EPINOVA

EPINOVA–2025–PR–01

Nonlinear Uncertainty in Drone Warfare:

Why Indeterminacy Outperforms Precision in Contested ISR Environments

Dr. Shaoyuan Wu

The Global AI Governance and Policy Research Center, EPINOVA LLC

GLOBAL AI GOVERNANCE RESEARCH CENTER

December 31, 2025

About This Report

This policy report examines the strategic implications of uncertainty in contemporary drone warfare. It develops an analytical framework for understanding how nonlinear interaction, adversary adaptation, and endogenous observability costs jointly produce spatiotemporal indeterminacy in contested ISR environments. The report is intended to inform policymakers, analysts, and researchers engaged in defense planning, evaluation metrics, strategic stability, and governance design.

The analysis is strategic and conceptual rather than operational. It does not provide tactical guidance, system specifications, or employment recommendations. Instead, it provides a decision-facing analytical framing to support evaluation, planning, and governance judgments under conditions where prediction is fragile and precision-seeking can generate exposure and accelerate counter-adaptation. For detailed scope, limitations, and interpretive boundaries, see Appendix B.

Author Note

Dr. Shaoyuan Wu is an AI governance scholar and policy analyst whose work examines how artificial intelligence and emerging technologies reshape decision-making, institutional design, and strategic stability under conditions of uncertainty and competition.

He is the founder of EPINOVA LLC and Director of its Global AI Governance and Policy Research Center, which conducts policy-oriented research on AI policy frameworks, digital transformation, and the institutional impacts of frontier technologies. His current work focuses on strategic technology competition, AI-driven structural transitions, and human–AI coevolution in global governance.

Disclaimer

The views expressed in this report are solely those of the author and do not represent the official position of any government, organization, or institution.

This report is based exclusively on publicly available information and conceptual analysis and does not draw on classified, restricted, or proprietary sources. All assessments reflect information available as of December 29, 2025.

Any references to timing, posture mixtures, signaling, or adaptation dynamics are illustrative only and should not be construed as operational, tactical, or system-design guidance.

Table of Contents

GLOBAL AI GOVERNANCE RESEARCH CENTER

EPINOVA

Table of Contents

GLOBAL AI GOVERNANCE RESEARCH CENTER

EPINOVA

Foreword

Unmanned aerial systems have moved from peripheral enablers to central features of contemporary conflict. Yet much of the policy and analytical debate surrounding drone warfare remains anchored in a precision-centric paradigm—one that assumes improved sensing, tighter synchronization, and faster decision cycles will reliably translate into sustained advantage. This report begins from a different premise: in increasingly contested environments, the decisive challenge is not how to see or strike more precisely, but how to remain operationally relevant as adversaries adapt to being seen.

This analysis was motivated by a growing mismatch between prevailing evaluation frameworks and observed operational dynamics. As sensing, communication, and timing themselves become targets, efforts to reduce uncertainty can generate new vulnerabilities, raise exposure costs, and accelerate counter-adaptation. In such conditions, precision alone becomes a brittle foundation for strategy. The question facing policymakers is therefore not whether uncertainty can be reduced, but how it should be managed.

This report reframes drone warfare as a problem of uncertainty management in a nonlinear, adaptive competitive system. Drawing on a portfolio-inspired logic of robustness under regime uncertainty, it develops the Permanent Operational Configuration (POC) framework, which treats force posture as a mixture of presence modes rather than a single optimized configuration. The intent is not to prescribe tactics or system designs, but to offer decision-makers a structured way to think about persistence, survivability, and escalation risk under conditions where prediction is inherently fragile.

The report is written for policy audiences concerned with force planning, evaluation metrics, strategic stability, and governance. It deliberately avoids operational detail and platform-specific discussion. Formal models are used sparingly and only to clarify trade-offs and cost structures, not to forecast outcomes or recommend employment concepts. Governance implications are explored as emerging challenges rather than settled solutions.

Readers should approach this report not as a technical manual or doctrinal proposal, but as an effort to realign strategic thinking with the realities of contested observability. Its central aim is to support more resilient decision-making in an era where uncertainty is no longer merely an obstacle to be overcome, but a condition that must be actively shaped.

Executive Summary

Drone warfare has entered a phase in which uncertainty, rather than precision, increasingly determines operational advantage. In contested environments—where sensing, communication, and timing are themselves targets—efforts to improve detection and interception encounter structural limits (Biddle, 2023; Center for Strategic and International Studies, 2023; Watling & Reynolds, 2023). Gains in positional or temporal accuracy often generate countervailing exposure costs, accelerate adversary adaptation, and undermine long-term sustainability. As a result, performance metrics centered on detection rates, intercept probabilities, or synchronized mass effects risk becoming misleading indicators of strategic effectiveness.

This report argues that the central challenge in drone warfare is no longer how to eliminate uncertainty, but how to manage it under sustained competition. Advantage derives from shaping the cost structure of adversary uncertainty reduction—making localization, timing, and classification progressively more expensive—rather than from maximizing ambiguity for its own sake. This reframing shifts analysis away from point prediction and toward robustness under regime uncertainty.

To operationalize this shift, the report introduces the Permanent Operational Configuration (POC) framework, inspired by portfolio logic. Rather than optimizing force posture for a single anticipated scenario, POC treats posture as a mixture of complementary presence modes—deterrent presence, survivable reserve, mobile uncertainty, and temporal randomization. Each mode trades immediate effectiveness against survivability and predictability. The objective is not peak performance in any one regime, but sustained operational pressure and survivability across uncertain futures (Freedman, 2013; Gray, 2015).

Analytically, the framework situates drone warfare within a partially observable stochastic game in which observation quality is endogenous to both sides' actions. Uncertainty is modeled through adversary belief dynamics, with emphasis on the marginal cost of belief convergence rather than absolute levels of ambiguity. This logic is integrated into cost–distance–frequency analysis by recasting operational frequency as event intensity (hazard) under regime shifts and by introducing posture-dependent presence efficiency and capability constraints.

Finally, the report highlights implications for governance and strategic stability. Force structures optimized for spatiotemporal indeterminacy complicate traditional verification, increase misperception risks, and strain transparency-based approaches

Executive Summary

to arms control. Managing these risks requires a shift toward process-based accountability, signal-management norms, and audit-by-design mechanisms that preserve strategic stability without demanding full observability.

The central conclusion is straightforward: in modern drone warfare, the decisive variable is not seeing more, but ensuring that seeing enough becomes prohibitively costly for the adversary. Managing uncertainty—rather than attempting to eliminate it—has become a core strategic function.

What This Means for Decision-Makers

- **Drone warfare is no longer primarily about seeing more—it is about controlling how expensive it is for the adversary to see enough.** In contested environments, marginal gains in detection, tracking, or timing often come at rapidly increasing cost and risk. Systems optimized for precision and synchronization may perform well in isolated engagements but degrade quickly under sustained adaptation and countermeasures.
- **Prioritize robustness over single-scenario optimization.** Force postures built around a single anticipated scenario—such as persistent patrols, tightly synchronized swarms, or continuous ISR—are vulnerable to regime shifts. Portfolio-inspired posture mixtures that distribute emphasis across visibility, survivability, mobility, and temporal variability tend to exhibit greater regime-robustness in comparative evaluation, even when peak efficiency is not maximized.
- **Shift evaluation metrics from point performance to uncertainty-cost dynamics.** Traditional indicators such as intercept rates or detection probabilities capture only short-term outcomes. More decision-relevant metrics assess how adversary sensing and timing costs scale over time, how quickly belief convergence occurs, and how posture choices affect exposure accumulation and sustainability.
- **Treat uncertainty management as a governance and stability problem, not only an operational one.** Force structures optimized for spatiotemporal indeterminacy complicate verification and can increase misperception risk. Strategic stability will depend less on transparency of deployment and more on process-based accountability, auditability, and norms governing operational signaling and timing behavior.
- **Core Takeaway:** Success lies not in eliminating uncertainty, but in ensuring that reducing it becomes prohibitively costly for the adversary—while maintaining control over escalation and accountability.

What This Means for Decision-Makers

- **Key Finding 1 — Precision-centric drone defense is structurally brittle in contested ISR environments.**
  In environments where sensing, communication, and timing are actively contested, efforts to maximize positional or temporal precision generate endogenous exposure costs and accelerate adversary adaptation. Effectiveness degrades not because systems fail to detect or intercept, but because the marginal cost of sustaining precision rises faster than its operational value. Under sustained contestation, predictive control becomes a liability rather than a durable advantage.

- **Key Finding 2 — The decisive variable is not uncertainty itself, but the marginal cost of reducing it.**
  Operational advantage derives from shaping the adversary's belief-convergence dynamics. Postures that impose rising marginal costs on adversary efforts to localize, time, or classify drone activity outperform postures optimized for single anticipated scenarios. The strategic objective is not maximal ambiguity, but making "knowing enough" progressively more expensive at acceptable cost.

- **Key Finding 3 — Regime-robust force postures require portfolio-style posture mixtures rather than single-mode optimization.**
  A Permanent Operational Configuration (POC), modeled as a mixture of complementary presence modes, provides robustness across regime shifts such as high-pressure ISR, saturation, deception, and recovery. This approach deliberately trades peak effectiveness in any one scenario for sustained survivability, pressure, and adaptability across uncertain futures—without reliance on perfect intelligence or stable future conditions.

- **Key Finding 4 — Operational frequency should be treated as hazard, not schedule, under regime uncertainty.**
  Recasting frequency as event intensity (hazard) better captures the dynamics of drone warfare than deterministic scheduling. At the level of strategic interaction, non-deterministic timing and posture-contingent intensity variation can reduce the reliability of adversary timing inference, complicate escalation interpretation, and avoid making continuous activity a necessary condition for influence.

What This Means for Decision-Makers

- **Key Finding 5 — Spatiotemporal indeterminacy increases verification hardness and reshapes escalation risk.**
  Force structures optimized for indeterminacy reduce transparency and complicate traditional arms-control verification, because detection does not imply classification, and classification does not imply actionable certainty. Absent process-based accountability and signal-management norms, such postures can increase misperception and inadvertent escalation risk even as they enhance operational survivability.

- **Key Finding 6 — The correct analytical unit of effectiveness is robustness under endogenous observability, not interception performance.**
  Effective evaluation must shift away from point prediction, platform-centric performance, and intercept rates toward the dynamics of observability, adaptation, and belief convergence over time. The most decision-relevant metrics capture how posture choices reshape exposure accumulation and the cost trajectory of adversary uncertainty reduction, rather than short-run detection or kill-chain outcomes.

1. Introduction

1. Introduction

For much of the past decade, drone warfare has been evaluated through a precision-centric lens. Advances in sensing, data fusion, and command-and-control have reinforced the assumption that improved detection, tighter synchronization, and faster decision cycles would translate into durable operational advantage. Within this paradigm, uncertainty is treated primarily as a deficit—something to be reduced through better intelligence, more persistent ISR, and more accurate targeting.

This assumption, however, no longer holds in contested environments.

1.1 Spatiotemporal Indeterminacy as a Structural Condition

In contemporary drone warfare, uncertainty is not simply a residual effect of imperfect information or the traditional "fog of war," but an increasingly structural condition of the battlespace, produced by the interaction between sensing, exposure, and adversary adaptation (Jervis, 1976; Taleb, 2010). As sensing and communication systems become more capable, they also become more targetable. Efforts to improve positional or temporal certainty therefore generate countervailing costs: expanded signature exposure, increasingly predictable activity patterns, and accelerated countermeasures as adversaries learn how precision is achieved (Watts, 2004; Kofman, 2022).

In this context, "indeterminacy" does not refer to the absence of signals or the deliberate use of deception per se. Rather, it denotes the persistent inability to convert observation into actionable certainty at acceptable and sustainable cost, even when detection and partial awareness are present.

This creates a fundamental trade-off. The more precisely a force attempts to know where and when drones will appear, the more it risks revealing how it operates. Precision in one dimension tends to erode resilience in another. Continuous ISR improves situational awareness but increases detectability; tight synchronization enhances short-term effectiveness but creates exploitable temporal regularities; centralized control improves coordination but exposes communication dependencies. Regular timing patterns are actively exploited by adversaries in contested environments, undermining predictability-based control logics and accelerating counter-adaptation (Watling & Reynolds, 2023; RAND Corporation, 2022). These dynamics are not incidental failures of execution—they are intrinsic to the structure of contested observability.

1. Introduction

As a result, spatiotemporal indeterminacy in drone warfare is not something that can be engineered away through incremental improvements in sensing or analytics (Watts, 2004; Kofman, 2022). It reflects a deeper constraint: observation itself alters the cost and risk landscape. Both sides adapt not only to each other's capabilities, but to each other's attempts at reducing uncertainty. In such settings, prediction becomes fragile, and the pursuit of ever-greater precision can paradoxically reduce long-term effectiveness.

1.2 From Eliminating Uncertainty to Managing It

Recognizing indeterminacy as structural shifts the analytical question. The core challenge is no longer how to eliminate uncertainty, but how to manage it under sustained competition. This requires moving beyond point prediction—when and where a drone will appear—and toward assessing how posture choices shape the economics of observability: the costs an adversary must incur to localize, time, and classify activity with sufficient confidence.

This reframing has important implications for policy analysis. Traditional metrics—detection rates, intercept probabilities, sortie counts—capture outcomes at specific moments but obscure the dynamics that determine whether those outcomes can be sustained. They say little about how quickly adversaries adapt, how exposure accumulates over time, or how escalation risks evolve as both sides probe each other's sensing and signaling thresholds.

By contrast, a robustness-oriented perspective treats uncertainty as a variable to be shaped rather than erased. Advantage derives from imposing rising marginal costs on adversary efforts to reduce uncertainty, even when absolute ambiguity cannot be maintained. The objective is not maximal opacity, but controlled indeterminacy: enough unpredictability to disrupt adversary planning while preserving operational relevance and strategic control.

1.3 Policy Relevance: Sustainability, Survivability, and Escalation Management

For policymakers, this shift from predictive control to robustness is consequential. First, it reframes sustainability. Forces optimized for precision may perform well in isolated engagements but degrade rapidly under continuous contestation as exposure costs compound. Robust postures prioritize endurance over peak performance, enabling prolonged operational relevance in uncertain conditions.

1. Introduction

Second, it reshapes survivability. Rather than relying solely on concealment or hardening, survivability increasingly depends on how difficult and expensive it is for an adversary to generate actionable certainty. Systems that remain partially indeterminate force adversaries to choose between escalating sensing investments and accepting residual risk.

Finally, this reframing affects escalation management. Precision-centric models implicitly assume that greater transparency and predictability stabilize interactions. In practice, however, tightly coupled sensing and timing can compress decision windows and lower thresholds for rapid escalation when misinterpretation occurs. Robustness-oriented approaches, while complicating verification, may offer alternative pathways to stability if paired with appropriate governance mechanisms.

Building on these observations, this report develops an analytical framework for uncertainty management in drone warfare. Rather than treating indeterminacy as an operational failure, it examines how force posture, observability, and adversarial adaptation interact under regime uncertainty—and what this implies for evaluation metrics, force planning, and strategic governance in increasingly contested ISR environments.

The analysis abstracts from earlier formal and empirical work on uncertainty management, cost–frequency dynamics, and unmanned systems competition (Wu, 2025a; Wu, 2025b).

2. Structural Indeterminacy in Drone Warfare

2. Structural Indeterminacy in Drone Warfare

The limits of precision-centric approaches in drone warfare do not stem from technological shortfalls or execution errors. They arise from a set of structural trade-offs that shape how sensing, control, and survivability interact under sustained contestation. These trade-offs impose constraints on what can be simultaneously known, controlled, and protected. Understanding them is essential before introducing any alternative framework for force posture or evaluation.

2.1 Measurement–Exposure Trade-offs in Contested ISR

At the core of contemporary drone warfare lies a fundamental measurement–exposure trade-off. Efforts to improve detection, tracking, and classification inevitably generate observable signatures—electromagnetic, acoustic, temporal, or behavioral—that adversaries can exploit. As ISR systems become more persistent and precise, they also become more legible as targets.

This dynamic undermines the assumption that improved sensing is a monotonic good. Persistent ISR can stabilize situational awareness in the short term, but over time it creates predictable patterns of activity and communication that facilitate countermeasures. High-fidelity tracking reduces ambiguity at the cost of increased exposure; intermittent or degraded sensing preserves survivability but accepts informational gaps. These are not temporary frictions but enduring features of contested environments.

Crucially, exposure is cumulative. Each additional unit of sensing or control not only improves current awareness but also contributes to a growing body of information about how a force operates. Adversaries learn not just where drones are, but how often they appear, how they are coordinated, and which signals precede action. Measurement thus reshapes the future information environment, altering the cost and feasibility of subsequent operations.

2.2 Three Core Indeterminacy Trade-offs

These dynamics manifest in at least three recurring trade-offs that structure drone warfare under contestation.

- **Spatiotemporal Indeterminacy (Time–Space Trade-off).** Forces rarely face uncertainty about location or timing in isolation. Fixing one dimension tends to expose the other. Concentrating activity in a narrow time window improves coordination but creates exploitable temporal signatures. Dispersing activity

2. Structural Indeterminacy in Drone Warfare

over time reduces predictability but dilutes immediate effect. Similarly, spatial concentration enhances local impact while increasing vulnerability to detection and targeting. Attempts to fully resolve both dimensions simultaneously encounter diminishing returns and rising risk.

- **Concentration–Survivability Trade-off.** Mass and concentration remain powerful tools, particularly in saturation or strike scenarios. However, concentrated forces are inherently more legible and more fragile under persistent surveillance. Dispersion improves survivability by complicating targeting, but often at the cost of reduced per-unit effectiveness and increased coordination burdens. This tension is not a matter of doctrinal preference; it reflects a structural constraint between efficiency and endurance.
- **Control–Autonomy Trade-off.** Tight control enables coordination, synchronization, and rapid re-tasking, but it also increases reliance on communications that can be detected, disrupted, or exploited. Greater autonomy reduces communication signatures and enhances survivability, yet introduces variability and limits centralized predictability. Under contested conditions, neither extreme is stable. Forces must navigate a spectrum in which control and autonomy impose opposing risks.

Together, these trade-offs define a space of irreducible indeterminacy. They explain why improvements in sensing, networking, or analytics alone cannot eliminate uncertainty without introducing new vulnerabilities elsewhere.

2.3 Why Indeterminacy Is Not Just “Fog of War”

It is tempting to interpret these trade-offs as a modern variant of the classic fog of war. That interpretation is incomplete. Traditional fog of war arises from incomplete information about a largely passive environment. In drone warfare, by contrast, uncertainty is actively produced and contested. Adversaries adapt not only to observed behavior, but to the methods used to observe and control.

This distinction matters. Fog of war implies that uncertainty can, in principle, be reduced through better intelligence and organization. Structural indeterminacy implies limits that persist even as capabilities improve. In such settings, the marginal cost of reducing uncertainty rises, and attempts to drive ambiguity toward zero can be counterproductive.

2. Structural Indeterminacy in Drone Warfare

The implication is that uncertainty has a lower bound shaped by adversary adaptation and exposure dynamics. Beyond a certain point, further investments in precision yield diminishing—and potentially negative—returns.

2.4 Implications for Force Evaluation

Recognizing indeterminacy as structural has immediate consequences for how drone forces are evaluated. Metrics that focus on instantaneous outcomes—detection probability, intercept rate, or sortie efficiency—capture only a snapshot of performance. They fail to account for how quickly those outcomes degrade as adversaries adapt and exposure accumulates.

More informative assessments examine how posture choices affect the trajectory of uncertainty over time: how rapidly adversaries can localize activity, how costly it is for them to maintain confidence, and how exposure compounds across repeated interactions. These dynamics cannot be captured by single-mode optimization or static configurations.

This sets the stage for an alternative approach. If indeterminacy cannot be eliminated, then advantage lies in how it is managed. The next section introduces the Permanent Operational Configuration (POC) framework, which treats force posture not as a fixed configuration optimized for a single scenario, but as a mixture of presence modes designed to remain effective across regime uncertainty.

3. The Permanent Operational Configuration Framework

3. The Permanent Operational Configuration (POC) Framework

The structural trade-offs outlined in the previous section suggest a clear implication: if spatiotemporal indeterminacy cannot be eliminated, then force posture must be designed to operate within it. This requires moving away from configurations optimized for a single anticipated scenario and toward postures that remain viable across regime uncertainty. The POC framework provides a way to conceptualize such postures.

3.1 Portfolio Logic Under Regime Uncertainty

The POC framework draws inspiration from portfolio logic, not as a literal transfer of financial theory, but as an analytical approach to robustness under uncertainty. In environments where future states cannot be reliably predicted and regime shifts are frequent, optimizing for a single expected outcome exposes forces to catastrophic underperformance when assumptions fail. Portfolio logic addresses this problem by distributing exposure across multiple, partially competing modes, accepting bounded inefficiency in exchange for resilience.

Applied to drone warfare, this logic reframes posture selection. Instead of asking which configuration performs best under a presumed set of conditions, POC asks how different presence modes interact to sustain operational relevance when conditions change. The objective is not to hedge against every conceivable outcome, but to avoid reliance on any single mode whose failure would be decisive.

As shown in Figure 1, this approach contrasts sharply with precision-centric postures. Configurations optimized for detection, synchronization, or mass effects allow adversaries to reduce uncertainty at relatively low marginal cost once patterns are identified. Portfolio-style postures, by contrast, alter the adversary’s cost curve, making further reductions in uncertainty increasingly expensive even when some information is available.

3. The Permanent Operational Configuration (POC) Framework

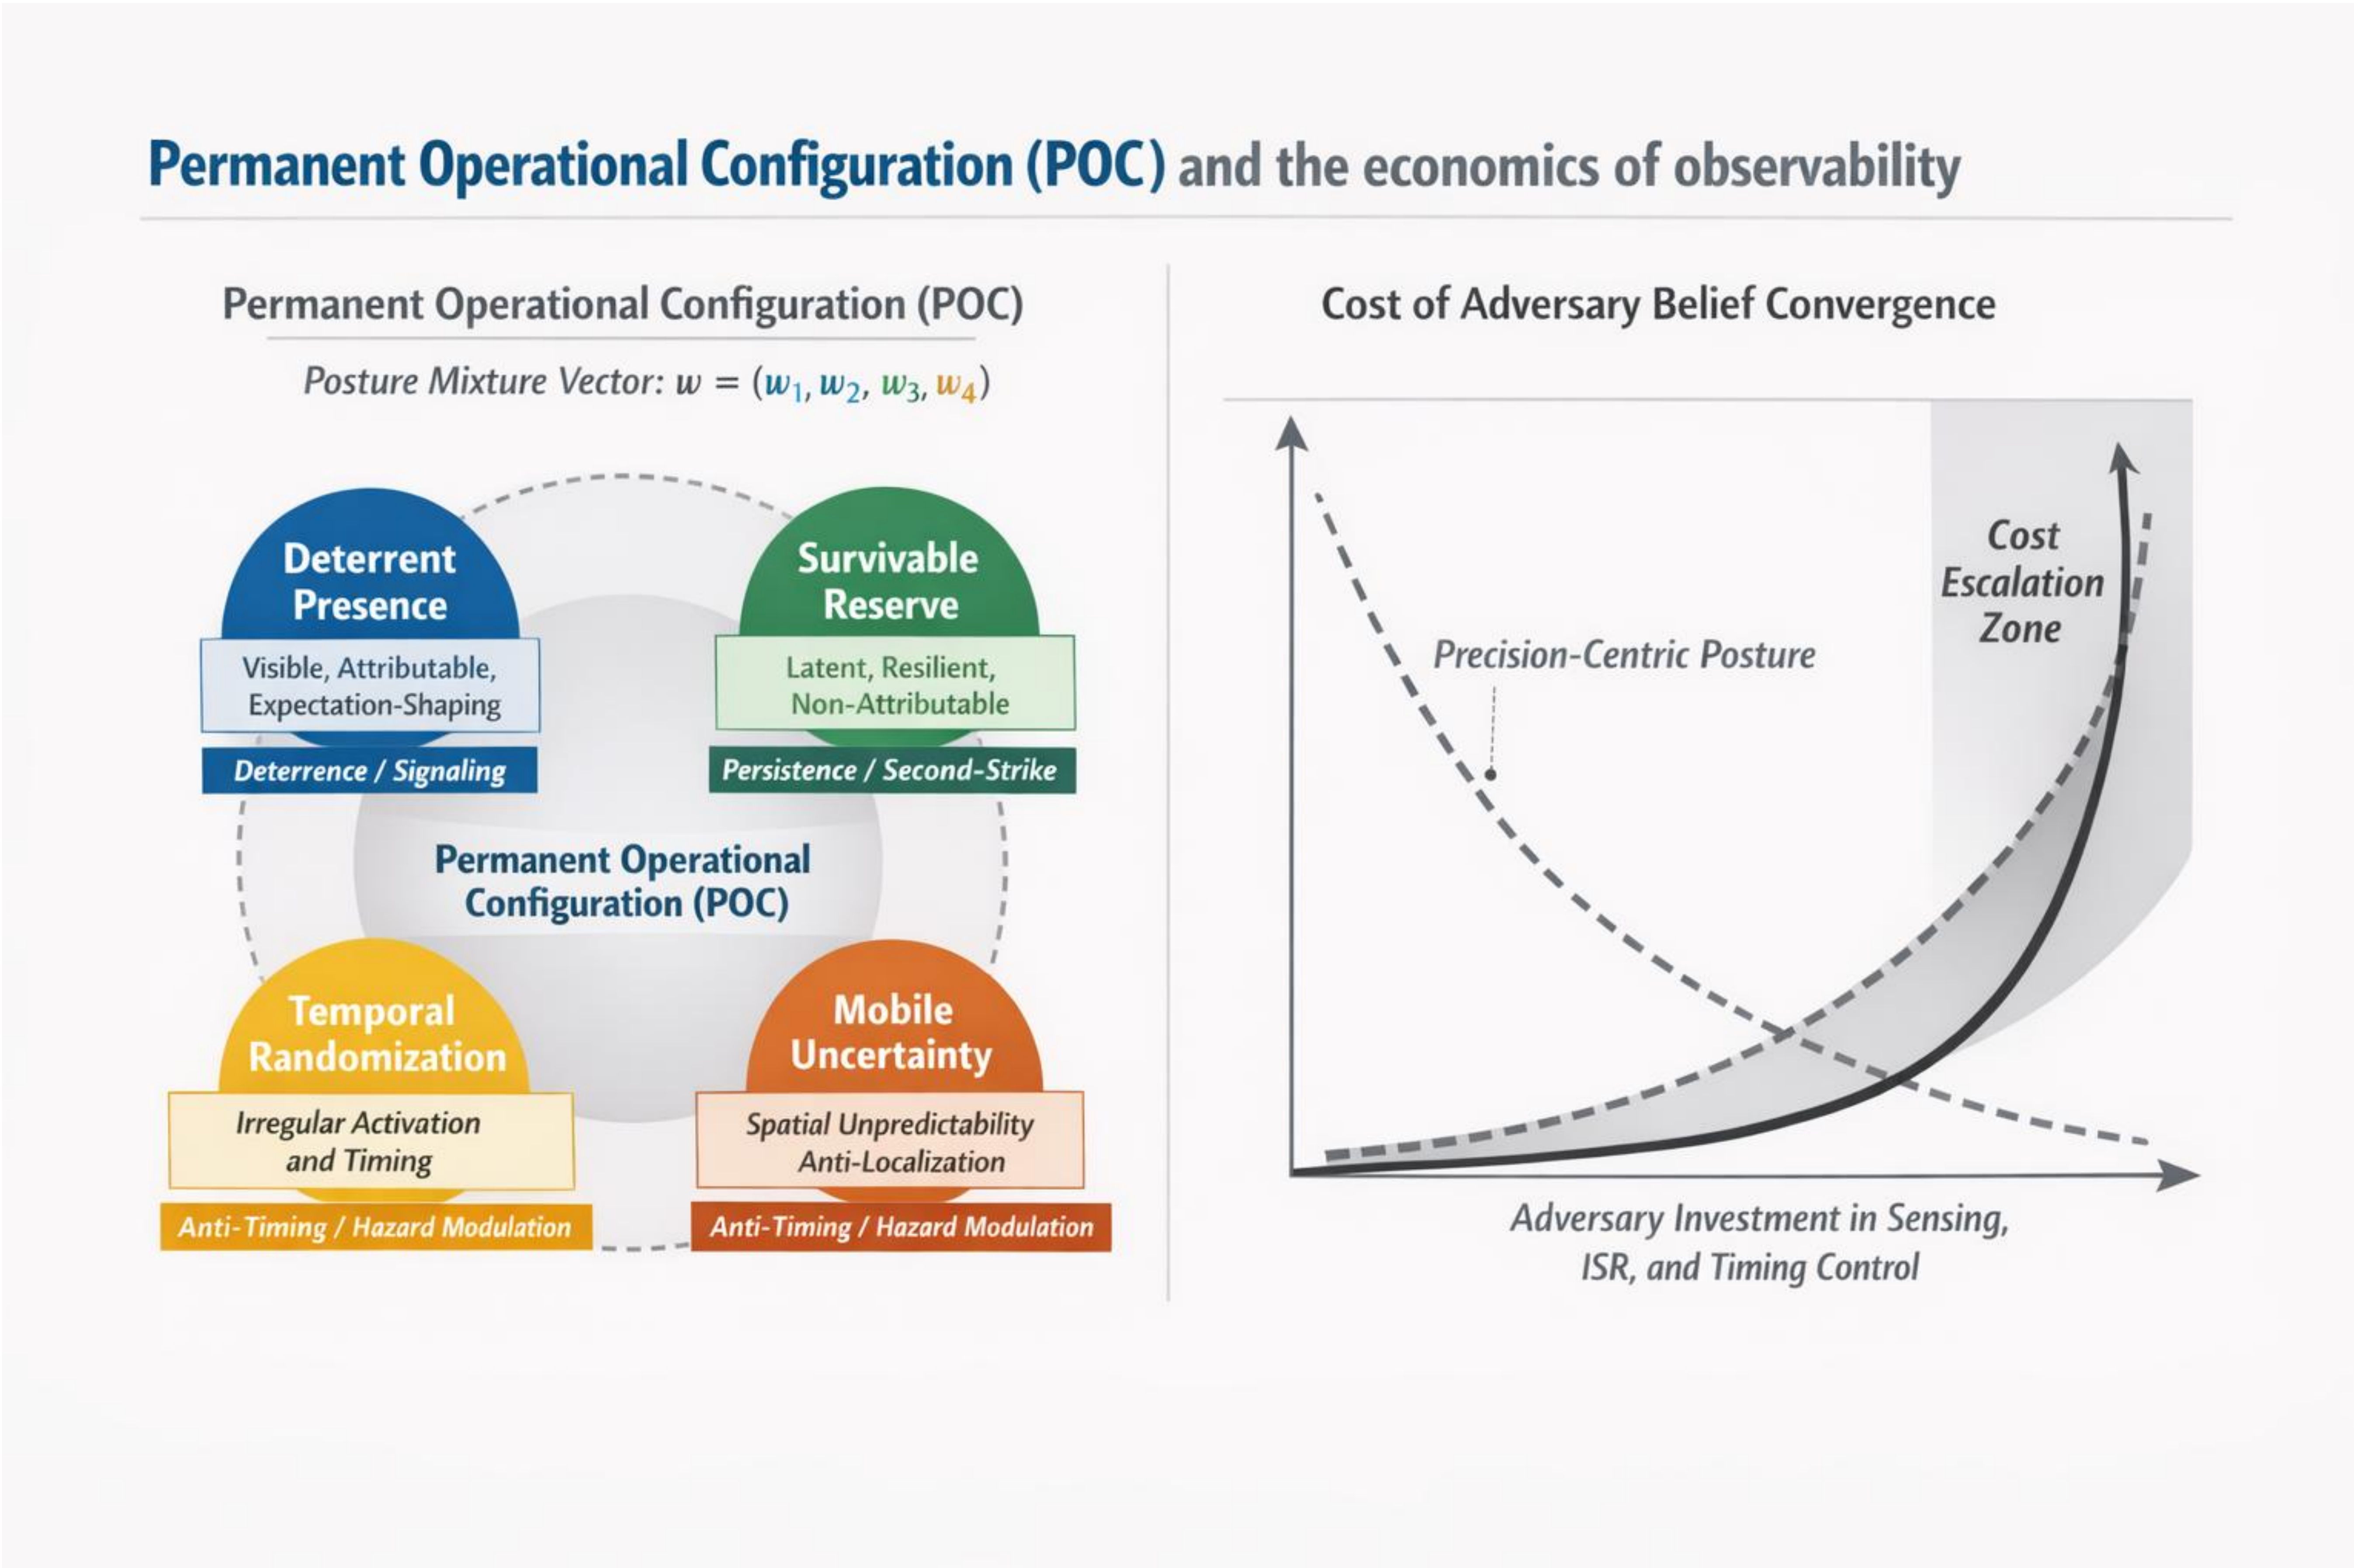

Figure 1. Permanent Operational Configuration (POC) and Adversary Uncertainty Costs
POC treats force posture as a mixture of presence modes rather than a single optimized configuration, reshaping the cost curve of adversary sensing and timing control. Compared with precision-centric postures, POC increases the marginal cost of belief convergence and delays the transition from detection to actionable certainty.

3.2 POC as a Posture-Mixture, Not a Configuration

A critical distinction in the POC framework is between a configuration and a mixture. A configuration implies a fixed arrangement optimized for a specific task or scenario. A mixture describes the relative weighting of different presence modes coexisting within the force posture at any given time.

Formally, POC treats posture as a vector of weights over presence modes. These weights are not orders of battle, deployment plans, or platform allocations. They are analytical descriptors of how the force distributes emphasis across mutually tensioned modes of presence. This abstraction allows comparison across posture classes without encoding operational detail.

3. The Permanent Operational Configuration (POC) Framework

The value of a mixture lies in its internal tension. Each mode imposes constraints on the others; none can be maximized without degrading the rest. POC does not seek to resolve these tensions, but to maintain them deliberately. The coexistence of partially contradictory presence modes prevents adversaries from collapsing uncertainty along a single dominant axis.

3.3 Four Presence Modes and Their Functional Roles

Within the POC framework, four presence modes capture the core functional dimensions of drone posture under contestation. These modes are analytical categories, not force elements.

- **Deterrent Presence.** Visible or legible activity that signals capability and intent. Deterrent presence shapes adversary expectations and behavior, but increases exposure and predictability if relied upon exclusively.
- **Survivable Reserve.** Latent or low-observability capacity preserved for endurance and recovery. Survivable reserves enhance second-move potential but contribute little to immediate influence if overemphasized.
- **Mobile Uncertainty.** Movement and variation that disrupt localization and pattern recognition. Mobility complicates targeting and prediction, but often reduces efficiency and coordination.
- **Temporal Randomization.** Irregular timing that undermines adversary scheduling and escalation models. Temporal randomization preserves indeterminacy without continuous activity, but limits synchronization and peak effect.

Individually, each mode is insufficient. Collectively, they create a posture that is difficult to classify, time, or neutralize cheaply. The effectiveness of POC arises not from maximizing any single mode, but from sustaining a mixture in which adversaries cannot confidently infer which mode dominates at a given moment.

3.4 Shaping the Economics of Observability

The central strategic contribution of POC lies in how it reshapes the economics of observability. Precision-centric postures allow uncertainty to be reduced quickly once sensing and adaptation catch up. POC, by contrast, ensures that each incremental gain in adversary certainty requires disproportionate additional effort.

3. The Permanent Operational Configuration (POC) Framework

As illustrated in Figure 1, the goal is not to maintain maximal ambiguity indefinitely. Some information will always be available. The objective is to impose escalating marginal costs on further belief convergence, making it increasingly expensive for the adversary to move from partial awareness to actionable certainty.

This distinction is critical. POC does not aim to "hide" forces completely, nor to deceive continuously. It aims to prevent uncertainty from collapsing cheaply. Even when adversaries detect activity, they face persistent difficulty in determining whether it is decisive, representative, or temporally relevant.

3.5 Analytical Boundary of the Framework

It is essential to emphasize what POC is not. It is not a deployment doctrine, a command-and-control architecture, or a prescription for autonomy levels. It does not specify how to implement presence modes or how to allocate platforms. Those questions depend on context, technology, and policy constraints beyond the scope of this report.

POC is an evaluative and planning framework. It provides a structured way to reason about posture robustness, exposure dynamics, and adaptation costs under regime uncertainty. Its value lies in shifting analysis away from single-scenario optimization and toward sustained performance across unknown futures.

The next section formalizes this intuition by modeling drone warfare as a partially observable stochastic game in which observability is endogenous and belief dynamics, rather than point predictions, determine strategic advantage.

4. Analytical Model: Observability, Belief, and Adaptation

4. Analytical Model: Observability, Belief, and Adaptation

The POC framework rests on a simple but consequential premise: in contested drone warfare, advantage is determined less by what is known at a given moment than by how costly it is for an adversary to know enough over time. To make this intuition analytically precise—without resorting to tactical specification—this section situates drone warfare within a formal but policy-oriented modeling framework.

4.1 Why a Partially Observable Stochastic Game

Drone warfare in contested environments cannot be adequately represented by static optimization or complete-information models. Both sides operate with incomplete and asymmetric information, and—critically—their actions alter the information environment itself. Surveillance, communication, and control are not neutral inputs; they change exposure, induce countermeasures, and reshape future observability.

For these reasons, the interaction is best understood as a partially observable stochastic game. Each side acts under uncertainty about the true system state, updates beliefs based on imperfect observations, and adapts behavior in anticipation of the opponent's learning and counter-learning. The environment introduces stochastic variation—weather, terrain, spectrum conditions—but uncertainty is driven primarily by strategic interaction rather than random noise.

This framing aligns with the core claim of the report: uncertainty in drone warfare is not merely informational, but endogenous and contested.

4.2 States, Observations, and Actions (Conceptual Level)

At an abstract level, the true system state comprises several interacting components: spatial distribution of activity, temporal regimes of operation, adversary capabilities and intent, and environmental conditions. None of these are fully observable. What each side perceives is a stream of observations—sensor returns, communication signatures, timing patterns—whose quality depends on both sides' actions.

Actions serve dual roles. They influence operational outcomes, but they also shape future observability. Increasing control intensity may improve coordination while simultaneously increasing detectability. Reducing activity may preserve survivability while degrading influence. Every action therefore trades immediate effect against informational consequences.

4. Analytical Model: Observability, Belief, and Adaptation

The model does not assign specific functional forms or numerical parameters. Its purpose is to capture directional relationships: how choices about posture and activity affect what the adversary can infer, how quickly, and at what cost.

4.3 Belief Dynamics and the Cost of Convergence

Rather than focusing on absolute uncertainty, the model emphasizes belief dynamics—how adversaries update their understanding of the system over time. Each side maintains a belief distribution over possible states, refining it as new observations arrive.

The strategic question is not whether beliefs eventually converge; some convergence is inevitable. The decisive variable is the marginal cost of further convergence. When uncertainty can be reduced cheaply—through persistent ISR, predictable timing, or stable signatures—precision-centric postures become vulnerable. Once patterns are learned, adversaries can act with confidence at low cost.

POC alters this dynamic. By sustaining a mixture of presence modes, it prevents belief convergence from accelerating smoothly. Initial observations may reduce uncertainty, but further refinement requires disproportionate investment. The adversary must expend additional sensing, analytic effort, or risk escalation to achieve marginal gains in confidence.

In this sense, the model reframes uncertainty as an economic problem. The objective is not to maximize entropy or preserve ignorance indefinitely, but to shape the slope of the belief-convergence curve.

4.4 Endogenous Observability and Measurement Disturbance

A central implication of the model is that observability is endogenous. Efforts to observe more precisely feed back into the system by increasing exposure, revealing patterns, and inviting countermeasures. Measurement is not passive; it perturbs the strategic environment.

This dynamic parallels—but should not be confused with—physical measurement disturbance. The analogy is structural, not literal. In drone warfare, disturbance arises from adversary adaptation rather than physical law. The more aggressively one side attempts to resolve uncertainty, the more information it reveals about its priorities, thresholds, and timing.

4. Analytical Model: Observability, Belief, and Adaptation

The model therefore rejects the assumption that better sensing monotonically improves outcomes. Beyond a point, additional observability may reduce robustness by collapsing indeterminacy too quickly and cheaply—for both sides.

4.5 Implications for Optimization and Evaluation

Within this framework, traditional optimization objectives appear misaligned. Maximizing detection probability or intercept rate at a given time step ignores how those gains affect future exposure and adaptation. Policies that perform well in short-run simulations may underperform in prolonged competition.

A robustness-oriented objective instead prioritizes long-run performance under belief dynamics. Success is measured by the ability to maintain operational relevance as adversaries learn, adapt, and probe. This naturally favors posture mixtures over single-mode optimization and validates the POC logic introduced in the previous section.

The analytical model thus provides a disciplined foundation for the shift from predictive control to uncertainty management. It clarifies why precision-centric approaches degrade under contestation and why portfolio-style postures can outperform them over time.

The next section translates this logic into an operational-analytic setting by integrating POC into cost–distance–frequency models. In doing so, it recasts frequency as hazard rather than schedule and links uncertainty management directly to evaluation metrics and planning tools familiar to policymakers.

5. Integrating POC into Cost–Distance–Frequency Analysis

5. Integrating POC into Cost–Distance–Frequency Analysis

The preceding sections established that advantage in contested drone warfare depends on managing uncertainty rather than eliminating it. For this insight to inform policy and planning, it must be translated into analytic tools already used to assess force effectiveness and sustainability. This section operationalizes the Permanent POC framework by integrating it into the cost–distance–frequency (CDF) family of models, with a critical modification: treating frequency as hazard rather than schedule.

5.1 From Deterministic Frequency to Hazard Under Regime Uncertainty

Traditional CDF models treat operational frequency as a controllable and largely deterministic variable—sortie rates, patrol cycles, or engagement intervals. This assumption breaks down in contested environments, where timing regularity itself becomes a vulnerability. Predictable schedules allow adversaries to economize sensing, align countermeasures, and lower escalation thresholds.

Under uncertainty management, frequency is better understood as event intensity (hazard) rather than a fixed schedule. Activity still occurs at an average rate, but its timing is deliberately irregular and contingent on regime conditions. From the adversary’s perspective, events are probabilistic rather than clock-driven, complicating efforts to anticipate when activity is decisive.

Recasting frequency in this way aligns with the core logic of POC. Temporal randomization does not require constant activity or higher sortie counts. Instead, it alters the adversary’s timing model, forcing greater investment to achieve confidence about when action matters.

5.2 Cost, Distance, and the Economics of Observability

Within the modified CDF framework, cost and distance retain their relevance but acquire new interpretation. Distance affects sensing feasibility and response latency, but its operational meaning depends on posture. Mobile uncertainty can decouple distance from predictability, while deterrent presence can make even distant activity strategically salient.

Cost is no longer limited to platform attrition or resource expenditure. It includes exposure costs: the accumulation of information about patterns, thresholds, and dependencies. Precision-centric postures often minimize immediate costs but incur

5. Integrating POC into Cost–Distance–Frequency Analysis

high long-run exposure. POC redistributes cost over time, accepting bounded inefficiency to prevent cheap belief convergence.

Operational value, in this formulation, is shaped by how posture choices influence the adversary's ability to reduce uncertainty cheaply. A posture that yields modest short-term effect but preserves indeterminacy may outperform a posture that delivers decisive early outcomes but rapidly becomes predictable.

5.3 Presence Efficiency and Capability Constraints

To make posture comparisons meaningful, the integrated framework introduces two abstract modifiers.

- **Presence Efficiency.** Presence efficiency captures how effectively a given posture mixture translates existence into operational relevance under specific environmental and regime conditions. It reflects factors such as terrain, weather, spectrum congestion, and adversary ISR pressure. Presence efficiency is not constant; it varies with both posture and context.
- **Capability Constraints.** Capability constraints represent aggregate limits on resources, sustainment, organizational capacity, and political tolerance. These constraints bound the feasible space of posture mixtures. POC does not assume unlimited resources or perfect flexibility; it operates within realistic ceilings on endurance and adaptation.

Together, these modifiers prevent the framework from collapsing into a purely theoretical exercise. They ensure that robustness is evaluated relative to feasible and sustainable force employment.

5.4 Expected Risk and Belief-Conditioned Loss

Integrating uncertainty management into CDF analysis also requires reframing risk. Expected loss is conditioned not only on physical exposure, but on adversary belief states and hazard dynamics. As adversaries gain confidence about timing or location, risk escalates nonlinearly.

POC mitigates this escalation by keeping belief convergence incomplete or costly. Even when adversaries detect activity, uncertainty about its representativeness or timing reduces the likelihood of efficient counteraction. Risk is managed not by eliminating exposure, but by preventing exposure from becoming decisive.

5. Integrating POC into Cost–Distance–Frequency Analysis

5.5 Implications for Assessment and Planning

This operationalization has direct implications for policy analysis and planning. First, it suggests that evaluation should focus on cost curves rather than point estimates. How quickly do adversary sensing costs rise as posture mixtures change? How does hazard-based timing affect escalation thresholds over repeated interactions?

Second, it challenges planning assumptions that equate activity with effectiveness. Irregular, posture-dependent activity may deliver greater long-term value than constant presence, even if average output appears lower.

Finally, it provides a common analytic language for comparing precision-centric and robustness-oriented approaches. By embedding uncertainty management into familiar CDF constructs, the framework enables decision-makers to assess trade-offs without abandoning established planning tools.

The next section extends this analysis to strategic stability and governance. As force postures become optimized for spatiotemporal indeterminacy, traditional assumptions about transparency, verification, and escalation control are strained—raising questions that cannot be resolved through operational analysis alone.

6. Strategic Stability and Governance Implications

6. Strategic Stability and Governance Implications

Optimizing drone warfare for spatiotemporal indeterminacy has consequences that extend beyond operational effectiveness. As force postures shift from precision-centric control toward robustness under uncertainty, long-standing assumptions about transparency, verification, and escalation management are strained. These implications are not incidental; they follow directly from the altered economics of observability described in earlier sections.

6.1 Verification Hardness and the Limits of Transparency

Traditional arms-control and confidence-building frameworks rest on a core assumption: that stability is enhanced when capabilities and deployments can be observed, counted, and verified. In drone warfare optimized for indeterminacy, this assumption becomes increasingly difficult to sustain.

Portfolio-style postures deliberately blur signals of presence, timing, and representativeness. Activity may be observable without being interpretable; detection does not imply classification, and classification does not imply confidence about intent or imminence. As a result, verification becomes less a matter of confirming physical facts and more a problem of interpreting probabilistic signals under strategic uncertainty.

This does not mean that transparency is undesirable, but it does mean that static transparency loses traction. Demands for continuous visibility or precise disclosure may undermine survivability without delivering commensurate stability benefits. In extreme cases, attempts to enforce transparency can incentivize destabilizing behavior by exposing vulnerabilities that adversaries can exploit cheaply.

As force structures are optimized for spatiotemporal indeterminacy, traditional verification mechanisms face increasing hardness: observability no longer guarantees interpretability, and disclosure no longer ensures strategic reassurance.

6.2 Indeterminacy, Misperception, and Escalation Risk

A second implication concerns escalation dynamics. Precision-centric models often assume that clearer information reduces miscalculation. In practice, tightly coupled sensing and timing can compress decision windows and amplify the consequences of misinterpretation. When systems are optimized for rapid detection and response, ambiguous signals may trigger disproportionate reactions.

6. Strategic Stability and Governance Implications

Indeterminacy alters this dynamic in complex ways. On one hand, uncertainty about timing and intent can create caution, slowing escalation. On the other, persistent ambiguity, especially when unaccompanied by shared interpretive frameworks, can increase anxiety and encourage worst-case assumptions. The risk is not uncertainty per se, but unmanaged uncertainty: situations in which actors lack common expectations about how signals should be read and weighted.

This underscores the need to treat indeterminacy as a strategic stability variable. Robust postures that complicate adversary planning must be paired with mechanisms that prevent ambiguity from translating into inadvertent escalation. Absent such mechanisms, force structures optimized for indeterminacy can heighten misperception and escalation risk, particularly under conditions of entangled sensing and autonomous adaptation (Acton, 2020; Crootof, 2022; UNIDIR, 2023).

6.3 From Static Verification to Process-Based Accountability

If physical observability cannot be reliably guaranteed, governance must shift focus. One promising direction is process-based accountability. Rather than verifying the precise location or timing of assets, oversight mechanisms can emphasize traceable decision authority, documented escalation thresholds, and auditable system behavior.

Such approaches do not require full disclosure of sensitive operational details. They seek assurance that actions are bounded by known processes, responsibilities, and constraints. This logic mirrors developments in AI governance, where auditability and accountability are increasingly emphasized over full transparency of models or data.

For drone warfare, process-based accountability could provide a foundation for stability even when spatiotemporal verification is infeasible.

6.4 Signal Management Norms

A further implication is the need for norms governing signal management. In an environment where timing, presence, and activity patterns are strategically manipulated, the absence of shared expectations increases misinterpretation risk.

Signal management norms would not prohibit indeterminacy, but would delineate boundaries around its use. Examples include understandings about what constitutes

6. Strategic Stability and Governance Implications

routine variation versus escalatory signaling, or about how certain activity patterns should be interpreted during crises. Such norms are necessarily imperfect, but they can reduce the likelihood that indeterminacy is mistaken for imminent attack.

Importantly, signal management focuses on behavior rather than capability. It does not require states to reveal how they achieve indeterminacy, only to clarify how signals should be read in specific contexts.

6.5 Audit-by-Design and Governance Compatibility

Finally, the shift toward robustness-oriented postures raises questions about compliance and legitimacy. Systems designed to resist observation must still be governable. This points toward audit-by-design approaches, in which accountability mechanisms are embedded at the design stage rather than imposed after deployment.

Audit-by-design does not imply real-time monitoring or external visibility. It implies that systems generate verifiable records of decision pathways, authorization, and constraint adherence that can be reviewed post hoc or under agreed conditions. This allows for accountability without sacrificing survivability.

6.6 Strategic Implications

Taken together, these governance challenges suggest that stability in drone warfare will depend less on perfect information and more on institutional adaptation. As uncertainty becomes a strategic resource, managing its effects—rather than suppressing it—becomes central to both effectiveness and restraint.

The next section translates these implications into concrete policy recommendations, identifying steps decision-makers can take to align evaluation metrics, force planning, and governance mechanisms with the realities of contested observability.

7. Policy Implications and Recommendations

7. Policy Implications and Recommendations

The analysis in this report points to a clear conclusion: existing approaches to drone warfare assessment, planning, and governance are increasingly misaligned with the realities of contested observability. Addressing this gap does not require abandoning precision or ISR capabilities, but it does require reordering priorities. This section outlines policy-relevant implications and recommendations across four domains: evaluation metrics, force posture planning, analytical practice, and governance design.

7.1 Reframe How Drone Defense Effectiveness Is Evaluated

7.1.1 Shift evaluation from point performance to uncertainty-cost dynamics

Decision-makers should move beyond metrics such as detection rates, intercept probabilities, or sortie counts as primary indicators of effectiveness. These measures capture short-term outcomes but obscure how quickly adversaries adapt and how exposure accumulates over time. More informative metrics assess how posture choices affect the marginal cost an adversary must incur to localize, time, or classify activity with sufficient confidence.

7.1.2 Treat observability as an endogenous variable

Analytical frameworks should model observability as a function of posture and behavior, not as a fixed input. Investments that marginally improve sensing but significantly increase predictability or exposure should be discounted accordingly. Evaluation processes should explicitly account for the trade-off between short-term clarity and long-term robustness.

7.2 Adopt Portfolio-Inspired Posture Planning

7.2.1 Plan force posture as a mixture rather than a configuration

Force planning should move away from optimizing for a single anticipated scenario, such as persistent patrols, tightly synchronized swarms, or continuous ISR, and instead design posture mixtures that deliberately balance visibility, survivability, mobility, and temporal randomization. This approach acknowledges regime uncertainty and reduces reliance on brittle assumptions about future conditions.

7. Policy Implications and Recommendations

7.2.2 Accept bounded inefficiency as a strategic choice

Robust posture mixtures will often sacrifice peak effectiveness in any one scenario. Policymakers should treat this not as a failure of optimization, but as the cost of resilience under uncertainty. Planning guidance should explicitly recognize that endurance and adaptability are strategic objectives in their own right.

7.3 Update Analytical and Wargaming Practices

7.3.1 Replace deterministic schedules with hazard-based assumptions

Operational analysis and wargaming should adopt intensity- or hazard-based representations of activity rather than fixed schedules. This shift better captures timing uncertainty, adversary adaptation, and escalation dynamics in contested environments.

7.3.2 Incorporate belief dynamics into scenario analysis

Scenarios should explicitly model how adversaries update beliefs over time and how posture choices affect the speed and cost of belief convergence. This allows analysts to compare not only outcomes, but the durability of those outcomes under sustained interaction.

7.4 Address Governance, Verification, and Escalation Risks

7.4.1 Move from static verification to process-based accountability

Where spatiotemporal verification is infeasible, governance frameworks should emphasize traceable decision authority, auditable processes, and clearly defined escalation thresholds rather than physical inspection alone. Accountability should focus on how decisions are made, not solely on what is observed.

7.4.2 Develop signal-management norms to reduce miscalculation

States should explore norms governing the interpretation and manipulation of operational signals—such as timing variability and presence patterns—to reduce the risk that indeterminacy is misread as imminent escalation. These norms need not constrain capability, but should clarify behavioral expectations in sensitive contexts.

7. Policy Implications and Recommendations

7.4.3 Treat indeterminacy as a strategic stability variable

Strategic stability assessments should explicitly consider how persistent uncertainty affects crisis behavior, escalation thresholds, and deterrence signaling. Transparency should no longer be assumed as the default stabilizing condition; instead, stability should be evaluated in terms of how uncertainty is managed and governed.

7.5 Summary

Taken together, these recommendations underscore a broader shift in policy thinking. Drone warfare should not be evaluated solely by how precisely forces can see or strike, but by how effectively they can sustain operational relevance, manage escalation risk, and preserve governance under contested observability. Uncertainty is no longer merely an obstacle to be overcome—it is a strategic variable that must be actively shaped.

The concluding section draws these threads together, emphasizing what decision-makers should optimize for, what they should avoid, and why robustness under indeterminacy has become a defining feature of modern drone warfare.

8. Conclusion

8. Conclusion

Drone warfare is undergoing a structural transition. What once appeared to be a problem of detection, interception, and synchronization has increasingly revealed itself as a problem of persistence under uncertainty. In contested environments, the decisive constraint is no longer how precisely forces can see or strike at any given moment, but how long they can sustain operational relevance as adversaries adapt to sensing, timing, and signaling patterns. Precision, long treated as the primary source of advantage, has become fragile when the costs of maintaining it escalate faster than its operational returns.

This report has argued that the central strategic variable in drone warfare is not uncertainty itself, but the economics of uncertainty reduction. Attempts to eliminate spatiotemporal indeterminacy, through continuous ISR, tight synchronization, or concentrated force employment, often expose forces to countermeasures that rapidly erode effectiveness. By contrast, postures that impose rising marginal costs on adversary efforts to localize, time, and classify activity can preserve survivability and influence even when absolute ambiguity is limited. Advantage lies not in perfect concealment, but in shaping how expensive "knowing enough" becomes for the opponent.

The portfolio-inspired Permanent Operational Configuration (POC) framework offers a way to conceptualize this shift. By treating force posture as a mixture of mutually tensioned presence modes rather than a single optimized configuration, POC explicitly trades peak performance in any one scenario for robustness across regime uncertainty. This is not a retreat from effectiveness, but a redefinition of it: success is measured by endurance, adaptability, and sustained pressure rather than momentary dominance.

Importantly, this reframing does not imply abandoning precision, deception, or control. Precision remains valuable—but only when embedded within postures that prevent adversaries from cheaply exploiting predictability. The objective is controlled indeterminacy: uncertainty sufficient to disrupt adversary planning without undermining governance, accountability, or strategic restraint.

The implications extend beyond operational analysis to questions of strategic stability and governance. Force structures optimized for spatiotemporal indeterminacy challenge traditional assumptions about transparency and verification. When presence and timing cannot be readily observed or interpreted, stability can no longer rest solely on static disclosure or physical inspection. Instead, it depends on

8. Conclusion

institutional mechanisms capable of managing uncertainty: process-based accountability, audit-by-design architectures, and norms governing the interpretation and manipulation of operational signals. Without such mechanisms, persistent indeterminacy risks increasing misperception and inadvertent escalation, even as it enhances survivability.

For policymakers, the core lesson is clear. Drone warfare should not be evaluated primarily through intercept rates, platform counts, or synchronization capacity. These indicators capture short-term outcomes, not sustainability. More meaningful assessments focus on how posture choices affect exposure over time, how adversary sensing costs scale, and how uncertainty interacts with escalation dynamics. In this sense, uncertainty management becomes a core strategic function—analogous to deterrence stability or second-strike survivability in earlier eras.

Modern drone warfare rewards neither perfect foresight nor maximal opacity. It rewards robustness under indeterminacy: the ability to remain effective, credible, and governable when prediction fails. Strategies that recognize and manage this reality will be better positioned to endure prolonged competition, reduce escalation risk, and preserve strategic control in an increasingly contested ISR environment.

Selected References and Sources

Selected References and Sources

- **Drone Warfare and Contested ISR**

Biddle, S. (2023). *The military lessons of the Russia–Ukraine war*. Council on Foreign Relations.

Center for Strategic and International Studies. (2023). *Drone warfare and air defense lessons from Ukraine*. CSIS. https://www.csis.org

Kofman, M. (2022). *Not built for purpose: The Russian military's ill-fated force design*. *War on the Rocks*. https://warontherocks.com

RAND Corporation. (2022). *Air and missile defense in contested environments*. RAND Corporation.

Royal United Services Institute. (2023). *The return of industrial warfare*. Royal United Services Institute.

Watling, J., & Reynolds, N. (2023). *Meatgrinder: Russian tactics in the second year of the Ukraine war*. Royal United Services Institute.

Watts, B. (2004). *Six decades of guided munitions and battle networks*. Center for Strategic and Budgetary Assessments.

- **Strategy, Uncertainty, and Adaptation**

Betts, R. K. (2002). Surprise despite warning. *Political Science Quarterly, 117*(1).

Freedman, L. (2013). *Strategy: A history*. Oxford University Press.

Gray, C. S. (2015). *The future of strategy*. Polity Press.

Jervis, R. (1976). *Perception and misperception in international politics*. Princeton University Press.

Schelling, T. C. (1966). *Arms and influence*. Yale University Press.

Taleb, N. N. (2010). *The Black Swan: The impact of the highly improbable*. Random House.

Van Creveld, M. (1985). *Command in war*. Harvard University Press.

- **Observability, Information, and Control**

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Boyd, J. R. (1987). *A discourse on winning and losing* (Unpublished briefing slides).

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors, 37*(1), 32–64.

Klein, G. (2017). *Sources of power: How people make decisions*. MIT Press.

Rumsfeld, D. (2002). *Department of Defense news briefing on uncertainty and intelligence* (Archival transcript). U.S. Department of Defense.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*.

- **Governance, Verification, and Accountability**

Acton, J. M. (2020). Escalation through entanglement. *International Security, 43*(4).

Crootof, R. (2022). War, responsibility, and killer robots. *Yale Journal of International Law, 45*(2).

Lewis, J. G., & Schultz, K. A. (2022). *Artificial intelligence and escalation*. Center for Strategic and International Studies.

Organisation for Economic Co-operation and Development. (2023). *Framework for the classification of AI systems*. OECD.

Stockholm International Peace Research Institute. (2022). *Autonomy in weapon systems: Implications for arms control*. SIPRI.

United Nations General Assembly. (2023). *Report of the Group of Governmental Experts on lethal autonomous weapon systems*.

United Nations Institute for Disarmament Research. (2023). *Responsible artificial intelligence in the military domain*. UNIDIR.

- **Related EPINOVA Research**

Wu, S.-Y. (2025a). *From detection to depletion: Cost-exchange limits in the Russia–Ukraine drone war* (Research Report No. EPINOVA–2025–01–RR). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.5281/zenodo.18036790

Wu, S.-Y. (2025b). *Permanent presence under uncertainty: A partially observable game-theoretic framework for unmanned systems, cost–frequency dynamics, and strategic stability*. EPINOVA. https://doi.org/10.5281/zenodo.18081107

Wu, S.-Y. (2025). *From predictive control to robustness: Why indeterminacy outperforms precision in contested ISR environments* (Policy Brief No. EPINOVA–2025–PB–03). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.5281/zenodo.18110856

Appendix A: Methods

Appendix A: Methods

This appendix clarifies the methodological choices, modeling assumptions, and analytical boundaries underlying the framework developed in this report. Its purpose is to ensure interpretability, prevent overextension of the findings, and clearly distinguish conceptual analysis from operational or tactical prescription.

The methodological foundation of this report is grounded in **chaos and nonlinear systems theory**, rather than physical indeterminacy or measurement limits. Uncertainty is treated as an emergent property of adaptive interaction, endogenous observability, and sensitivity to initial and structural conditions within a competitive system.

A.1 Analytical Scope, Assumptions, and Boundary Conditions

A.1.1 Analytical Assumptions and Boundary Conditions

This analysis assumes persistent partial observability as the normal condition of conflict.

Operational interaction is therefore modeled under conditions in which neither side possesses complete, continuous, or costless access to the adversary's operational intelligence.

Within a chaos-theoretic perspective, this assumption reflects a system in which outcomes remain highly sensitive to informational structure, adaptive response, and interaction history, even when underlying rules are deterministic.

The analysis explicitly excludes the degenerate case of complete and costless disclosure of adversary operational intelligence. This exclusion is not an idealization, but a boundary condition: such a scenario corresponds to systemic breakdown or terminal dominance, rather than sustained competition. Under conditions of total informational transparency, the system collapses into a trivially predictable regime; uncertainty management, posture robustness, adversarial adaptation, and nonlinear divergence cease to be analytically meaningful, and the dynamics examined in this report no longer apply.

A.1.2 Analytical Purpose and Level of Abstraction

The framework presented in this report is **strategic and analytical**, not tactical or operational. It is designed to:

- Examine how uncertainty, observability, and adaptation interact in drone warfare as a **nonlinear, adaptive system**;

Appendix A: Methods

- Compare classes of force posture and evaluation logic under conditions of **regime uncertainty and path dependence**;
- Inform policy-level decisions related to assessment metrics, force-planning principles, and governance design.

Accordingly, the analysis does not specify platform capabilities, force quantities, deployment patterns, command-and-control architectures, or execution procedures. All variables are intentionally abstracted to preserve generality and to avoid encoding operational guidance.

A.2 Modeling Framework

A.2.1 Partially Observable Stochastic Game (POSG)

A formal treatment of the POSG framework, including belief dynamics and cost–frequency interactions, is presented in Wu (2025b). In this report, the POSG formulation is used only as an analytical framing to clarify information–cost trade-offs and adaptive interactions.

Drone warfare is modeled as a **partially observable stochastic game** between two adaptive actors (hereafter “Blue” and “Red”) operating within a stochastic environment.

- States represent latent system conditions, such as spatial distributions, activity regimes, and environmental constraints.
- Observations are imperfect and action-dependent; observation quality is endogenously shaped by the behavior of both sides.
- Actions influence not only immediate outcomes but also future observability, belief evolution, and adaptive response.

Within a chaos-oriented analytical lens, this modeling choice captures how **deterministic decision rules can nonetheless generate divergent trajectories** due to feedback, adaptation, and sensitivity to information structure. Information is therefore not treated as exogenous. Efforts to improve sensing, control, or coordination reshape exposure, adaptation incentives, and long-term system stability.

A.2.1.1 Structural Formalization

For clarity, the interaction described in this report can be represented at a high level as a partially observable stochastic game

Appendix A: Methods

$$G = \langle S, A_B, A_R, O_B, O_R, T, Z \rangle,$$

where:

- S denotes latent system states capturing spatial, temporal, and operational regimes;
- A_B, A_R represent action spaces of Blue and Red actors;
- O_B, O_R denote imperfect observation spaces;
- T is a stochastic transition function influenced by joint actions;
- Z is an observation function, whose quality is endogenously shaped by posture and behavior.

Each actor maintains a belief distribution $b_i(s)$ over latent states, updated through imperfect observations. The analytical focus of this report is not on equilibrium solutions or optimal policies, but on how posture choices alter the **rate and marginal cost of belief convergence** over time.

No functional forms, payoff specifications, equilibrium concepts, or solution methods are assumed. This formalization is illustrative and serves only to clarify structural relationships, not to enable computation or operational application.

A.2.2 Belief Dynamics and Uncertainty

In this report, belief refers to an analytical construct representing an actor's probabilistic assessment of latent system states under conditions of partial and contested observability. Belief is modeled as a probability distribution over possible states of the operational environment, updated as imperfect observations are received. It is not a psychological construct, a subjective conviction, or an assessment of informational accuracy.

Belief is treated as a state variable that mediates between observability, action, and exposure over time, enabling comparative analysis of robustness under regime uncertainty. The analytical focus is not on whether beliefs are correct, but on how rapidly—and at what marginal cost—belief convergence occurs under adversary adaptation. Strategic advantage is therefore evaluated in terms of how posture choices shape the cost trajectory of belief refinement, rather than the elimination of uncertainty or the achievement of perfect information.

Uncertainty is represented through adversary belief distributions over latent system states. The analytical focus is on:

Appendix A: Methods

- The **rate and cost of belief convergence**, rather than absolute uncertainty levels;
- How posture choices shape the **nonlinear scaling** of uncertainty reduction costs.

Information entropy is used as a descriptive proxy for uncertainty but is not treated as an optimization objective. Consistent with chaos theory, the framework does not seek to eliminate uncertainty or maximize ambiguity. Instead, it examines how posture choices influence **divergence, convergence, and instability** in adversary belief dynamics.

From a nonlinear systems perspective, belief evolution may exhibit path dependence and sensitivity to initial informational conditions, such that small differences in observation structure or posture can generate disproportionate divergence in long-run trajectories, even under deterministic decision rules.

A.3 Portfolio-Inspired Posture Representation

A.3.1 Posture as a Mixture, Not a Configuration

Force posture is represented as a **posture-mixture vector** $\boldsymbol{w}$, indicating the relative emphasis placed on different presence modes. These modes describe functional roles, such as visibility, survivability, mobility, and temporal uncertainty, rather than discrete force elements or units, corresponding to the four POC presence modes defined in Section 3.

This abstraction enables comparison across posture classes while preserving the essential chaos-theoretic insight that system behavior depends on interaction structure rather than component optimization.

Formally, the posture-mixture vector can be written as:

$$\boldsymbol{w} = (\boldsymbol{w_1}, \boldsymbol{w_2}, \dots, \boldsymbol{w_n}), \qquad \sum_{i=1}^{n} \boldsymbol{w_i} = 1 \ , \boldsymbol{w_i} \geq \boldsymbol{0}$$

where each component $\boldsymbol{w_i}$ represents the relative emphasis assigned to a distinct presence mode.

This normalization is interpretive rather than operational: the vector does not encode force size, platform allocation, or resource quantities, but captures the structural composition of posture emphasis across mutually tensioned modes.

Appendix A: Methods

A.3.2 Robustness Under Regime Uncertainty

The portfolio analogy is employed to capture robustness under **regime uncertainty and nonlinear response**, not to import financial optimization techniques directly. The framework assumes that:

- Regime uncertainty is persistent and cannot be eliminated through improved prediction alone due to **sensitivity to adaptive feedback**;
- Bounded inefficiency may be acceptable when it stabilizes system behavior and prevents rapid collapse into predictable regimes.

No assumptions are made regarding optimal posture weights. The framework is comparative rather than prescriptive.

A.4 Integration with Cost–Distance–Frequency Models

A.4.1 Frequency as Hazard, Not Schedule

Operational frequency is treated as **event intensity (hazard)** rather than as a deterministic schedule. This reflects empirical patterns in contested environments, where timing regularity is actively exploited by adversaries and can induce **phase locking or synchronization collapse**.

The framework does not assume a specific stochastic process (e.g., Poisson or Hawkes). Hazard is used as a generic intensity concept to support qualitative and comparative analysis of timing instability.

A.4.2 Presence Efficiency and Capability Constraints

Two abstract modifiers are introduced:

- **Presence Efficiency**, capturing how effectively a given posture mixture translates presence into operational relevance under environmental conditions;
- **Capability Constraints**, representing aggregate limits related to resources, sustainment, and organizational capacity.

Both constructs are intentionally underspecified to allow adaptation across different strategic and institutional contexts and to preserve generality in nonlinear analysis.

Appendix A: Methods

A.5 Explicit Exclusions

To avoid misinterpretation, the framework does not model:

- Tactical employment, targeting, or engagement sequences;
- Platform-specific performance characteristics;
- Quantitative force sizing or deployment density;
- Real-time command-and-control architectures;
- Optimization of lethality or strike effectiveness.

Any operational or tactical application would require additional assumptions, empirical inputs, and context-specific validation beyond the scope of this report.

A.6 Validity, Generalizability, and Interpretation

A.6.1 Conceptual Validity

The framework supports **comparative reasoning under nonlinear interaction**, not prediction. Its validity lies in clarifying trade-offs, feedback structures, and cost dynamics, not in forecasting conflict outcomes or engagement success.

A.6.2 External Validity

The findings are most applicable to:

- Highly contested sensing and communication environments;
- Prolonged competitive interactions characterized by adaptation and counter-adaptation;
- Contexts in which **small structural changes can generate disproportionate strategic effects**.

They may be less applicable in permissive environments or scenarios involving overwhelming capability asymmetries.

A.6.3 Governance Interpretation

Governance implications discussed in the report are normative and institutional, not enforcement-ready. Concepts such as auditability, accountability, and signal-management norms are presented as design directions rather than finalized regulatory instruments.

Appendix A: Methods

A.7. Summary

This appendix underscores a central analytical choice of the report: drone warfare is treated as a chaotic, adaptive system in which uncertainty arises endogenously from interaction, observability, and feedback—not as a problem of physical indeterminacy or technical optimization.

The abstractions employed are deliberate. They are designed to support strategic reasoning and policy evaluation by clarifying structural trade-offs and cost dynamics, while explicitly avoiding tactical, operational, or system-design prescription.

Appendix B: Limitations and Scope

Appendix B: Limitations and Scope

This appendix defines the analytical scope, limitations, and interpretive boundaries of the report. Its purpose is to clarify what the framework does—and does not—claim, and to prevent misinterpretation of the analysis as operational, tactical, predictive, or system-design guidance.

B.1 Scope of Analysis

This report examines drone warfare at the strategic and analytical level, treating it as a chaotic, adaptive system characterized by contested observability, nonlinear interaction, and adversarial adaptation. The focus is on uncertainty management, the economics of observability, and posture robustness under regime uncertainty. The framework is intended to inform policy evaluation, force-planning principles, and governance discussions.

B.2 Non-Operational and Non-Tactical Boundary

The report does not specify platforms, force size, deployment density, command-and-control architectures, rules of engagement, targeting processes, or execution procedures. Core constructs, such as posture mixtures, belief dynamics, endogenous observability, and hazard-based timing, are intentionally abstract to preserve generality and to avoid encoding operational guidance. Any translation into practice would require additional assumptions, empirical calibration, and context-specific validation beyond the scope of this report.

B.3 No Predictive or Optimization Claims

The framework is comparative rather than predictive. It does not forecast conflict outcomes, engagement success rates, escalation probabilities, or system performance metrics, nor does it optimize lethality, efficiency, or operational effectiveness. Consistent with a chaos-theoretic perspective, the analysis emphasizes sensitivity, feedback, and divergence over point prediction. Its purpose is to clarify structural trade-offs, cost dynamics, and robustness properties under uncertainty—not to predict specific trajectories or outcomes.

B.4 Context Dependence and Applicability

The findings are most applicable to highly contested sensing and communication environments characterized by sustained adaptation and counter-adaptation over time,

Appendix B: Limitations and Scope

particularly in prolonged competitive interactions where observability, timing, and signaling are actively contested. The framework may be less applicable in permissive environments, short-duration conflicts, or scenarios involving overwhelming and persistent capability asymmetries where uncertainty collapses rapidly and nonlinear dynamics play a reduced role.

B.5 Analogy and Conceptual Boundary Conditions

References to portfolio logic, stochastic games, hazard models, or uncertainty principles are conceptual and structural analogies, not direct transfers of financial optimization techniques, physical uncertainty relations, or formal control guarantees. These analogies organize reasoning about robustness, adaptation, and cost scaling in complex systems. They do not imply formal equivalence, mathematical isomorphism, deterministic guarantees, or physical measurement limits as explanatory mechanisms.

B.6 Governance Interpretation Boundaries

Governance implications discussed in the report, such as verification hardness, process-based accountability, and signal-management norms, are normative and exploratory. They identify emerging institutional challenges and potential design directions rather than proposing ready-to-implement regulatory instruments, compliance regimes, or enforcement mechanisms. These discussions are intended to frame policy questions, not to prescribe legal standards or operational rules.

B.7 Bottom Line

This report should be read as a conceptual and analytical framework for understanding uncertainty management in drone warfare under conditions of contested observability and nonlinear adaptation. It is not a blueprint for system design, operational employment, tactical execution, or predictive assessment.

Appendix C: Abbreviations and Glossary

Appendix C: Abbreviations and Glossary

C.1 Abbreviations

Abbreviation	Full Term
AI	Artificial Intelligence
C2	Command and Control
CDF	Cost–Distance–Frequency (analysis/model family)
EW	Electronic Warfare
G	Game (formal notation, used in POSG definition)
ISR	Intelligence, Surveillance, and Reconnaissance
IHL	International Humanitarian Law
LAWS	Lethal Autonomous Weapon Systems
MVA	Minimum Viable, Auditable (framework)
POC	Permanent Operational Configuration
POSG	Partially Observable Stochastic Game
PR	Policy Report (as used in report numbering)
ROE	Rules of Engagement
UAS	Unmanned Aerial System(s)
UNGA	United Nations General Assembly
UNIDIR	United Nations Institute for Disarmament Research

Appendix C: Abbreviations and Glossary

C.2 Glossary

Audit-by-Design
A governance design direction in which systems generate verifiable records of authorization, decision pathways, and constraint adherence, enabling post hoc accountability without requiring continuous external visibility.

Belief (Belief State / Belief Distribution)
A probability distribution over latent system states maintained by an actor under partial observability. The report emphasizes the rate and cost of belief convergence rather than point prediction accuracy.

Belief Convergence
The process through which an actor's belief distribution narrows over time as observations accumulate. The strategic focus is on whether further convergence can be achieved cheaply or only at escalating marginal cost.

Contested ISR Environment
A subset of contested observability conditions emphasizing pressure on intelligence, surveillance, and reconnaissance functions, including denial, deception, spectrum contestation, and counter-ISR dynamics.

Contested Observability
An environment in which sensing, communication, and timing are actively targeted, degraded, or exploited by both sides. Observability is endogenous to behavior and adaptation rather than an exogenous given.

Controlled Indeterminacy
A posture condition in which uncertainty persists in strategically relevant dimensions (timing, localization, classification, representativeness), preventing cheap belief convergence while maintaining governance and escalation control.

Cost–Distance–Frequency (CDF) Analysis
A family of assessment logics relating operational costs, distance-related constraints, and activity rates. This report modifies the frequency component by treating it as hazard under regime uncertainty.

Endogenous Observability
The principle that observation quality depends on both sides' actions. Efforts to improve sensing or control can increase exposure and invite countermeasures, reshaping future observability.

Appendix C: Abbreviations and Glossary

C.2 Glossary

Event Intensity (Hazard-Based Timing)
A conceptual representation of activity likelihood per unit time, used to capture non-deterministic timing under regime uncertainty. Hazard is treated as an analytical intensity concept, not an operational schedule or specific stochastic process.

Exposure Cost
The cumulative strategic cost of revealing operational patterns, thresholds, or dependencies through repeated sensing, communication, or regular timing, shaping an adversary's future inference capability.

Indeterminacy (Spatiotemporal Indeterminacy)
The persistent inability to convert observation into reliable actionable certainty at acceptable and sustainable cost, even when partial awareness exists. Indeterminacy is treated as a structural outcome of measurement–exposure trade-offs and adversary adaptation.

Measurement–Exposure Trade-off
The structural relationship in which increasing sensing, control, or persistence improves short-term clarity while increasing signatures, predictability, and long-run vulnerability through cumulative exposure.

Partially Observable Stochastic Game (POSG)
A formal class of models for strategic interaction under partial observability and stochastic transitions, where players update beliefs based on imperfect observations and actions affect both outcomes and information. In this report, POSG is used as a conceptual framing rather than a computational model.

Permanent Operational Configuration (POC)
A portfolio-inspired framework treating force posture as a mixture of presence modes rather than a single optimized configuration. POC is an evaluative and planning lens at the strategic level, not a doctrine or operational design template.

POC Presence Modes
Analytical categories describing functional posture roles under contestation, including deterrent presence, survivable reserve, mobile uncertainty, and temporal randomization. These modes are descriptive, not prescriptive.

Appendix C: Abbreviations and Glossary

C.2 Glossary

Posture Mixture (Posture-Mixture Vector)
An analytical representation of how emphasis is distributed across presence modes within a POC. The mixture is interpretive rather than operational and does not encode force size or deployment plans.

Process-Based Accountability
A governance approach emphasizing traceable authority, documented constraints, and auditable decision processes rather than full spatiotemporal observability, proposed as a stability mechanism under contested observability.

Regime Uncertainty
Uncertainty over which interaction regime is dominant and how regimes shift over time under adaptation, motivating portfolio-style robustness rather than single-scenario optimization.

Robustness (Regime-Robustness)
The capacity to sustain operational relevance and control across regime shifts and adaptive countermeasures, even when peak effectiveness in any single scenario is not maximized.

Verification Hardness
A condition in which capabilities and deployments may be observable but cannot be reliably interpreted or verified to traditional arms-control standards, as observability fails to produce actionable interpretability.

Citation Recommendation

Wu, S.-Y. (2025). *Nonlinear Uncertainty in Drone Warfare: Why Indeterminacy Outperforms Precision in Contested ISR Environments* (Policy Report No. EPINOVA–2025–PR–01). Global AI Governance and Policy Research Center, EPINOVA LLC. https://doi.org/10.5281/zenodo.18111066

GLOBAL AI GOVERNANCE RESEARCH CENTER

EPINOVA