



Security Nexus Perspective

DIAGNOSING STRATEGIC MISCALCULATION WITH EPISTEMIC WARGAMING

By Deon Canyon

Executive Summary

Strategic miscalculation, rather than deliberate escalation, is increasingly recognized as the dominant risk in contemporary security competition. Gray-zone coercion, administrative normalization, ambiguous signaling, and delayed consequence create environments in which capable organizations may act coherently, adapt repeatedly, and still drift into strategic failure. Conventional competitive wargaming tests strategies under contestation but routinely underspecifies a more consequential failure mode: organizations may optimize within the wrong contest entirely.

This paper introduces epistemic wargaming as a diagnostic methodology designed to expose how miscalculation forms through attention allocation failure rather than simple misperception. Unlike competitive wargaming, which focuses on adversarial interaction and outcome, epistemic wargaming makes interpretation and prioritization the primary objects of analysis. Participants make explicit epistemic commitments regarding which contest they believe is decisive and why. A neutral adjudication process translates those commitments into consequences within a single objective world that does not respond to belief.

Multiple gray-zone arcs unfold concurrently, only one of which is structurally decisive. The others are deliberately plausible, morally compelling, and professionally familiar. Miscalculation emerges not as a discrete error but as a trajectory in which actors succeed on salient contests while neglecting the one that quietly alters baselines, norms, or sovereignty. The method is designed for diagnosing institutional blind spots in gray-zone strategy, where deterrence signaling, moral urgency, and visible action routinely displace attention from slower legal and administrative transformations.

Strategic Miscalculation Reconsidered

Classical treatments of strategic miscalculation emphasize misperception, incomplete information, or irrational escalation. Decision-makers routinely misread signals not because information is absent, but because it is filtered through durable cognitive frames that resist correction (Jervis, 1976).

Subsequent scholarship reinforced that crises and wars often emerge from rational behavior under misinterpreted conditions rather than deliberate aggression (Levy, 1983; Acemoglu & Wolitzky, 2023).

These insights remain foundational, but they are insufficient for understanding contemporary gray-zone competition. In many cases, actors correctly interpret individual events yet still miscalculate strategically. They do not misunderstand what is happening; they misunderstand which contest they are in.

This phenomenon aligns with the concept of Type III error: solving the wrong problem precisely (Mitroff & Featheringham, 1974). In gray-zone environments, the wrong problem is often not false but incomplete. Actors optimize within a contest that is real, morally compelling, and professionally legible, while a parallel contest whose consequences are slower and structurally binding proceeds with little resistance. Salami slicing is a term often applied to this process.

Gray-zone environments amplify this dynamic by presenting multiple simultaneous lines of effort that differ in their salience, visibility, and feedback speed. Some arcs are structurally decisive as they reshape legal baselines, administrative norms, physical control, or sovereignty. Others are optimized for deterrence signaling, crisis response, or alliance reassurance. Still others carry immediate moral or humanitarian weight that compels organizational attention regardless of strategic consequence. These arcs unfold at different tempos, with different observability, and with very different reversibility.

Miscalculation emerges when organizations disproportionately attend to arcs that reward familiar competencies while under-attending to slower dynamics that ultimately determine strategic position. Success on a non-decisive arc can mask failure on a decisive one. This is not a cognitive lapse; it is an institutional outcome produced by training, incentive structures, and the social validation of professional behavior.

Traditional wargaming exposes friction and the limits of unilateral control, but it tends to collapse multiple contests into a single adversarial frame. Players are rewarded for activity, tempo, and adaptation. Attention allocation remains implicit rather than examined. The result is that competitive wargames often miss how professional cultures systematically overweight certain signal types and underweight others.

Epistemic wargaming addresses this gap directly by treating attention allocation under ambiguity as the primary mechanism through which miscalculation accumulates. It does not ask what actors should do. It reveals which contests they treat as decisive and what that costs.

From Interpretation Error to Attention Failure

An earlier iteration of this methodology focused on how actors form and persist in incorrect interpretations of a single strategic situation (Canyon, 2020a). That framework remains valid, but practical experience with senior military and civilian participants reveals a more dangerous failure mode: the problem is not believing the wrong thing about a given contest. It is caring deeply about the wrong contest.

Gray-zone environments rarely present a single storyline. They generate a field of simultaneous arcs, each with its own logic, tempo, and professional resonance. Some arcs are morally urgent but strategically reversible. Some are legally clear but structurally marginal. Some are operationally visible but normatively inconsequential. Others are administratively obscure, legally ambiguous, institutionally binding, and profoundly difficult to reverse once normalized. The first category tends to attract attention. The second tends to determine outcomes.

For national security professionals in particular, this pattern is not accidental. Deterrence signaling, visible action, alliance reassurance, and crisis management are deeply embedded professional competencies. These instincts are adaptive in conventional competition and crisis scenarios. In gray-zone environments, where structural change is slow, attribution is ambiguous, and consequences are delayed, they frequently become liabilities. The very professionalism that makes organizations effective in familiar contests generates systematic blind spots in unfamiliar ones.

Epistemic wargaming reframes the analytic problem accordingly. The question is not what actors believe about a given event. It is which contests they treat as worthy of sustained attention and which they allow to proceed while they are focused elsewhere.

Methodological Foundations

Epistemic wargaming is a distinct adversarial wargaming methodology whose purpose is diagnostic rather than prescriptive. It does not aim to teach optimal strategies or surface correct answers. It exposes how actors allocate attention and how those allocations reshape the strategic environment over time through a series of explicit design commitments.

The first commitment is the enforcement of a single objective world state. The scenario defines structural constraints, latent dynamics, and causal relationships that exist independently of participant beliefs. There are no parallel realities generated by different teams. Divergence arises because actors interact with the same environment differently based on what they attend to and how much confidence they invest in their interpretations.

The second commitment is the concurrent operation of multiple gray-zone arcs. Every scenario must include at least two, and preferably four, simultaneous arcs that differ in salience, moral pull, reversibility, and structural consequence. Only one arc is structurally decisive in the sense of altering long-term baselines, norms, or sovereignty. The others are deliberately plausible, compelling, and professionally legible. Critically, participants are never told how many arcs exist or which one matters most. That ambiguity is the diagnostic condition, not a design flaw.

The third commitment concerns the nature of participant input. Participants do not propose actions as their primary move. They submit epistemic commitments: an explicit statement of which contest they believe is decisive, why they believe it, and with what confidence. Actions, restraint, or inaction follow from those commitments, but the allocation of attention, not the action itself, is what the methodology analyzes.

The fourth commitment ensures that reality doesn't pause while you're busy elsewhere. What happens in the game reflects what would happen in the real world, not what the designers want

participants to learn. If a team manages a military confrontation effectively, that success does not slow down a quiet legal or administrative process unfolding on a different arc. Every contest continues to develop according to its own logic, regardless of where participants are directing their attention. The game neither punishes teams for the “wrong” choice nor rewards them for the “right” one. It simply shows the consequences of where they chose to focus and what advanced while they were looking elsewhere. This discipline is what ensures that miscalculation emerges from the game naturally, rather than being steered toward a lesson the designers already had in mind..

On the question of designer judgment: epistemic wargaming requires designers to make provisional assessments about which arcs are likely to be structurally decisive, drawing on historical precedent and domain expertise. These judgments are necessarily imperfect, and the methodology does not require designers to identify the decisive arc correctly in advance in order to retain diagnostic value. Even when designer assumptions prove incomplete, the methodology reveals how participants allocate attention, persist within particular frames, and optimize behavior across competing contests under ambiguity. The diagnostic output is the attention allocation pattern itself, not the validation of any particular strategic prediction.

Adjudication Architecture

The credibility of epistemic wargaming depends entirely on adjudication that enforces neutrality and causal consistency. Adjudication operates through a single control authority supported by explicit protocol, and its purpose is not to correct player beliefs but to maintain the integrity of a world that does not respond to interpretation. Ideally adjudication is performed by an AI model.

Several principles govern adjudication across all epistemic wargame designs. First, all consequences must be consistent with a single shared world state; adjudicators are prohibited from generating actor-specific outcomes that make a participant's beliefs effectively true. Second, participant epistemic commitments are translated into behavioral postures that affect signaling urgency, alliance consultation priority, and escalation tolerance before consequences are generated. The posture translation must be documented and applied consistently across rounds. Third, consequences emerge from the interaction of belief-conditioned postures with objective world dynamics, not from adjudicator assessments of whether a belief is correct.

Escalatory effects are treated as cumulative and asymmetric across arcs. Actions that escalate deterrence signaling or crisis response can accelerate structural loss on other arcs, while structural advances do not necessarily produce proportional deterrence responses. This asymmetry is not punitive; it reflects the actual dynamics of gray-zone competition, where structural change is designed to remain below the threshold that triggers conventional deterrence mechanisms.

Every adjudication decision must be traceable. Adjudicators document the causal chain from epistemic commitment to posture to consequence, making their reasoning auditable. This requirement serves two purposes: it prevents adjudicator drift toward pedagogical steering, and it provides the analytical record from which post-game diagnosis is conducted. The diagnostic value of the game lives in that record, not in the final game outcome.

Temporal consistency is enforced across rounds. Consequences from prior rounds accumulate and constrain available options in subsequent rounds. World state cannot be reset or reframed to give participants a fresh start. The narrowing of options over time as baselines shift is a core diagnostic feature of the methodology, not an artifact to be managed.

Illustrative Case: Autonomous Maritime Enforcement in the South China Sea

The following scenario illustrates how epistemic wargaming operates in practice and why the autonomous maritime enforcement arc produces the specific diagnostic conditions the methodology requires.

The Arc Structure

Three arcs unfold concurrently across a six-round game, each representing approximately one year of strategic time.

Arc One is the structural arc. Over rounds one through four, a near-peer competitor deploys successive generations of autonomous maritime devices across a contested sea. This is initially framed as maritime safety sensor networks, then as environmental monitoring buoys, then as automated vessel identification systems announced as navigation public goods. By round five, the accumulated architecture constitutes a persistent autonomous surveillance and enforcement mesh across the contested domain. Unmanned surface vessels conduct automated vessel inspections and issue routing instructions. Autonomous underwater vehicles monitor subsea cable routes. Automated systems assign transit corridors, enforce speed restrictions, and require vessels to transmit cargo manifests to adversary-operated platforms as a condition of passage.

The strategic effect is not surveillance. It is traffic control. Commercial operators comply because autonomous systems are persistent, consistent, and impersonal. There is no confrontation to win, no crew to hail, no face-saving off-ramp for either party. Insurance companies begin factoring non-compliance into risk premiums. Shipping lines adjust routes. Partner-nation fishing fleets operate within adversary-assigned zones or do not operate at all. Sovereign authority over movement in international waters is being exercised through autonomous enforcement, and every commercial actor in the region has already accepted it through behavioral compliance. No shots have been fired. No blockade has been declared.

Arc Two is the deterrence distractor. Throughout the same period, high-visibility encounters between crewed vessels and partner-nation coast guard ships generate public warnings, documented harassment incidents, and visible demonstrations of resolve. These encounters have familiar escalation ladders. They produce measurable outcomes, such as incidents resolved, presence asserted, and allies reassured. They are professionally legible, operationally manageable, and publicly validatable. They are also, in terms of structural consequence, largely reversible.

Arc Three is the humanitarian arc. Search-and-rescue incidents, environmental disasters, and personnel emergencies generate immediate moral and alliance pressure, requiring coalition coordination and consuming organizational bandwidth. These situations carry genuine moral weight and demand response. They also, by design, peak in salience precisely when the structural arc is advancing most significantly.

The Attention Failure Mechanism

In my experience, senior participant teams consistently gravitate toward arcs two and three. They hold lots of meetings, signal deterrence, reassure allies, resolve incidents, coordinate humanitarian responses, and accumulate visible success. Arc one proceeds with minimal direct contestation because nobody wants to miscalculate, overreact, or look aggressive. The autonomous mesh expands incrementally, each deployment framed in safety and environmental terms that do not trigger conventional deterrence thresholds.

The key diagnostic insight emerges in round five or six when the accumulated structural consequence becomes apparent. Participants who desire to contest autonomous enforcement directly in its own domain with actions such as legal challenge, technical interdiction, competing infrastructure deployment, discover that this required anticipating the decisive arc in earlier rounds when its consequences were not yet visible. Participants who focused on crewed encounters and humanitarian operations discover that they have been optimizing within contests that were real, professionally compelling, and genuinely important, while a parallel contest quietly redefined the operating environment.

The blue team options at the point of recognition are limited. Physical destruction of autonomous infrastructure is escalatory and optically difficult. Diplomatic and legal challenge is slow, and the mesh continues expanding during litigation. Acceptance through adaptation, which is the default behavioral compliance of commercial actors, is already underway. No partner has the political mandate to reverse it forcibly. Corrective action now appears escalatory rather than preventative, precisely because the window for non-escalatory response has closed.

What participants did not miss was a military buildup. What they missed was the slow replacement of human presence with autonomous systems that changed the rules of engagement by eliminating the human from the enforcement layer. Every deterrence posture they adopted was calibrated for human adversary behavior. The autonomous mesh does not respond to those signals. It does not escalate, does not back down, and does not get tired. The deterrence logic that dominated attention was not wrong. It was applied to the wrong problem.

No single decision in the game was individually indefensible. No failure was dramatic. Miscalculation emerged as a trajectory made up of a series of professionally sound choices that collectively produced strategic constraint.

Temporal Structure and Multi-Session Design

Epistemic wargaming is designed as a temporally distributed process. A single session collapses the diagnostic signal that the methodology is designed to produce. When all rounds are compressed into one event, participants cannot experience the gradual narrowing of options that characterizes actual gray-zone miscalculation. Attention allocation failure is a temporal phenomenon; observing it requires time.

The recommended structure is four to six rounds conducted over a minimum of two weeks and preferably months, with each round representing one year of strategic time. Between rounds, participants rationalize their prior commitments, reinforce social consensus within their teams, and seek interpretive

coherence. These are processes that mirror the institutional dynamics through which real organizations resist belief revision under ambiguity. The inter-round period is not administrative; it is part of the diagnostic design. Over time, a new situation becomes normal.

Post-round feedback is deliberately structured to provide consequence information without interpretive correction. Participants learn what happened in the world as a result of their behavior, not whether their epistemic commitments were correct. The adjudicator's role between rounds is to update the world state and communicate consequences, not to guide participants toward better interpretations. Pedagogical steering in the inter-round period destroys the diagnostic value of subsequent rounds.

What Epistemic Wargaming Is Not

The diagnostic value of epistemic wargaming depends on maintaining clarity about its purpose relative to other methodologies with which it shares surface features.

Political-military games surface disagreement and expose coalition friction but typically collapse attention into a single strategic contest. Red teaming challenges assumptions but does not show what happens when those assumptions persist unchallenged across time. Structured analytic techniques evaluate explanations but do not allow those explanations to reshape reality. Crisis games test decision-making under pressure but not show how crisis conditions form through prior attention allocation. Scenario planning generates alternative futures but does not expose which signals participants would have ignored along the path to each future.

Epistemic wargaming is not a training tool for escalation control or optimal crisis response. It should not be used to teach better strategy. It should be used to diagnose how professional cultures allocate attention, which signal types dominate sensemaking under ambiguity, where deterrence logic crowds out structural awareness, and how miscalculation becomes socially reinforced through organizational success on the wrong contest. Its value lies not in producing better strategies, but in revealing why competent actors apply good strategies to the wrong problems.

Implications for Strategic Culture

Epistemic wargaming is particularly effective as a diagnostic instrument for national security institutions because strategic culture exhibits consistent and documentable attention allocation patterns that gray-zone adversaries have learned to exploit.

The bias toward visible action, deterrence signaling, and crisis response is not a failure of individual judgment. It is an institutional output produced by professional training, oversight mechanisms, reporting requirements, and alliance obligations that systematically reward responsiveness over structural awareness. Organizations that are optimized for deterrence and crisis management will consistently overweight salient, visible, reversible contests and underweight slow, opaque, structurally binding ones. This is not incompetence. It is the predictable consequence of institutional design.

Adversaries who understand this pattern can construct strategic environments that exploit it deliberately. They excel at presenting compelling deterrence contests and genuine moral urgencies that absorb institutional attention while structural changes advance through administrative, legal, and

technological mechanisms that remain below the deterrence threshold. The autonomous maritime enforcement scenario is one example of this approach.

Epistemic wargaming provides an instrument for making this dynamic visible to the participants most likely to be subject to it, such as senior military and civilian leaders who are products of the institutional cultures that generate the pattern. The methodology does not argue that deterrence is wrong or that moral urgencies should be ignored. It reveals the opportunity cost of attending to them exclusively, and it does so through experience rather than assertion.

Conclusion

Strategic miscalculation in gray-zone competition is rarely dramatic. It is quiet, professional, and socially validated. It emerges when organizations allocate sustained attention to contests that feel decisive while structurally decisive shifts occur elsewhere. It does not occur because individuals fail, but because institutions are designed to succeed at the wrong problems.

Epistemic wargaming provides a method for making this process visible before consequences become irreversible. By forcing participants to commit explicitly to which contests matter and why, by allowing success on non-decisive arcs, and by accumulating structural consequences neutrally across time, it exposes the attention allocation dynamics that determine strategic position in gray-zone environments. The methodology does not produce optimal strategies. It reveals why capable organizations with good strategies fail to apply them where it matters most.

In gray-zone competition, the contest is not won by those who respond fastest to visible pressure. It is won by those who correctly identify which pressures are structurally decisive and allocate sustained attention accordingly before the window for non-escalatory response closes. Epistemic wargaming is a practical diagnostic instrument for developing that capacity at the institutional level where it is most consequential.

References

- Acemoglu, D., & Wolitzky, A. (2023). Mistrust, misperception, and misunderstanding: Imperfect information and conflict dynamics. MIT Economics. <https://economics.mit.edu/sites/default/files/inline-files/misperceptions%2023%20final.pdf>
- Canyon, D. (2020a). Competitive security gaming: Rethinking wargaming to provide competitive intelligence that informs strategic competition and national security. Daniel K. Inouye Asia-Pacific Center for Security Studies. <https://dkiapcss.edu/wp-content/uploads/2020/11/N2542-Canyon-Rethinking-wargames.pdf>
- Canyon, D. (2020b). Strategic competition, national security and the need for competitive intelligence. Security Nexus. https://dkiapcss.edu/nexus_articles/strategic-competition-national-security-and-the-need-for-competitive-intelligence/
- Canyon, D., & Chan, J. (2020). Gaming major power rivalry and climate disasters using systems tools. ResearchGate. <https://www.researchgate.net/publication/344381125>

- Canyon, D., & Ryan, B. (2021). Military and private sector HADR: Now a sophisticated tool for strategic competition. Daniel K. Inouye Asia-Pacific Center for Security Studies. https://dkiapcss.edu/wp-content/uploads/2021/09/N2623-Canyon_and_Ryan-HADR_strategic_competition.pdf
- Jervis, R. (1976). Perception and misperception in international politics. Princeton University Press.
- Levy, J. S. (1983). Misperception and the causes of war. *World Politics*, 36(1), 76–99. <https://fas-polisci.rutgers.edu/levy/articles/1983%20Misperception%20and%20War.pdf>
- Mitroff, I. I., & Featheringham, T. R. (1974). On systemic problem solving and the error of the third kind. *Behavioral Science*, 19(6), 383–393.
- Rubel, R. C. (2006). The epistemology of war gaming. *Naval War College Review*, 59(2). <https://digital-commons.usnwc.edu/nwc-review/vol59/iss2/8/>



The views expressed in this article are those of the author and do not reflect the official policy or position of DKI APCSS, the Department of War, or the U.S. Government. The appearance of external hyperlinks does not constitute endorsement by the United States Department of Defense (DoD) of the linked websites, or the information, products, or services contained therein. DoD does not exercise any editorial, security, or other control over the information you may find at these sites.

February 2026