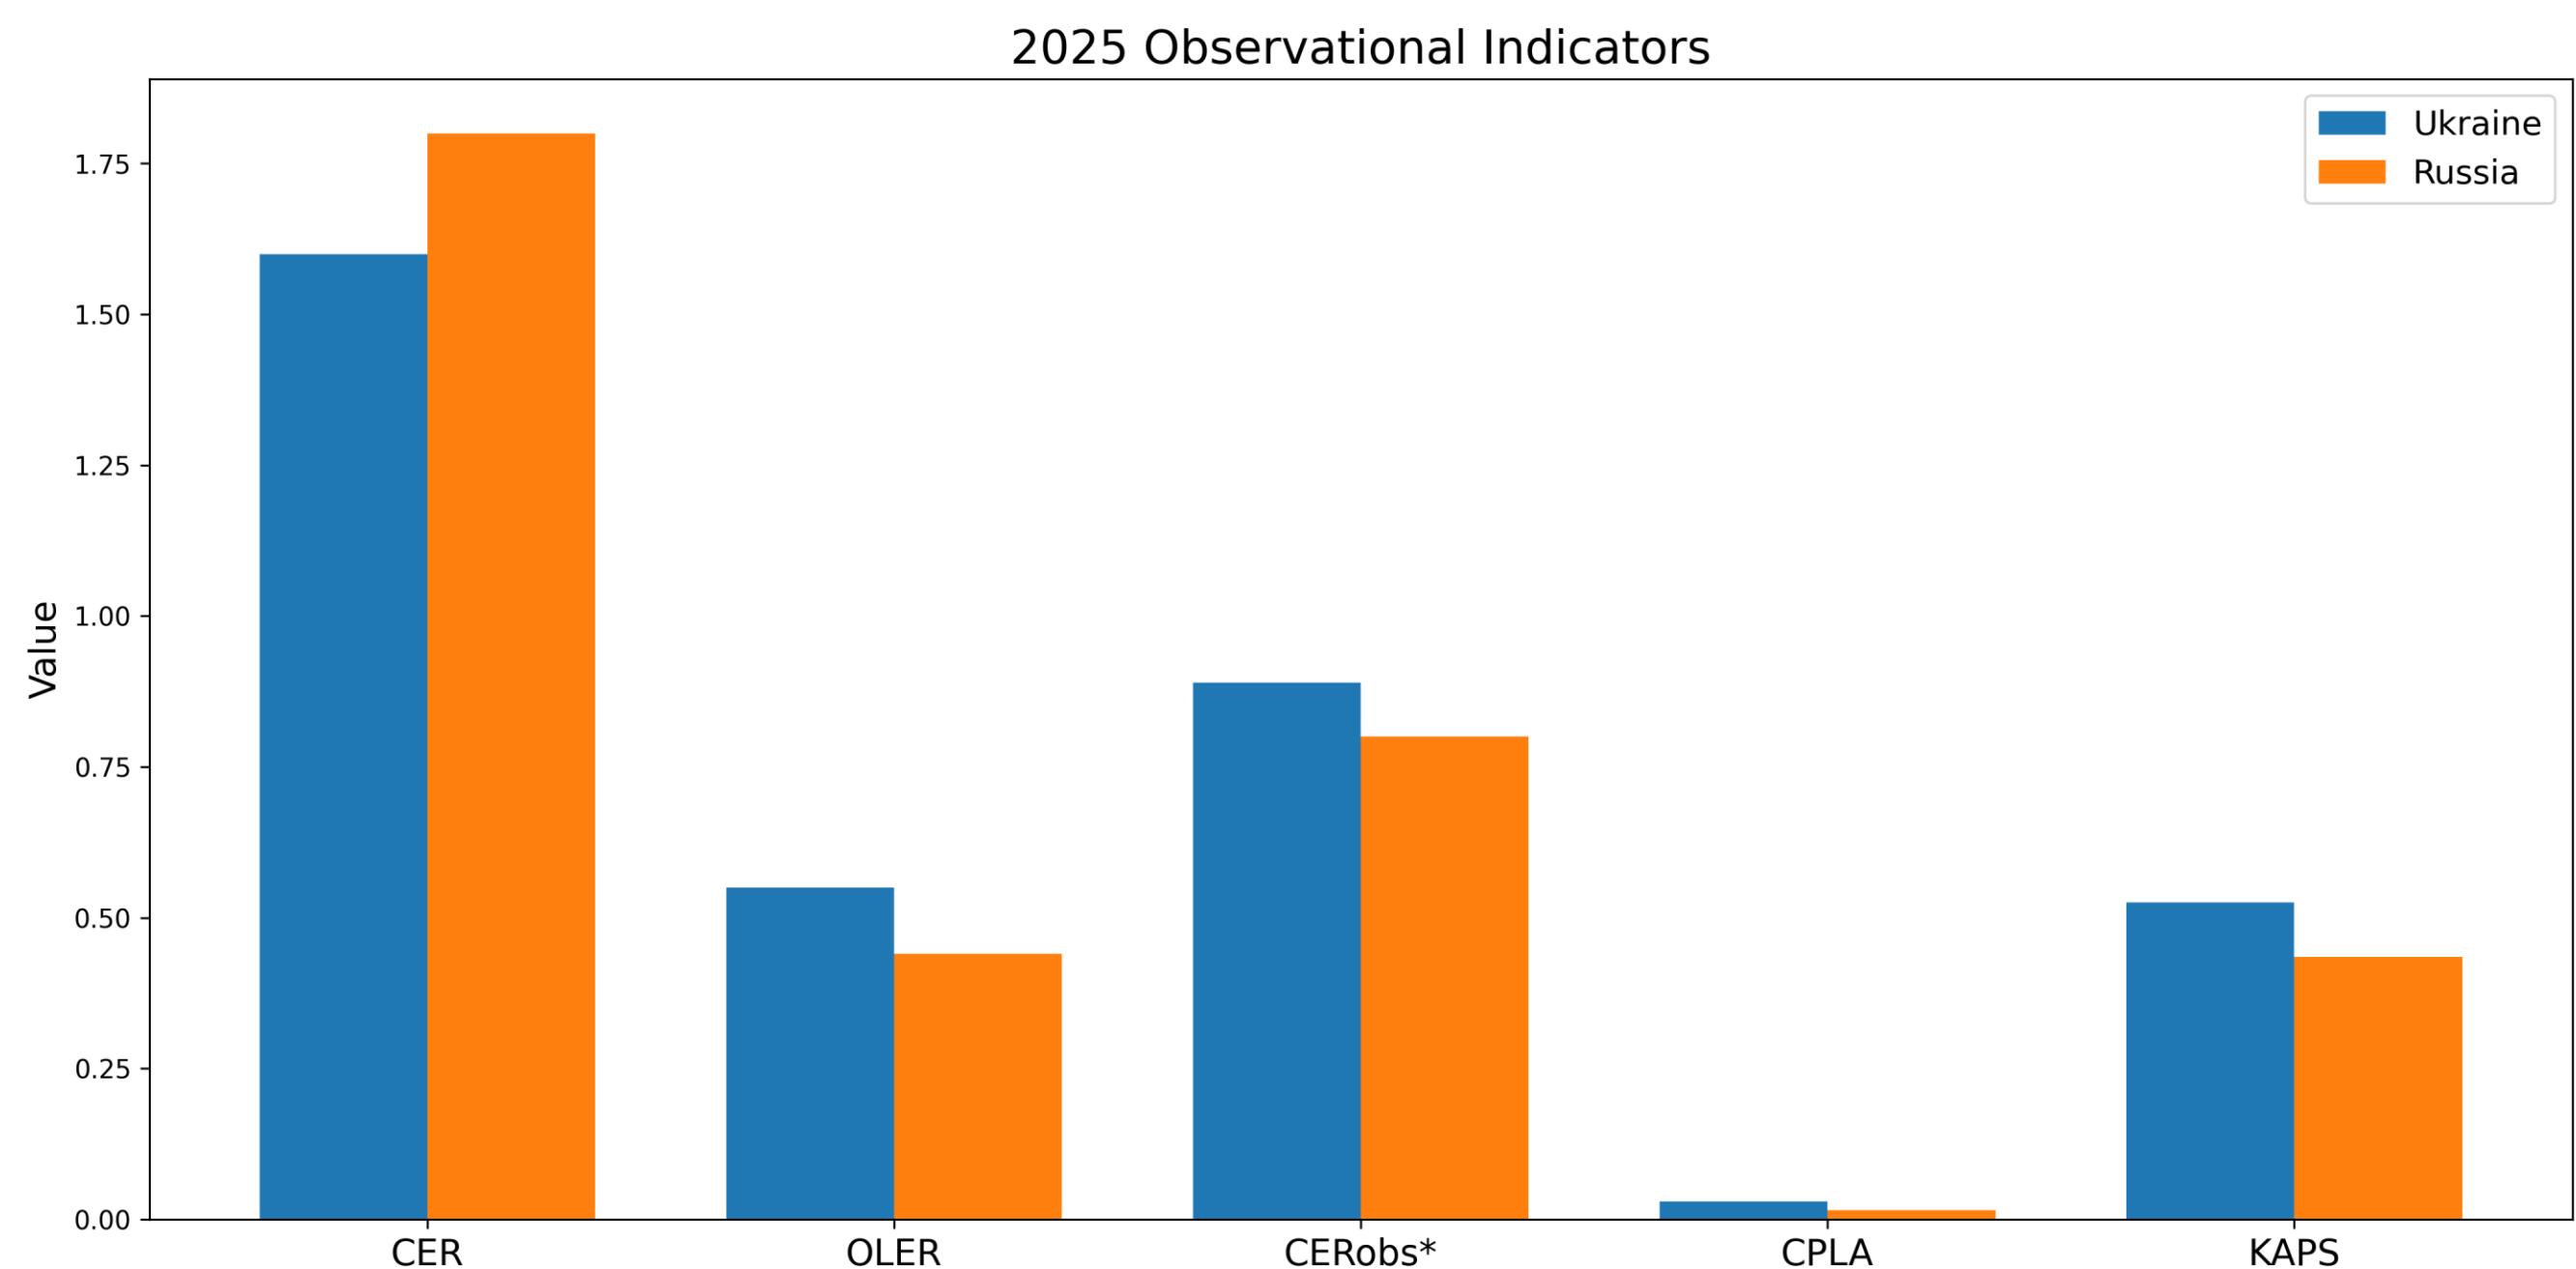


Figure 1c. Observational Indicators (2025)



4

Baseline Results (2023–2025)

4.3 Counterfactual Stress-Test Results (2025)

Table 1B reports counterfactual stress-test results for 2025 under a set of explicitly defined loss-amplification scenarios. All results in this subsection are scenario-based and do not represent observed outcomes.

Figure 2 visualizes these results as a sustainability decision frontier, plotting CPLA_{cf} against CER^*_{cf} to illustrate relative trade-offs across force-structure choices under fixed stress assumptions. The frontier highlights relative sustainability trade-offs and should not be interpreted as a normative or universally optimal defense configuration.

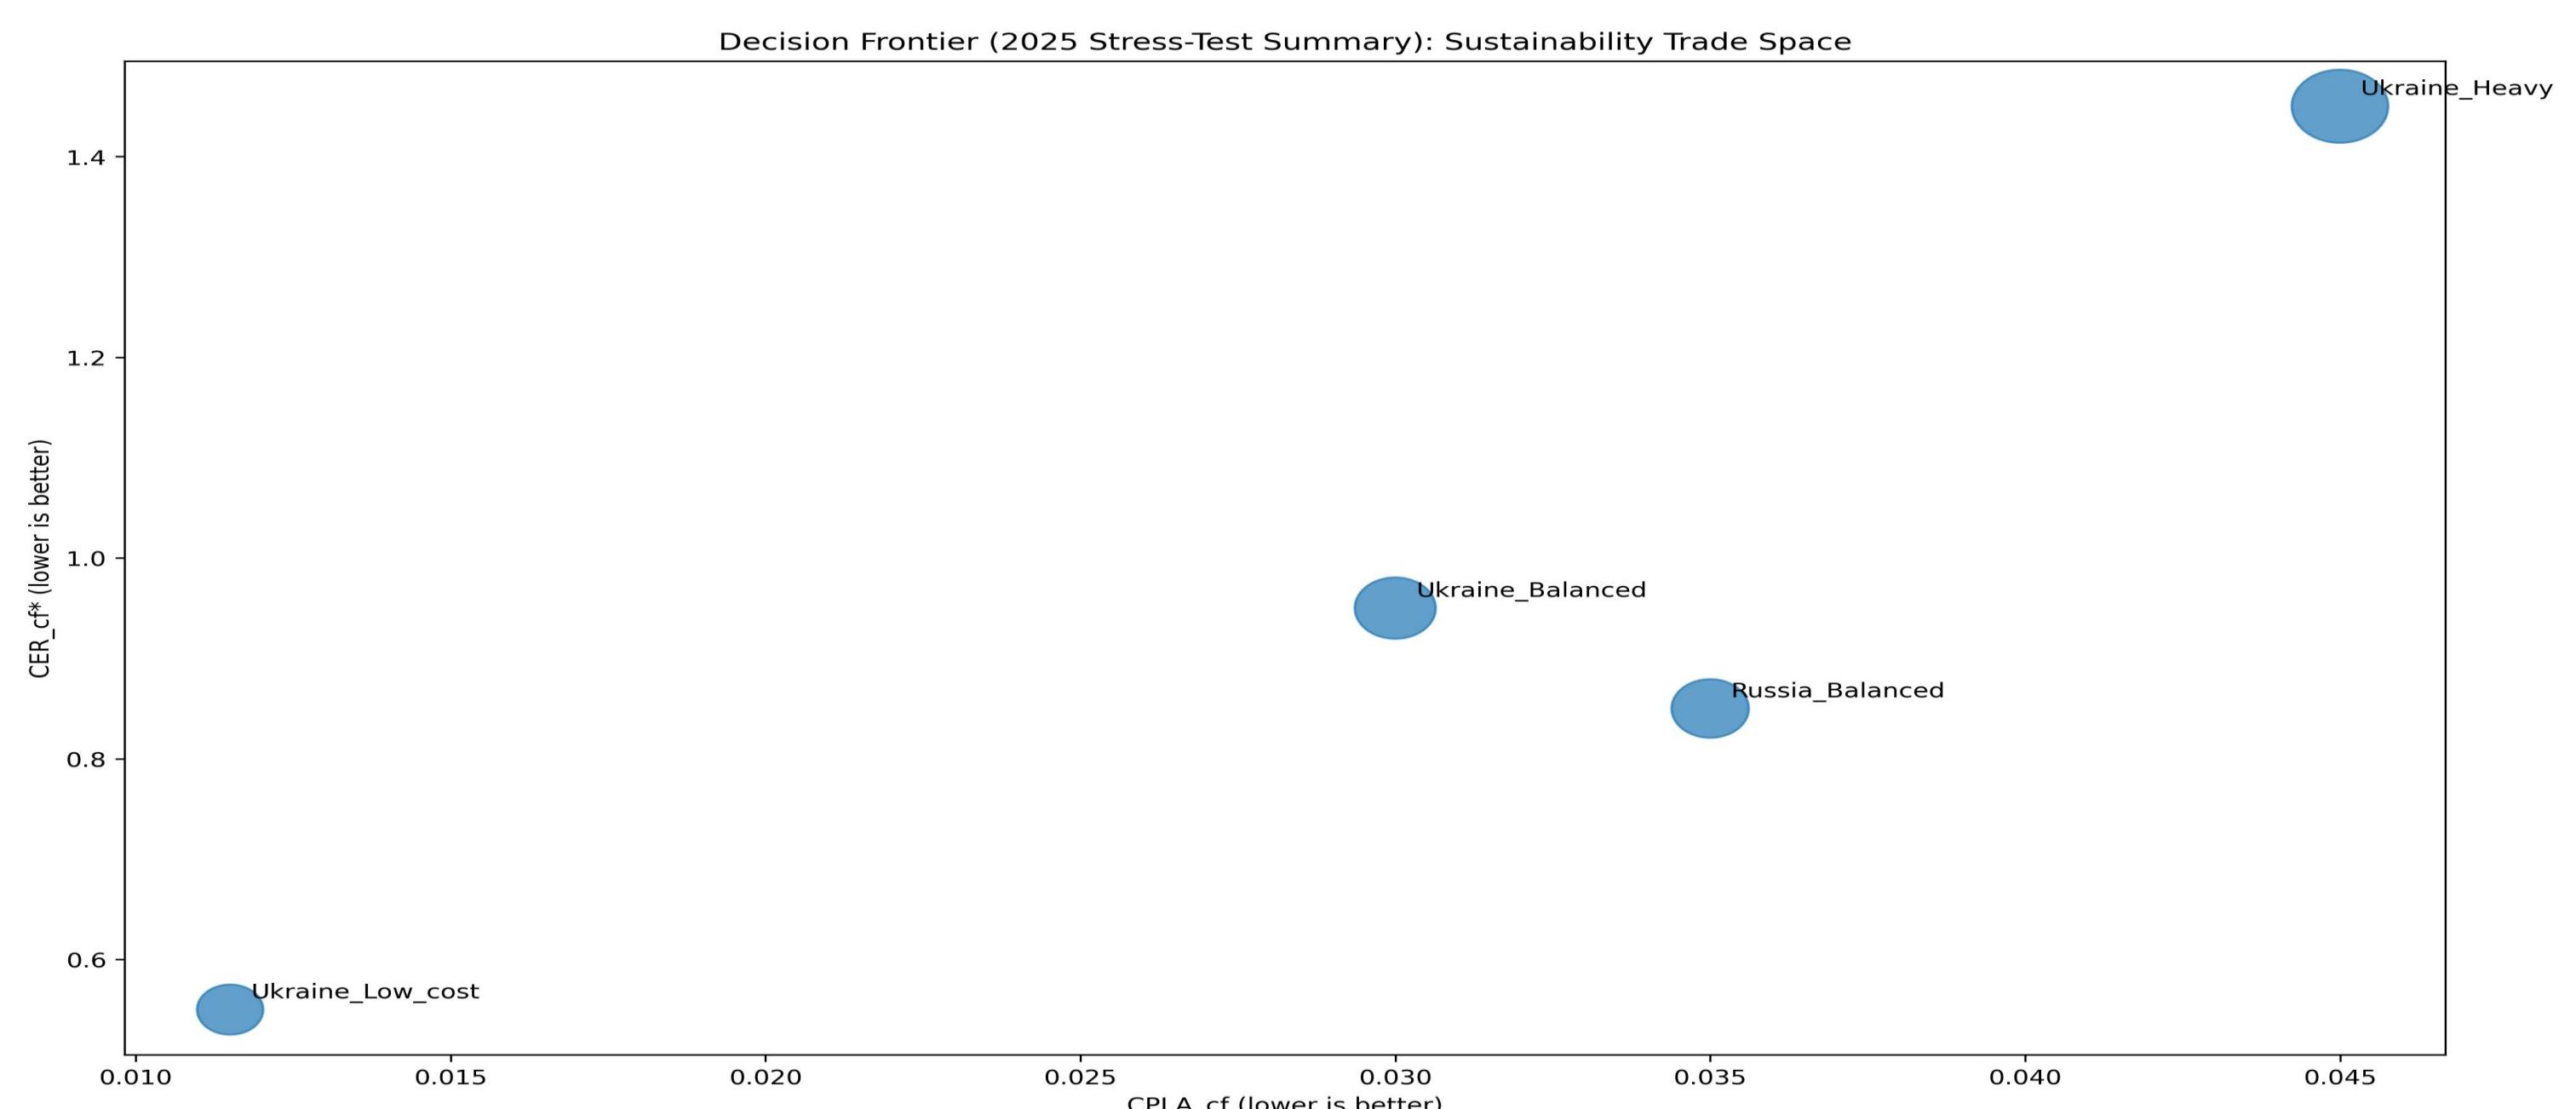
First, force-structure choice dominates sustainability under stress. Across identical loss-amplification scenarios, **low-cost-heavy configurations (S3)** consistently exhibit lower CER^*_{cf} and CPLA_{cf} than SAM-heavy defenses. As assumed loss amplification increases, SAM-heavy structures experience steeply rising marginal costs, while low-cost terminal-layer-heavy defenses maintain flatter cost–loss curves.

Second, economic fragility compounds faster than mission degradation. Even under elevated loss multipliers, differences across force structures manifest primarily in cost sustainability rather than immediate mission failure. This reinforces the central analytic claim: sustainability stress accumulates before mission collapse.

Third, balanced defenses occupy an unstable middle ground. S2 configurations outperform SAM-heavy structures but remain vulnerable to rapid CPLA escalation under stress. This suggests that balanced mixes may be operationally adequate in moderate conditions yet insufficiently robust for prolonged high-saturation environments.

Overall, Table 1B demonstrates that under stress, sustainability outcomes are shaped less by absolute interception capability than by the composition and scalability of the terminal defense layer.

Figure 2. Sustainability Decision Frontier under Drone Saturation: CPLA_{cf} vs. CER^*_{cf}



4

Baseline Results (2023–2025)

4.4 Findings

Three findings are central for policy and force-structure assessment under sustained drone saturation.

a) The 2025 CPLA “jump” reflects structural stress, not accounting error.

The sharp rise in CPLA observed in 2025 is a structural consequence of saturation dynamics. Effective attack volume increases dramatically, while defenders are compelled to rely on progressively more expensive interception methods to preserve acceptable loss suppression. As a result, marginal efficiency deteriorates even when OLER remains relatively stable. This divergence is the quantitative manifestation of the classic “missile-versus-drone” problem: tactical effectiveness can be sustained, but only at accelerating and ultimately unsustainable cost.

b) Low-cost terminal-layer defenses are most robust under stress.

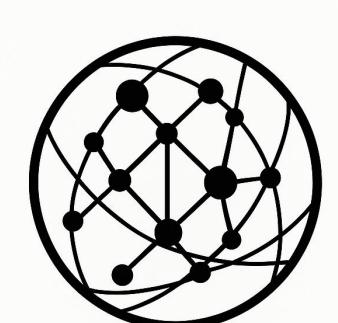
Across all counterfactual scenarios, S3 configurations consistently outperform SAM-heavy defenses on both cost and composite indicators. Crucially, low-cost terminal layers are not merely cost-saving in benign conditions; they are structurally stabilizing under stress. This finding has direct policy relevance for force design under prolonged saturation.

c) KAPS anchors analysis in strategic outcomes.

While CER and CPLA reveal sustainability erosion, **KAPS** confirms that mission-critical functions can remain operational despite incomplete interception. In 2025, Ukraine’s KAPS remains above 0.4, indicating continued power and energy-system functionality. The principal risk, therefore, is not immediate operational failure but long-term viability driven by accelerating resource consumption.

A limited robustness check incorporating peak-salvo effects ($\beta = 0.15$, not shown) further worsens CPLA outcomes under SAM-heavy force structures while preserving the qualitative ranking of S1–S3 configurations. This confirms that the core findings are robust to moderate saturation-intensity adjustments.

Interpretation and Limitations



5

Interpretation and Limitations

5.1 Interpretation of Baseline Results (2023–2025)

The two cases stress different dimensions of sustainability: volume-driven depletion on the Ukrainian side, and value-density-driven loss amplification on the Russian side.

a) “Shooting bullets at mosquitoes”: why sustainability precedes intercept rates

Under conditions of sustained drone saturation, intercept or shoot-down rates are a fundamentally misleading performance metric when used in isolation. A defending force can preserve high nominal interception rates in the short term by relying on **faster, scarcer, and more expensive interceptors**, but this comes at the cost of rising defensive expenditures (C_{def}), declining magazine depth, and increased operational and maintenance burdens. In such cases, the problem shifts rapidly from the tactical to the strategic level.

Accordingly, this report defines “effectiveness” as a **joint condition**: maintaining acceptable loss suppression (**OLER**) and mission outcomes (**KAPS**) **within tolerable cost bounds** (**CER** and **CPLA**). From this perspective, **CPLA** and **CER_{obs}*** often reveal sustainability limits **earlier than intercept rates**:

- Rising **CPLA** indicates that the marginal cost required to avoid one additional unit of critical-asset functional loss is increasing rapidly.
- Rising **CER_{obs}*** indicates that cost-exchange pressure and residual functional loss are deteriorating simultaneously—even when interception performance appears stable.

This divergence explains why defenses can appear tactically successful while moving steadily toward economic and logistical exhaustion.

b) Structural signals from the Ukrainian case: scale-driven saturation and the dominance of unit cost curves

The observational baseline (Table 1A) shows that Ukraine’s attack intensity (**N_{total}** and **N_{attack}**) increased by orders of magnitude from 2023 to 2025*, pushing the air-defense system into a **high-frequency, high-density, high-consumption steady-state pressure regime**. In this regime, long-run outcomes are no longer determined by the technical ceiling of individual engagements, but by whether the defense structure can shift the bulk of engagements to **low-cost terminal layers**.

- Where low-cost terminal defenses (AAA, EW, interceptor drones) are insufficient, defenders are forced to expand the share of expensive SAM interceptions, causing **CER and CPLA to deteriorate rapidly** and creating a classic cost-exchange trap.
- Even when OLER does not collapse, rising CPLA signals that comparable levels of functional suppression are being maintained at **increasing marginal cost**, which over time translates into inventory constraints and readiness fragility.

The central structural implication is therefore not whether interception can continue, but whether critical functions can be preserved at **declining marginal cost**, which is why KAPS and CPLA are treated as co-equal analytic anchors.

5

Interpretation and Limitations

c) Structural signals from the Russian case: deep strike and the value density of functional loss

The Russian case—proxied through energy and refining infrastructure—illustrates a different mechanism. Ukrainian deep strikes are often directed at **high-value nodes and bottlenecks**, where a single successful hit can generate outsized functional spillovers through production losses, supply disruptions, repair cycles, and secondary effects.

As a result, even with N_{attack} far below the scale of Russia's Shahed campaign, deep strikes can impose **nonlinear pressure** at the critical-asset level. In this context:

- **KAPS** provides a closer approximation to strategic effect than interception counts.
- Deterioration in **OLER** reflects cumulative functional erosion rather than isolated damage events.
- **CER** remains relevant, but its policy implications point more directly toward **node hardening, redundancy, and recovery capacity**, rather than interceptor substitution alone.

This pattern reflects the high value density of targeted assets rather than defensive failure per se.

d) Joint interpretation rules: avoiding single-indicator fallacies

Counter-drone performance should be interpreted through a—cost, suppression, and outcome—rather than through any single indicator: **three-axis framework**

- **CER** answers whether defense is economically asymmetric relative to attack.
- **OLER** captures observed suppression of critical-asset functional loss under comparable pressure.
- **$\text{CER}_{\text{obs}}^* = \text{CER} \cdot \text{OLER}$** indicates whether cost pressure and residual loss are deteriorating simultaneously.
- **CPLA_{obs}** measures the marginal cost of avoiding one unit of functional loss and is the most direct resource-allocation signal.
- **KAPS** anchors the analysis to strategic outcomes by indicating whether core functions are preserved, thereby avoiding “intercept-rate illusions.”

Policy conclusions should therefore be drawn from **consistent signals across indicators**. For example, rising CPLA and **$\text{CER}_{\text{obs}}^*$** —even with stable OLER—typically indicate that the system is approaching a sustainability boundary. Conversely, controlled CER, improving OLER, and stable or rising KAPS suggest that structural adjustments are producing genuine strategic gains.

5

Interpretation and Limitations

5.2 Practical Use Guidance

Recommended use cases

- **Cross-year trend detection:** prioritize systematic movements in CER_{obs}^* and $CPLA_{obs}$ to identify approaching cost-inventory saturation points.
- **Force-structure comparison:** run identical annual data under S1/S2/S3 to assess marginal CPLA improvements and verify whether KAPS is preserved or improved.
- **Critical-asset prioritization:** disaggregate KAPS by asset class (e.g., electricity vs. refining) to support budget-to-target matching and resilience planning.

Not recommended (common analytic traps)

- Do not base conclusions on a single CER or intercept rate.
- Do not interpret RDNA as drone-attributable loss.
- Do not ignore heterogeneity in N or decoy proportions.
- Do not treat counterfactual outputs (CLER, $CPLA_{cf}$) as observed reality.

5.3 Key Caveats: What the Baseline Can and Cannot Claim

a) Attribution limits: RDNA and energy-loss proxies are not drone-exclusive

- Ukraine's RDNA estimates reflect comprehensive national damage from multiple sources and are used solely to provide a stable macro-level context, not to attribute losses to drone operations.
- Russian energy-loss proxies cover only selected sectors and facilities and represent a conservative window into deep-strike effects rather than total national loss.

b) Observational vs. counterfactual boundaries: CLER must not substitute for OLER

Loss suppression is explicitly separated into:

- **OLER (observational layer)**, derived from critical-asset functional loss proxies; and
- **CLER (counterfactual layer)**, determined by scenario multiplier MMM and used exclusively for stress testing.

Accordingly, CLER—and derived indicators such as CER_{cf}^* and $CPLA_{cf}$ —must not be interpreted as real-world measurements. Their correct use is **comparative**, under a fixed **M**, across defense structures and years.

c) Cost uncertainty: CER and CPLA are highly structure-sensitive

Defensive cost depends not only on how many engagements occur, but on **how they are handled** (SAM vs. AAA vs. EW vs. interceptor drones). Attack costs likewise vary with platform

5

Interpretation and Limitations

type, domestic substitution, decoy share, and scale effects. In the absence of fully auditable force-structure data, CER should be interpreted as a **directional signal**, not a precise exchange rate. This is why S1/S2/S3 scenarios are used to explicitly expose structural sensitivity.

d) Heterogeneity in N: distinguishing N_{total} from N_{attack} is essential

Public counts often aggregate attack drones, decoys, reconnaissance platforms, and heterogeneous payloads. Treating N as an intensity proxy is necessary and reasonable, but only when combined with the dual-track $N_{\text{total}}/N_{\text{attack}}$ framework and sensitivity analysis over decoy share **d**.

e) Temporal lag and data gaps: provisional treatment of 2025 is methodologically required

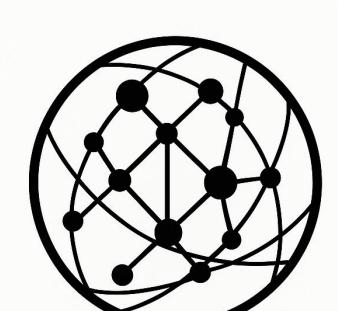
Where authoritative annual assessments are incomplete, 2025 values are reported as ranges and explicitly marked as provisional (* / **). These values preserve model continuity and trend analysis but do not constitute final factual judgments.

5.4 Roadmap for Refinement

To move from a comparability-first MVA baseline toward higher-confidence estimates suitable for stronger policy claims, three data improvements should be prioritized:

- **Auditable force-structure data:** annual shares and cost ranges for SAM, AAA, EW, and interceptor-drone engagements.
- **Functional loss and recovery curves:** outage duration, repair cycles, substitution costs, and cascading effects to upgrade OLER and KAPS from point estimates to resilience trajectories.
- **Wave- and target-type-disaggregated attack data:** separating decoys from attack payloads and **tactical** from deep strikes to reduce N heterogeneity and improve causal inference.

Policy Options



GLOBAL AI
GOVERNANCE
RESEARCH CENTER

6

Policy Options

This section translates the analytic findings into executable policy options. Each option is framed in terms of **mechanism, expected indicator movement, key risks, and enabling conditions**, allowing decision-makers to evaluate trade-offs under sustained saturation conditions rather than optimizing for intercept or shoot-down rates alone.

6.1 Mission-Based Defense: Prioritize Critical Assets and Resilience (KAPS-First)

a) Mechanism

Accept incomplete interception in noncritical areas while concentrating protection, hardening, redundancy, and rapid recovery capacity on power and fuel infrastructure. Emphasize bypass routes, spare parts, and repair speed.

b) Expected indicator movement

- **KAPS** ↑ even if intercept rates and **OLER** do not materially improve.
- **CPLA** ↓ if resilience investments avert large functional losses at moderate cost.

c) Key risk

Political and psychological costs of tolerated leakage; adversary adaptation toward newly prioritized nodes.

d) Enablers

Asset criticality mapping, prepositioned spares, rapid repair logistics, and clear public-communication strategies to manage expectations.

6.2 Rebalance the intercept mix toward low-cost terminal layers (S2 → S3).

a) Mechanism

Shift a larger share of engagements against low-speed, low-RCS drones from high-end SAM systems to lower-cost terminal layers, including AAA, EW, and interceptor drones. Reserve expensive interceptors primarily for high-speed or high-value threats.

b) Expected indicator movement

- **CER** ↓, **CPLA** ↓ due to lower marginal engagement costs.
- **CER^{*}_{obs}** ↓ even if **OLER** remains flat, reflecting improved sustainability.
- **KAPS** stable or modestly ↑ if critical assets remain protected.

c) Key risk

S3 dominance is conditional on sufficient detection and cueing performance; absent this, leakage risks increase. Leakage may increase if detection, classification, and track continuity are insufficient to cue low-cost effectors in time.

6

Policy Options

d) Enablers

Dense low-altitude sensing, sensor fusion and cueing networks, and standardized engagement authorities enabling rapid use of low-cost effectors without SAM-level authorization delays.

6.3 Invest in Detection and Track Continuity (“Seen” → “Engageable”)

a) Mechanism

Increase low-altitude sensor density and improve track continuity in cluttered environments; reduce sensor-to-shooter latency to expand usable engagement windows.

b) Expected indicator movement

- Primary improvement in **OLER** ↓ and **KAPS** ↑.
- **CPLA** ↓ only if incremental suppression gains outweigh added sensing and integration costs.

c) Key risk

Diminishing returns in heavy EW environments and rising system-integration complexity.

d) Enablers

Common data standards, rapid re-tasking procedures, resilient sensor fusion architectures, and explicit EW deconfliction rules.

6.4 Magazine Depth and Sustainment Strategy (Endurance as a Capability)

a) Mechanism

Treat counter-UAS not as a finite interceptor problem but as an endurance portfolio encompassing stockpiles, production throughput, maintenance cycles, training pipelines, and replacement timelines.

b) Expected indicator movement

- Stabilizes C_{def} volatility and reduces cost spikes during peak months.
- Improves robustness of **CER** and prevents abrupt **CPLA** jumps under surge conditions.

c) Key risk

Budget competition with other priorities and industrial base constraints.

d) Enablers

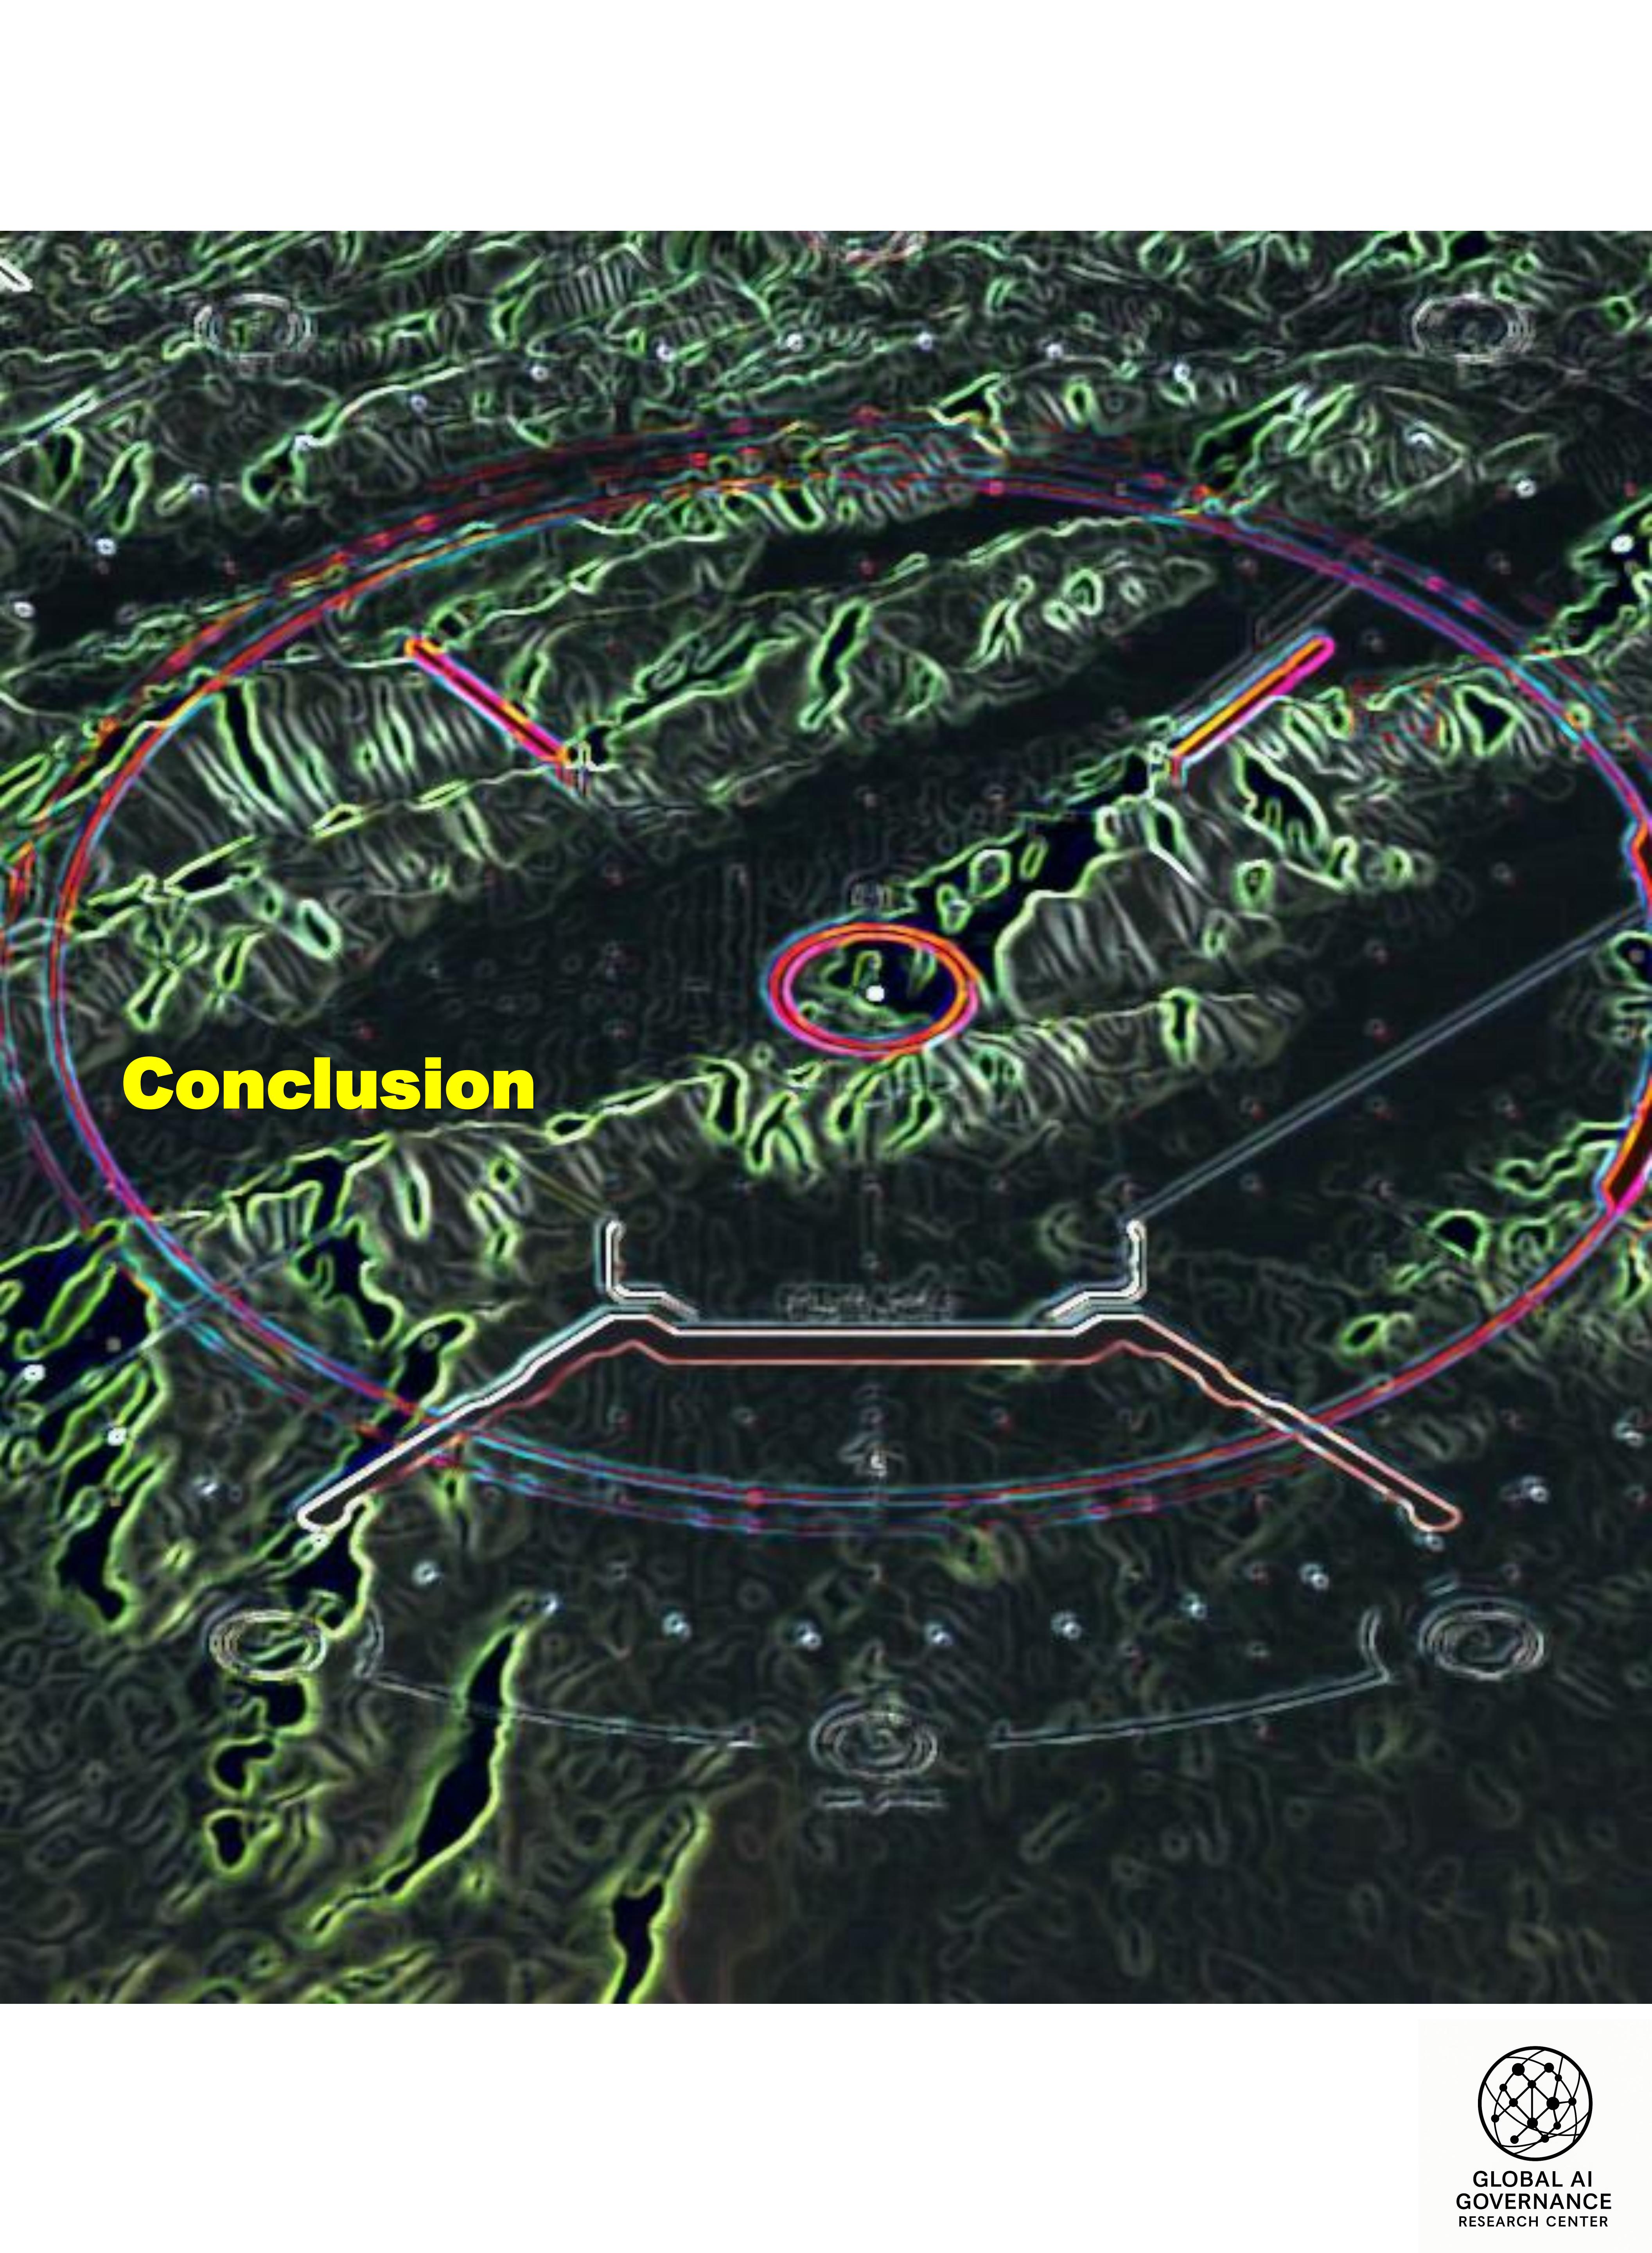
Multi-year procurement authorities, modular system design, diversified suppliers, and surge-production planning.

6 Policy Options

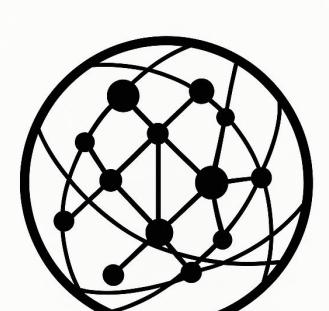
6.5 Recommended Decision Criterion

Select the policy portfolio that **minimizes CPLA** subject to maintaining **KAPS $\geq K_{\min}$** (a defined mission threshold) across a specified stress band for key uncertainty parameters (**d, β , M**).

This criterion operationalizes the central finding of the report: under saturation, **strategic effectiveness is defined by sustainable loss avoidance and mission preservation, not by maximal interception performance**.



Conclusion



Conclusion

By the end of 2025, the Russia–Ukraine war has pushed drone warfare from a tactical adjunct into a **system-level driver of military sustainability**. Low-cost, mass-producible, expendable, and networked unmanned platforms now impose persistent pressure on air-defense systems across three binding constraints: **physical limits** (low-altitude clutter and weak signatures degrading detection and track continuity), **temporal limits** (mismatches among sensor refresh, data fusion, authorization, and interceptor timelines under saturation), and **resource limits** (magazine depth, operational tempo, and fiscal sustainability). Under these conditions, evaluation frameworks centered on intercept or shoot-down rates become increasingly misleading. High intercept performance can coexist with declining sustainability and therefore cannot be treated as a reliable proxy for long-term effectiveness.

By using a **Minimum Viable, Auditable (MVA)** baseline grounded in publicly available data and conservative proxies, this report demonstrates that the central policy challenge in the current conflict is not whether defenses can intercept more threats, but whether they can **sustain an acceptable cost–loss exchange while preserving critical national functions over prolonged attrition**. To address this challenge, the analysis shifts from intercept-centric metrics toward sustainability- and outcome-oriented indicators: **CPLA (Cost per Loss Avoided)** and **CER*** as measures of economic viability, complemented by **KAPR/KAPS** as indicators of mission-level effectiveness. Together, these metrics translate the intuitive “shooting mosquitoes with bullets” dilemma into a structured, comparable, and decision-relevant framework. **CPLA and CER*** **function as early-warning indicators of defense exhaustion, often years before intercept performance visibly degrades**. Crucially, when defenses are forced to rely on high-cost interceptors against low-value threats, **CPLA and CER*** tend to deteriorate well before intercept rates visibly decline, providing an early-warning signal of structural exhaustion.

The broader implication of the Russia–Ukraine drone contest is therefore not the superiority of any single unmanned platform or air-defense system, but a **fundamental shift in the cost structure and governance logic of warfare**. Tactically, effective counter-drone defense increasingly depends on layered architectures centered on low-cost terminal interception, multi-sensor fusion, and high-frequency command-and-control adaptation. Strategically, counter-UAS capability must be treated as an **endurance capability**, the ability to sustain mission-relevant functions under continuous pressure, rather than as a measure of peak interception performance. At the normative and governance level, the expanding use of AI-assisted detection, targeting, and automated engagement continues to strain existing frameworks of responsibility, proportionality, and meaningful human control, shifting these debates from abstract principles toward institutional stress testing under real operational constraints.

The limits of this assessment must also be emphasized. Because the analysis relies on open-source data and proxy variables, 2025 loss and cost estimates remain provisional, and their absolute values should not be interpreted as definitive. Advancing from directional insight toward higher-confidence policy guidance requires progress in three areas: **(1)** auditable data on interceptor mix and unit costs across SAM, AAA, EW, and interceptor drones; **(2)** asset-level

Conclusion

functional loss and recovery curves that capture resilience dynamics rather than point damage; and (3) attack data disaggregated by wave intensity, target type, and decoy composition. Improvements along these dimensions would allow the present MVA framework to support more robust budgetary trade-offs, force-structure decisions, and strategic planning under sustained drone saturation.

While grounded in the Russia–Ukraine case, the analytical framework developed here is intentionally **portable**. Its logic applies to other theaters confronting low-cost saturation threats, including maritime, littoral, and mixed-domain environments, where **sustainability, rather than interception alone, will increasingly define military effectiveness**.

References

- ACLED. (n.d.). *Ukraine Conflict Monitor: Data and methodology*. <https://acleddata.com>
- Atlantic Council. (2025, January 16). *Ukraine's escalating air attacks bring Putin's invasion home to Russia*. <https://www.atlanticcouncil.org/blogs/ukrainealert/ukraines-escalating-air-attacks-bring-putins-invasion-home-to-russia/>
- Australian Army Research Centre. (2025, August 19). *Translating Ukraine lessons for the Pacific Theatre*. <https://researchcentre.army.gov.au/library/occasional-papers/translating-ukraine-lessons-pacific-theatre>
- Business Insider. (2025, November 11). *Ukraine is fighting a war of innovations, and AI may be helping it down Russian drones*. <https://www.businessinsider.com/ukraine-russia-drones-ai-shahed-geran-anti-air-defense-2025-11>
- Center for European Policy Analysis (CEPA). (2023, September 27). *An urgent matter of drones: Lessons for NATO from Ukraine*. <https://cepa.org/comprehensive-reports/an-urgent-matter-of-drones/>
- Center for Strategic and International Studies (CSIS). (2024, October 23). *Assessing Russian firepower strikes in Ukraine*. <https://www.csis.org/analysis/assessing-russian-firepower-strikes-ukraine>
- Center for Strategic and International Studies (CSIS). (n.d.). *Russian firepower strike tracker: Analyzing missile attacks in Ukraine*. <https://www.csis.org/programs/futures-lab/projects/russian-firepower-strike-tracker-analyzing-missile-attacks-ukraine>
- Center for Strategic and International Studies (CSIS). (2025, May 13). *Drone saturation: Russia's Shahed campaign*. <https://www.csis.org/analysis/drone-saturation-russias-shahed-campaign>
- Center for Strategic and International Studies (CSIS). (2025, September 9). *Russia's massed strikes: The strategy of coercion by salvo*. <https://www.csis.org/analysis/russias-massed-strikes-strategy-coercion-salvo>
- Congressional Research Service. (2025, March 31). *Department of Defense counter unmanned aircraft systems (R48477)*. https://www.congress.gov/crs_external_products/R/PDF/R48477/R48477.2.pdf
- De Cubber, G. (2025). *Standardized evaluation of counter-drone systems: Methods, technologies, and performance metrics*. *Drones*, 9(5), 354. <https://www.mdpi.com/2504-446X/9/5/354>
- Department of Defense. (2023, January 25). *DoD Directive 3000.09: Autonomy in weapon systems*. <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>
- Department of Defense. (2024, December 5). *Fact sheet: Strategy for countering unmanned systems*. <https://media.defense.gov/2024/Dec/05/2003599149/-1/-1/0/FACT-SHEET-STRATEGY-FOR-COUNTERING-UNMANNED-SYSTEMS.PDF>
- European Commission. (2024, February 14). *Updated Ukraine recovery and reconstruction needs assessment*. https://ec.europa.eu/commission/presscorner/detail/en/ip_24_801
- Federal Aviation Administration. (2024, February 5). *UAS detection and mitigation systems ARC: Final report*. https://www.faa.gov/regulations_policies/rulemaking/committees/documents/media/UAS-Detection-Mitigation-Systems-ARC_Final-Report_02052024.pdf

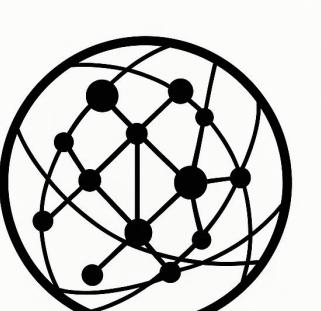
References

- Gong, J., et al. (2023). *Improved radar detection of small drones using Doppler signal-to-clutter ratio detector*. *Drones*, 7(5), 316. <https://www.mdpi.com/2504-446X/7/5/316>
- Government of Ukraine, World Bank, European Commission, & United Nations. (2024). *Third Ukraine rapid damage and needs assessment (RDNA3): February 2022–December 2023*. <https://ukraine.un.org/sites/default/files/2024-02/UA%20RDNA3%20report%20EN.pdf>
- Government of Ukraine, World Bank, European Commission, & United Nations. (2025). *Fourth Ukraine rapid damage and needs assessment (RDNA4): February 2022–December 2024*. <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099022025114040022>
- Government Accountability Office. (2022). *Counter-drone technologies*. <https://www.gao.gov/assets/gao-22-105705.pdf>
- Government Accountability Office. (2025, June 18). *Army modernization: Air and missile defense efforts*. <https://www.gao.gov/assets/gao-25-107491.pdf>
- Human Rights Watch & Harvard Law School International Human Rights Clinic. (2023). *Review of the 2023 U.S. policy on autonomy in weapons systems*. <https://humanrightsclinic.law.harvard.edu/wp-content/uploads/2023/02/Review-of-the-2023-US-Policy-on-Autonomy-in-Weapons-Systems.pdf>
- Institute for Science and International Security. (2025, December 1). *Monthly analysis of Russian Shahed-136 deployment against Ukraine*. <https://isis-online.org/isis-reports/monthly-analysis-of-russian-shahed-136-deployment-against-ukraine>
- Institute for the Study of War. (2025). *Russian force generation and technological adaptations updates*. <https://understandingwar.org>
- International Institute for Strategic Studies. (2025, April 14). *Russia doubles down on the Shahed*. <https://www.iiss.org/online-analysis/military-balance/2025/04/russia-doubles-down-on-the-shahed/>
- Joint Air Power Competence Centre. (2020). *A comprehensive approach to countering unmanned aircraft systems*. <https://www.japcc.org/wp-content/uploads/A-Comprehensive-Approach-to-Countering-Unmanned-Aircraft-Systems.pdf>
- Joint Chiefs of Staff. (2012). *JP 3-01: Countering air and missile threats*. <https://defenseinnovationmarketplace.dtic.mil/wp-content/uploads/2018/02/JointDoctrine-CounteringAirandMissileThreats.pdf>
- Khan, M. A., Menouar, H., Eldeeb, A., Abu-Dayya, A., & Salim, F. D. (2022). *On the detection of unauthorized drones: Techniques and future perspectives*. *IEEE Sensors Journal*. <https://doi.org/10.1109/JSEN.2022.3171293>
- National Institute of Standards and Technology. (2020). *Measuring and comparing small unmanned aircraft systems: Test methods – Introduction*. <https://www.nist.gov/system/files/documents/2020/07/06/NIST%20sUAS%20Test%20Methods%20-%20Introduction%20%282020B1%29.pdf>
- Open Source Munitions Portal. (n.d.). *Shahed-136 series*. <https://osmp.ngo/model/shahed-136-series/>
- Reuters. (2025, October 16). *Inside Ukraine's drone campaign to blitz Russia's energy*. <https://www.reuters.com/graphics/UKRAINE-CRISIS/RUSSIA-ENERGY/gdpzbxkgwpw/>

References

- Reuters. (2025, December 13). *Ukraine's Odesa suffers major blackouts after Russian attack.* <https://www.reuters.com/business/aerospace-defense/ukraines-odesa-suffers-major-blackouts-after-russian-attack-2025-12-13/>
- Sky News. (2025, November 4). *Russian “suicide drone” launches quadruple this year.* <https://news.sky.com/story/russian-suicide-drone-launches-quadruple-this-year-13463263>
- United Nations Peacekeeping. (2025). *Guidelines on counter unmanned aircraft systems.* <https://resourcehub01.blob.core.windows.net/%24web/Policy%20and%20Guidance/corepeacekeepingguidance/Thematic%20Operational%20Activities/Military/2025.16%20Guidelines%20on%20Counter%20Unmanned%20Aircraft%20Systems.pdf>
- U.S. Army. (2017). *ATP 3-01.81: Counter-unmanned aircraft system techniques.* https://aviation-assets.info/wp-content/uploads/ARN3099_ATP-3-01x81-FINAL-WEB.pdf
- U.S. Army. (2025). *FM 3-01: U.S. Army air and missile defense operations.* https://rdl.train.army.mil/catalog-ws/view/100.ATSC/C01CC9C1-DA1C-4D5E-A6EB-5FCFAE218DCD-1398170439966/FM_3_01wc1.pdf
- World Bank Open Knowledge Repository. (2024–2025). *Ukraine rapid damage and needs assessments (RDNA3–RDNA4).* <https://openknowledge.worldbank.org>

Appendix



GLOBAL AI
GOVERNANCE
RESEARCH CENTER

Appendix A: Abbreviations and Glossary

AAA: Anti-Aircraft Artillery.

AI: Artificial Intelligence.

CER: Cost Exchange Ratio.

CER*: Composite Cost–Loss Indicator.

CLER: Counterfactual Loss Exchange Ratio.

CPLA: Cost per Loss Avoided.

C2: Command and Control.

EW: Electronic Warfare.

GNSS: Global Navigation Satellite System.

IR: Infrared.

KAPR: Key Asset Preservation Ratio.

KAPS: Key Asset Preservation Score.

LER: Loss Exchange Ratio.

MVA: Minimum Viable, Auditable.

N: Attack Intensity (Proxy).

OLER: Observational Loss Suppression Ratio.

OSINT: Open-Source Intelligence.

RCS: Radar Cross Section.

SAM: Surface-to-Air Missile.

SCR: Signal-to-Clutter Ratio.

UAS / UAV: Unmanned Aerial System / Unmanned Aerial Vehicle.

Appendix A: Abbreviations and Glossary

Attack Intensity: The scale and frequency of inbound drone-related activity, used as a pressure proxy rather than a direct measure of effectiveness.

Cost Sustainability: The ability of a defense system to maintain acceptable performance without incurring economically or logically prohibitive costs over time.

Critical Assets: Infrastructure or systems whose functional degradation would produce disproportionate strategic, economic, or societal effects (e.g., electricity grids, fuel refineries).

Decoy Share (d): The estimated proportion of inbound aerial objects that do not carry effective attack payloads and are intended to exhaust defensive resources.

Functional Loss: Degradation in the operational capacity of a critical asset, measured through mission-relevant proxies such as blackout duration or production downtime.

Interceptor Drone: A low-cost unmanned platform designed to engage and destroy other drones, typically used as a terminal-layer defensive measure.

Loss Suppression: The reduction of functional degradation relative to a reference or baseline level of attack pressure.

Mission Outcome: The preservation or degradation of essential national or military functions, distinct from tactical engagement or intercept counts.

Saturation: A condition in which attack volume or frequency exceeds the defender's capacity to respond efficiently, leading to rising costs or leakage.

Terminal Defense Layer: The final defensive layer close to the protected asset, typically involving guns, short-range systems, EW, or interceptor drones.

Appendix B: Tables

Table 1A. Observational Sustainability and Mission Outcomes (MVA Baseline)

Side	Year	N _{total}	N _{attack}	Defense Mix	CER	OLER	CER [*] _{obs}	CPLA _{obs}	KAPS
Ukraine (vs Shahed)	2023	3,500	2,450	S2	1.2–1.5	0.32–0.36	0.38–0.54	0.002–0.004	0.48–0.52
Ukraine (vs Shahed)	2024	10,000	7,000	S2	1.3–1.6	0.34–0.38	0.44–0.61	0.003–0.006	0.45–0.50
Ukraine (vs Shahed)	2025*	45,000	31,500	S2	1.6–2.0	0.40–0.48	0.64–0.96	0.010–0.020	0.40–0.47
Russia (vs UA deep strike)	2023	10	7	S2	1.1–1.4	0.45–0.55	0.50–0.77	0.05–0.12	0.55–0.60
Russia (vs UA deep strike)	2024	81	57	S2	1.3–1.6	0.48–0.58	0.62–0.93	0.02–0.05	0.52–0.58
Russia (vs UA deep strike)	2025**	120	84	S2	1.4–1.8	0.50–0.60	0.70–1.08	0.02–0.04	0.50–0.55

Notes: Attack-intensity values (N_{total}, N_{attack}) are not numerically symmetric across sides and reflect different attack modes (high-frequency saturation vs. low-frequency node targeting); they are not intended for direct cross-side comparison.

a) Reporting conventions

- CER and CPLA are reported as ranges.
- OLER and KAPS are asset-level proxy results derived from observed functional-loss indicators.
- Cross-side comparisons are mechanism- and trend-focused rather than numerically commensurate, due to differences in loss proxies and asset definitions.

b) Sensitivity scope

Reported ranges reflect a bounded sensitivity grid over three parameters only:

- **Decoy share** $d \in \{0.2, 0.3, 0.4\}$;
- **Attack unit cost** u_{atk} ; and
- **Defensive unit-cost mix** u_{def} , implied by force-structure scenarios **S1/S2/S3** (low / base / high). These ranges are not statistical confidence intervals.

c) Units

- CER (defined in Eq. (10)) and CER^{*} (defined in Eq. (11)) are dimensionless ratios.
- OLER and KAPS are normalized scores on [0,1].
- CPLA_{obs} (defined in Eq. (9)) is reported as USD per unit of functional loss avoided, where the unit is defined by the selected loss proxy (e.g., USD per blackout-hour avoided).

d) Asterisks.

- * indicates values that are provisional or partially reported, reflecting incomplete annual coverage or reliance on interim public data for the most recent period.
- ** indicates values that are highly provisional, based on limited or ongoing reporting and therefore subject to greater uncertainty.

Asterisked values are included to preserve temporal continuity and trend analysis under the MVA standard, but should not be interpreted as finalized empirical estimates.

Appendix B: Tables

Table 1B (a). Unified counterfactual assumptions

Scenario	M
Low	1.3
Base	1.7
High	2.3

Table 1B (b). Counterfactual Stress Test Results (2025)

Side	Defense Mix	M	CER	CLER	CER [*] _{cf}	CPLA _{cf}
Ukraine	S1 SAM-heavy	1.3 / 1.7 / 2.3	2.5–3.0	0.43–0.59	1.1–1.8	0.03–0.06
Ukraine	S2 Balanced	1.3 / 1.7 / 2.3	1.6–2.0	0.43–0.59	0.7–1.2	0.02–0.04
Ukraine	S3 Low-cost	1.3 / 1.7 / 2.3	0.9–1.2	0.43–0.59	0.4–0.7	0.008–0.015
Russia	S2 Balanced	1.3 / 1.7 / 2.3	1.4–1.8	0.43–0.59	0.6–1.1	0.02–0.05

Notes:

a) Scenario status

All results in Table 1B are counterfactual and derived exclusively for stress-testing purposes under explicitly stated assumptions. They do not represent observed outcomes.

b) Indicator construction

- **CLER** (defined in Eq. (12)) is derived solely from the scenario loss multiplier **M** and reflects assumed loss amplification under degraded defensive conditions.
- **CER^{*}_{cf}** (defined in Eq. (13)) is reported to illustrate relative sensitivity across defense structures under identical counterfactual stress and must not be interpreted as an observed or empirical marginal cost.

c) Stress-parameter interpretation

Ranges in Table 1B reflect variation over the stress multiplier $M \in \{1.3, 1.7, 2.3\}$ under each force-structure scenario. Lower bounds correspond to $M = 1.3$ and upper bounds to $M = 2.3$. Intermediate values ($M = 1.7$) fall within the reported ranges.

d) Interpretation boundary

Results in Table 1B are intended for comparative analysis across force-structure scenarios (S1/S2/S3) under fixed stress parameters. They should not be compared numerically with observational indicators reported in Table 1A.

Appendix C: Decision Rules for Interpreting Counter-Drone Performance under Saturation

This appendix presents a structured decision framework for interpreting counter-drone performance under conditions of sustained saturation. The framework integrates cost sustainability, loss suppression, and mission outcomes, and is explicitly designed to prevent misinterpretation based on single indicators, such as intercept or shoot-down rates, when assessing system-level effectiveness.

C.1 Core Indicators

- **CER (Cost Exchange Ratio):** Assesses whether defense remains economically efficient relative to attack.
- **OLER (Observational Loss Suppression Ratio):** Measures the extent to which functional losses of critical assets are suppressed under observed attack pressure.
- **CER^{*}_{obs} (Composite Cost–Loss Indicator):** Evaluates whether cost pressure and residual functional loss are jointly worsening or improving over time.
- **CPLA_{obs} (Cost per Loss Avoided):** Captures the marginal cost required to prevent one unit of critical-asset functional loss.
- **KAPS (Key Asset Preservation Score):** Assesses whether essential national and military functions (e.g., power continuity, fuel supply) are being preserved.

C.2 Decision Matrix

Indicator Pattern	Strategic Interpretation	Policy Implication
Low CER + Low OLER + High KAPS	Defense is both cost-effective and mission-effective.	Maintain current force structure; prioritize sustainment and readiness.
High CER + Low OLER + High KAPS	Defense remains effective but is economically stressed.	Reduce reliance on high-cost interceptors; expand low-cost terminal layers.
Low CER + High OLER + Low KAPS	Defense is inexpensive but operationally ineffective.	Invest in detection, C2 integration, and terminal-layer effectiveness.
High CER + High OLER + Low KAPS	Defense approaching saturation or systemic failure.	Reprioritize protection toward critical assets; restructure defense concept.

Appendix C: Decision Rules for Interpreting Counter-Drone Performance under Saturation

C.3 Interpretive Guidance

- **Intercept rates alone are insufficient.**

High shoot-down performance can coexist with deteriorating sustainability when achieved through increasingly expensive interceptors.

- **CPLA serves as an early-warning signal.**

Rising CPLA_{obs} often precedes visible degradation in KAPS or OLER, indicating impending resource or inventory exhaustion.

- **KAPS anchors analysis to strategic outcomes.**

Preservation of power continuity and fuel supply may justify tolerating incomplete interception in lower-priority areas.

- **CER^{*}_{obs} captures structural trends.**

Persistent increases in CER^{*}_{obs} signal compounding cost pressure and residual loss, even when individual indicators appear stable.

C.4 Common Analytic Errors to Avoid

- Treating CER or intercept rate as standalone measures of effectiveness.
- Interpreting national-level damage totals (e.g., RDNA estimates) as drone-specific losses.
- Ignoring defense-mix composition (e.g., SAM vs. AAA vs. EW vs. interceptor drones).
- Reading counterfactual stress-test outputs (e.g., CER^{*}_{cf}, CPLA_{cf}) as observed reality rather than scenario bounds.