

Permanent Presence under Uncertainty: A Partially Observable Game-Theoretic Framework for Unmanned Systems, Cost–Frequency Dynamics, and Strategic Stability

Author: Dr. Shaoyuan Wu

ORCID: <https://orcid.org/0009-0008-0660-8232>

Affiliation: Global AI Governance and Policy Research Center, EPINOVA LLC

Date: 2025-12-29

Abstract

The operational advantage of unmanned systems is increasingly derived not from superior precision or predictive accuracy, but from their ability to sustain presence while resisting low-cost opponent inference. This paper develops a formal framework that conceptualizes permanent presence as a strategic condition rather than a platform-centric capability, reframing it as a portfolio under uncertainty. Building on a Partially Observable Stochastic Game (POSG) formulation, the analysis models the structural tension between observability, operational tempo, and survivability, and extends a Cost–Distance–Frequency (CDF) framework by treating frequency as an adversary-facing random variable rather than a controllable constant. The paper further explores the implications of this shift for strategic stability, verification, and AI-enabled arms control, arguing that uncertainty itself has become a central stabilizing—and destabilizing—factor in contemporary unmanned conflict.

1. Introduction

Contemporary analyses of unmanned and AI-enabled military systems overwhelmingly focus on moments of engagement: detection, targeting, interception, and strike. Whether framed in terms of kill chains, decision speed, or cost-exchange ratios, effectiveness is typically assessed at the point of contact. Implicit in this perspective is the assumption that strategic value is generated primarily through episodic superiority—brief intervals in which one side sees more clearly, reacts faster, or strikes harder than the other (Schelling, 1966; Jervis, 1976).

This paper begins from a different premise: in many contemporary operational environments, the decisive variable is not episodic superiority, but continuous presence as a strategic condition.

In contested domains characterized by saturation, attrition, and rapid adaptation, unmanned systems increasingly derive their strategic value not from single decisive

actions, but from their ability to remain—to exert pressure, shape adversary expectations, and constrain decision space over time. Presence, rather than engagement, becomes the dominant currency of influence.

However, presence in such environments is inherently fragile. Any persistent pattern—whether spatial, temporal, or electromagnetic—can be observed, inferred, and eventually targeted. As a result, the challenge is no longer how to deploy unmanned systems more precisely, but how to sustain presence without collapsing into predictability.

This leads to a critical conceptual shift. Permanent presence does not mean continuous activity, fixed deployment, or constant signaling. Instead, it refers to the **credible persistence of capability under conditions of adversarial observation**, where neither location nor timing can be cheaply resolved. Presence, in this sense, is less a physical condition than an informational one, shaped by how adversaries perceive, infer, and allocate resources in response.

The strategic problem, therefore, is not how to maximize visibility or control, but how to design systems and doctrines such that adversaries face increasing marginal costs when attempting to determine where, when, or whether a capability will manifest. The objective is not perfect concealment—an unrealistic goal—but the maintenance of an uncertainty structure that resists rapid collapse.

This paper terms this condition **Permanent Presence under Uncertainty**.

To formalize this concept, the analysis moves beyond platform-centric or engagement-centric models and instead treats presence as a portfolio across multiple modes of existence—visible and latent, mobile and static, regular and randomized. The interaction between these modes is inherently adversarial and informational, making it unsuitable for deterministic or fully observable frameworks. Accordingly, the paper adopts a Partially Observable Stochastic Game (POSG) formulation to capture how both sides act, observe, and adapt over time.

By reframing unmanned operations around permanent presence rather than episodic action, the paper advances three claims. First, survivability and effectiveness are increasingly governed by uncertainty management rather than precision maximization. Second, frequency and tempo should be modeled as adversary-facing random variables rather than controllable constants, requiring a reconfiguration of standard cost–distance–frequency models. Third, this shift has direct implications for strategic stability, verification, and AI governance, as systems optimized for non-localization and non-timing challenge existing arms control and accountability paradigms.

Permanent presence is analytically distinct from forward deployment, continuous ISR, or deterrence-by-denial. It does not require constant activity, visibility, or contact, but rather the credible persistence of capability whose spatial and temporal manifestation cannot be cheaply resolved by an adversary.

2. A POSG Model of Permanent Unmanned Presence

2.1 Core Intuition and Strategic Logic

The analogy to an “uncertainty principle” in unmanned conflict does not imply epistemic blindness. Rather, it reflects a structural tradeoff:

- Increased observability and coordination improve short-term efficiency but raise exposure.
- Reduced observability and randomized tempo degrade efficiency but enhance survivability.
- Adversaries respond not statically, but through adaptive investment in sensing and inference.

These dynamics naturally map onto a Partially Observable Stochastic Game, in which both sides act under uncertainty and can modify the information structure through their actions (Shapley, 1953; Hansen et al., 2004). A single-agent Partially Observable Markov Decision Process (POMDP) formulation is insufficient, as uncertainty here is not exogenous noise but the outcome of strategic interaction, where both sides actively manipulate observability and inference costs (Kaelbling et al., 1998).

2.2 State Space

Let the true system state at time t be:

$$S_t = (X_t, M_t, \Theta_t, \Omega_t)$$

where:

- X_t : spatial state (location or distribution, partially hidden from the adversary)
- M_t : mission or tempo state (activity level, operational windows, rhythm)
- Θ_t : adversary capability and intent state (threat intensity, strike timing, resource levels)
- Ω_t : environmental state (weather, terrain masking, spectrum background, logistical constraints)

The central structural tension lies between X_t and M_t : fixed and predictable rhythms facilitate localization and timing attacks, while stochastic rhythms reduce efficiency but increase persistence.

2.3 Observation Model

Neither player observes S_t directly. Instead, observations are generated as:

$$O_t \sim Z(S_t, a_t^{Blue}, a_t^{Red})$$

Crucially, the observation function Z depends on **both sides' actions**. Measures that improve internal coordination, such as tighter communication or synchronized behavior, may simultaneously increase detectability. Thus, information quality is endogenous to strategy, not an exogenous parameter.

2.4 Adversarial Belief State and Entropy

The adversary maintains a belief distribution over states:

$$b_t^{Red}(S) = P(S_t = S \mid O_{1:t}, a_{1:t-1})$$

Strategic success is not defined by maximizing or freezing belief entropy indefinitely, which is unrealistic. Instead, the objective is to ensure that **reductions in entropy require disproportionate adversarial investment**.

Entropy provides a natural metric:

$$H(b_t^{Red}) = - \sum_S b_t^{Red}(S) \cdot \log b_t^{Red}(S)$$

Entropy is used here as a generic, model-agnostic proxy for adversarial uncertainty, without loss of generality regarding alternative divergence or dispersion measures (Cover & Thomas, 2006). More broadly, the use of entropy reflects the premise that information acquisition and uncertainty reduction are costly and subject to strategic tradeoffs (Sims, 2003).

Permanent presence, in this sense, is achieved when adversarial belief convergence becomes increasingly expensive.

2.5 From Permanent Investment to Permanent Configuration

Drawing inspiration from permanent portfolio theory, unmanned presence can be conceptualized as a weighted configuration over abstract “existence modes”:

$$w = (w_1, w_2, w_3, w_4), \sum_i w_i = 1$$

where:

- w_1 : deterrent presence (visible, expectation-shaping)
- w_2 : survivable reserve (hidden, second-capability persistence)
- w_3 : mobile uncertainty (spatial ambiguity)

- w_4 : temporal randomization (rhythmic ambiguity)

The optimization objective is long-term survivability and sustained effect rather than single-episode payoff:

$$\max_{\pi^{Blue}} E\left[\sum_{t=1}^T \gamma^t \cdot (U_t^{Blue} - \lambda \cdot L_t^{Blue})\right]$$

where utility captures abstract mission effects, losses include exposure and depletion, and discounting reflects the value of continued existence.

3. Extending the Cost–Distance–Frequency Framework

3.1 Frequency as an Uncertain Variable

Traditional CDF models treat frequency as a controllable operational parameter. Under adversarial inference, frequency should instead be modeled as a random variable from the opponent’s perspective:

$$F_t \sim D(\mu_t, \sigma_t; \text{regime})$$

Here, σ_t captures temporal uncertainty. Increasing variance, rather than mean intensity, becomes a strategic lever.

3.2 Adversarial Reconnaissance Investment

Adversaries may invest resources r_t to reduce uncertainty:

$$\sigma_{t+1} = \sigma_t \cdot e^{-\kappa r_t} + \eta$$

The irreducible lower bound η represents a structural uncertainty floor. Strategic design aims not to eliminate predictability, but to ensure diminishing returns on adversarial investment.

3.3 Incorporating Belief Entropy into Operational Value

Operational value can be reformulated as:

$$\text{Value} = V(D, F) - C - \alpha \cdot \text{Exposure} + \beta \cdot H(b_t^{Red})$$

Here, adversarial uncertainty itself becomes a positive strategic output.

3.4 Presence Efficiency, Constraints, and Risk

Presence Efficiency:

$$PE_t = g(w, \Omega_t) \in [0,1]$$

Capability Constraint:

$$K(w) \leq K_{max}$$

Expected Risk:

$$ER_t = E[L_t | b_t^{Red}, F_t]$$

Together, these elevate CDF from an engineering cost model to a stochastic control framework under strategic opposition.

4. Implications for AI Governance and Arms Control

4.1 Strategic Stability and Information Structure

Stability is no longer governed primarily by platform counts or deployment zones, but by persistence of second-capability presence, difficulty of verification, and heightened misperception risk. Stability thus becomes an information-structural problem rather than a purely material one (Glaser, 1997; Acton, 2013).

4.2 Verification under Uncertainty

Systems optimized for non-localization and non-timing undermine traditional verification regimes. Transparency may paradoxically reduce survivability, forcing governance to shift toward **process-based compliance** rather than static inventories.

4.3 Emerging Normative Directions

Four governance levers emerge:

- (1) **Risk-of-Miscalculation Controls**
- (2) **Assured Human Accountability** through traceable responsibility chains
- (3) **Signal Management Norms** regulating manipulative informational behavior
- (4) **Audit-by-Design** requirements embedded at the system-design stage

These challenges mirror those found in AI governance—traceability, auditability, and accountability—transposed from algorithms to adversarial autonomous presence (Kroll et al., 2017; Raji et al., 2020).

5. Conclusion

Unmanned systems are reshaping conflict not by perfecting prediction, but by resisting it. By formalizing permanent presence as a portfolio under partial observability and adversarial adaptation, this paper demonstrates how uncertainty

itself becomes a strategic resource. The resulting shift—from precision dominance to uncertainty management—has profound implications for operational design, cost modeling, and global governance regimes. Future arms control and AI governance frameworks must therefore grapple not only with what systems can do, but with how—and whether—they can be known.

Reference

- Acton, J. M. (2013). *Reclaiming strategic stability*. Carnegie Endowment for International Peace.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Wiley.
- Glaser, C. L. (1997). The security dilemma revisited. *World Politics*, 50(1), 171–201. <https://doi.org/10.1017/S0043887100012843>
- Goldblat, J. (2002). *Arms control: A guide to negotiations and agreements*. SAGE Publications.
- Hansen, E. A., Bernstein, D. S., & Zilberstein, S. (2004). Dynamic programming for partially observable stochastic games. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 709–715). AAAI Press.
- Jervis, R. (1976). *Perception and misperception in international politics*. Princeton University Press.
- Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134. [https://doi.org/10.1016/S0004-3702\(98\)00023-X](https://doi.org/10.1016/S0004-3702(98)00023-X)
- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165(3), 633–705.
- Osborne, M. J., & Rubinstein, A. (1994). *A course in game theory*. MIT Press.
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). ACM. <https://doi.org/10.1145/3351095.3372873>
- Schelling, T. C. (1966). *Arms and influence*. Yale University Press.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10), 1095–1100. <https://doi.org/10.1073/pnas.39.10.1095>
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3), 665–690. [https://doi.org/10.1016/S0304-3932\(03\)00029-1](https://doi.org/10.1016/S0304-3932(03)00029-1)

Van Benthem, J. (2011). *Logical dynamics of information and interaction*. Cambridge University Press.

Recommended Citation:

Wu, S.-Y. (2025). *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>