

Nobody Left to Deter

Command Fragmentation in Iran, the Sequencing Problem,
and the Nuclear Risk You Are Not Pricing

Dr. Gregory S. Carmichael, DBA

Lt. Col., USAF (Ret.) • Chief Executive Officer, Quantum Reserve Capital
San Juan, Puerto Rico

Paper 02 • Geopolitical Risk Series • 23 July 2026

Abstract

The public debate over Iranian nuclear weapons asks the wrong question. It asks *whether* Iran acquires a weapon. The variable that actually governs risk is *under what institutional conditions* it does so. This paper argues that the sequencing of nuclear acquisition relative to command consolidation matters more than acquisition itself, and that current conditions are producing the dangerous ordering.

A nuclear state that consolidates authority first — clear succession, subordinated commands, an identifiable decision-maker, a survivable arsenal — converges toward the cautious behavior observed in every historical case of nuclear acquisition. A nuclear state that acquires weapons *during* institutional collapse, with fragmented authority and a leadership cadre selected by attrition, does not. These are not adjacent risks. They are different distributions.

The paper states the deterrence case at full strength, including a correction of three widely circulated claims about Twelver Shia eschatology that do not survive contact with the sources, and then argues that the case does not survive 2026. Its entire evidentiary base is drawn from a decision-making population that has been substantially eliminated. Iranian senior command has been struck at least three times between June 2025 and April 2026; mean tenure in the senior IRGC commands examined here has collapsed by roughly an order of magnitude. The Islamic Revolutionary Guard Corps maintains a distributed ground architecture explicitly designed to act independently after decapitation.

The synthesis is that ideology's operative role is structural, not motivational: under acute threat, deliberative capacity degrades and decision-makers revert to over-learned defaults, while attrition selects survivors for hardness rather than caution. The resulting risk requires no one to decide to start a nuclear war.

A partial test is available in the one domain where this institution's coherence can be observed rather than inferred. Public U.S. reporting on Iranian cyber operations against critical infrastructure returns a split verdict: operational infrastructure durability has collapsed on a curve resembling command tenure, the campaign has not paused through three decapitation waves, and implants demonstrably outlive the commands that placed them — while the unit responsible is simultaneously improving its tradecraft. Four indicators are proposed for instrumentation, one of them publicly observable in near-real time, together with explicit falsification conditions. The paper closes on the market implication: the energy risk premium prices supply interruption, not command failure, and the outcome distribution is bimodal in a way that renders its own expected value nearly uninformative.

Keywords: escalation, deterrence, command and control, nuclear proliferation, Iran, threat rigidity, tail risk, energy markets.

JEL Classification: D74, D81, F51, H56, Q34.

Disclosed conflicts: The author is Chief Executive Officer of Quantum Reserve Capital, which operates a paper-traded macro research portfolio with modeled exposure to energy and precious-metals volatility. No live capital is deployed against the views expressed here. The author is a retired United States Air Force officer and a former Department of Defense advisor. This paper reflects the author's independent analysis and is not written for, on behalf of, or at the direction of any government.

Handling: Open-source analysis. No classified, proprietary, or non-public information is used, referenced, or implied. All factual claims derive from public reporting and published scholarship as cited.

Standing caveat: Probability estimates here are structured judgment, not computation — offered to make reasoning falsifiable and to force explicit comparison between branches, not to be read as model output. The confidence intervals around them are wide.

1 The Wrong Question

The question dominating public discussion of Iran is whether it acquires a nuclear weapon. That question is not unimportant, but it is the wrong organizing variable, and treating it as the primary one has produced a debate in which both camps talk past the actual risk.

Consider two states. Both possess a small nuclear arsenal. Both are governed by a revolutionary regime with maximalist rhetoric. Both face a nuclear-armed adversary with a declared preventive-strike doctrine at close range.

In the first, authority is consolidated. There is an identifiable head of state whose control over the arsenal is unambiguous, a functioning chain of command, a general staff with decades of accumulated institutional memory, and a weapons complex sufficiently dispersed and hardened that the state does not fear losing it in a first strike.

In the second, the head of state was killed five months ago. His successor was installed in days under pressure from the security services, and the constitutionality of the process is contested. The senior military command has been struck three times in ten months. The officers now holding operational authority have been in post for weeks. The ground forces operate under a doctrine written explicitly to permit independent action when higher command is unreachable. The arsenal, such as it is, is small, poorly dispersed, and understood by everyone including its owners to be vulnerable.

These two states have the same capability and radically different risk profiles. The first is the configuration in which nuclear acquisition has historically produced caution. The second is the configuration in which it produces catastrophe. The variable separating them is not capability, ideology, or intent. It is **institutional coherence at the moment of acquisition**.

At the Dinner Table

We are talking about enrichment percentages. Which is a fact about centrifuges. Which is a fact about metallurgy. None of it is a fact about the people who have to decide, at three in the morning, whether to destroy everything.

This paper makes three claims.

First, the standard case for Iranian deterrability is stronger than its critics allow. It rests on a long behavioral record, a hard historical test the regime passed at enormous ideological cost, and a reading of Shia doctrine that is more accurate than the one circulating in popular commentary. Section 3 states that case at full strength and corrects three specific claims about Twelver eschatology that do not survive contact with the sources.

Second, that case does not survive 2026. Its entire evidentiary base is drawn from a decision-making population that has been substantially eliminated. This is a sample-selection problem, and it is fatal to the inference rather than merely inconvenient to it. Section 4 develops the argument.

Third, the resulting risk operates through a mechanism that requires no one to decide to start a nuclear war. Under acute threat, deliberative capacity degrades and decision-makers revert to over-learned defaults; simultaneously, attrition selects survivors for revolutionary commitment rather than strategic caution. Sections 5 through 7 develop this into a falsifiable thesis with named indicators.

Section 8 states the estimates and the conditions that would prove them wrong. Section 9 argues that the current energy risk premium prices supply interruption but not command failure, and that the two have different distributions.

2 Where the Escalation Trap Leads

The strategic frame for what follows is Robert Pape's escalation trap, and it is worth stating precisely because it is frequently reduced to the truism that wars escalate.

Pape's argument is narrower and sharper. An escalation trap forms when early tactical success fails to produce the intended political result. The rational response would be to reconsider the theory of victory. The observed response is almost always to interpret failure as insufficient force. Threats intended to deter become tests of will, and each round makes withdrawal more politically costly than continuation. The mechanism is not irrationality. It is that reputational and sunk-cost commitments accumulate faster than strategic learning.

Pape has publicly staged the Iran case as follows. Stage 1 is leadership-decapitation bombing, in which targets are struck and leaders killed but the regime consolidates rather than fractures. Stage 2 is horizontal escalation by that consolidated regime, principally the closure of the Strait of Hormuz. Stage 3 is the ground option undertaken to reopen it. Stage 4 is the emergence of Iran as what Pape terms a fourth center of world power, aligned with Russia and China against American regional primacy.

Two features of this framework matter for the argument here.

The first is that it has substantial predictive standing. Stages 1 through 3 were publicly specified before they occurred and then occurred in sequence. This is unusual and it earns the framework a hearing on stage 4.

The second is the fork Pape identifies at the end. The off-ramp is real but its content is unpleasant: the United States can exit the trap by accepting Iranian regional ascendancy, which on Pape's own account means accepting a nuclear Iran within a year or two. De-escalation and strategic defeat are, in this framing, the same event described from different angles.

The escalation trap does not have a good exit. It has a costly exit and a catastrophic one. This paper is about a risk that lives inside the *costly* exit — the one commonly described as de-escalation — and which is therefore systematically underexamined, because analytic attention concentrates on the branch everyone agrees is bad.

The conventional reading of the off-ramp is that it defers conflict rather than resolving it: Iran ascends, becomes nuclear, and a larger confrontation arrives later. That reading is half right, and Section 5 argues that the half it gets wrong is the important half.

3 The Deterrence Case, Stated at Full Strength

Serious analysis requires stating the opposing case in its strongest form rather than its most convenient one. The case that Iran is deterrable is considerably stronger than popular commentary allows, and three of the claims most often used against it are simply inaccurate.

3.1 The behavioral record

The Islamic Republic has existed since 1979 and has been in adversarial relationship with the United States and Israel for that entire period. Across forty-seven years it has conducted its conflict with Israel almost exclusively through proxies, has repeatedly absorbed direct attacks without proportionate response, and has consistently subordinated ideological objectives to

regime survival when the two conflicted.

The clearest instances are structural rather than rhetorical. Iran abandoned its position in Syria rather than expend the resources to hold it. It absorbed the killing of Qasem Soleimani in 2020 with a telegraphed, casualty-minimizing missile response. Its direct exchanges with Israel in 2024 were announced in advance through intermediaries, calibrated in size, and designed to be intercepted — the behavior of a state performing retaliation while managing escalation, not a state seeking a decisive outcome.

3.2 The 1988 test

The strongest single data point is the end of the Iran–Iraq war.

In July 1988 Ruhollah Khomeini accepted United Nations Security Council Resolution 598 and terminated a war he had personally framed in explicitly religious terms, had repeatedly sworn to continue until the liberation of Karbala, and in which martyrdom had been operational doctrine at the scale of human-wave assaults. He described the decision as more bitter than drinking poison. He made it anyway, because the alternative was the destruction of the state.

The significance is not that he chose peace. It is the identity of the chooser. This was the most ideologically committed leader the Islamic Republic will ever have, at the apex of revolutionary fervor, with unquestioned personal authority, choosing institutional survival over religious continuation and saying publicly how much it cost him. Any theory holding that the regime is indifferent to its own destruction must explain 1988, and none does.

3.3 What the doctrine actually says

Three claims about Twelver Shia eschatology circulate widely in strategic commentary. All three are inaccurate, and the inaccuracies run in the direction that makes catastrophe look more certain.

Claim one: the return of the Hidden Imam is contingent on the destruction of Israel. The classical Twelver sources on the signs preceding the return — the rising of the Sufyani, the cry from the heavens, the killing of the Pure Soul at Mecca, the swallowing of an army at Bayda — were compiled between the ninth and eleventh centuries. They predate the modern state of Israel by roughly a millennium and do not reference it. Modern Iranian popular preaching has retrofitted Israel into the narrative and that retrofit circulates widely, but it is homiletics rather than established doctrine, and it is contested within the seminaries.

Claim two: the return can be deliberately hastened by human action. This is the load-bearing error. In mainstream Twelver theology the timing of the return belongs to God alone. There is an old and explicit quietist tradition warning against attempts to force or predict it, and claims to know the timing are treated as a characteristic error rather than a legitimate program. The concept of *ta'jil*, hastening, exists in Shia practice as prayer and moral preparation. It is not a military doctrine.

The institutional history runs the same direction. The Hojjatieh Society, the notable apocalyptic current in modern Iranian religious life, was pressured into dissolution by Khomeini in 1983 — the revolutionary state suppressed the millenarians. Mahmoud Ahmadinejad's Mahdhist enthusiasms in the 2000s drew accusations of doctrinal deviation from within the establishment clergy. The Islamic Republic has consistently treated apocalyptic activism as a threat to itself, which is what one would expect of an institution whose own legal foundation

depends on the occultation continuing.

Claim three: mass casualties are religiously costless because the dead become martyrs. Martyrdom is genuinely valorized and this paper does not minimize it. But it does not do the work the argument requires. Martyrdom is conferred on those who die in a legitimate cause; it is not a jurisprudential indemnity for the party doing the killing, and the mass killing of Muslims is not martyrdom-conferring for the killer. A nuclear weapon used against Israel kills the Palestinian population Iran claims to champion and contaminates the third holiest site in Islam. Iranian clerical writing has never resolved this cleanly, because it does not resolve.

At the Dinner Table

The theology does not say what it is often claimed to say. The return of the Twelfth Imam is not a thing anyone believes they can schedule, and the version in which a regime bombs its way to the end of history is closer to a Western thriller plot than to anything in the source material.

3.4 Capability and geography

Two further constraints operate independently of intent.

Israel maintains a survivable second-strike capability. Iran, for a considerable period after any acquisition and possibly permanently, would not. A first strike that fails to eliminate retaliation is not a strategy; it is national suicide with additional steps. The asymmetry is not marginal and it is understood by both parties.

Geography compounds this. The territory in dispute is approximately twenty-two thousand square kilometers, and the population Iran claims to be liberating occupies and adjoins it. There is no targeting solution that destroys Israel and spares the Palestinians. This is not a moral observation; it is a physical constraint on the stated objective.

3.5 The empirical failure of the apocalyptic thesis

Finally, the apocalyptic-regime thesis has generated testable predictions and they have failed. The most prominent public instance was Bernard Lewis's 2006 argument in the *Wall Street Journal* that Iran might attempt an apocalyptic act on 22 August of that year, a date of eschatological significance. Nothing occurred. A thesis that generates dated predictions and misses them is not thereby refuted, but it accumulates a record, and this one's record is poor.

The case for Iranian deterrability rests on forty-seven years of consistent behavior, a hard test passed at extraordinary ideological cost, a correct reading of the doctrine, and physical constraints that operate regardless of intent. Anyone modeling Iran as an apocalyptic actor seeking the end of history is modeling something that has not been observed. That case should be taken seriously — and it is about to be substantially undermined, though not for any of the reasons its critics usually give.

4 Why That Case Does Not Survive 2026

4.1 The sample-selection problem

Every data point in Section 3's behavioral record — the 1988 ceasefire, the calibrated 2024 exchanges, the Soleimani response, the abandonment of Syria, forty-seven years of proxy-only warfare against Israel — was generated by a decision-making population that has been substantially eliminated.

This is not a caveat. It is a structural defect in the inference.

Behavioral evidence about an organization predicts future behavior only while the organization persists. What is being measured is not a property of the Iranian state in the abstract but a property of a particular set of people occupying particular roles with particular accumulated judgment. Remove those people and the base rate measures a population that no longer exists.

Assumptions Named

This section assumes that (i) institutional behavior is carried substantially by individuals rather than by formal structures alone, (ii) the reported attrition figures in Section 6 are approximately accurate, and (iii) replacement officers have materially less accumulated crisis judgment than those they replace. Assumption (i) is contestable — institutions do encode doctrine that survives personnel turnover — and Section 7.3 addresses the strongest version of that objection.

The 1988 case is more vulnerable than it first appears. Khomeini's decision was made by an intact institution, over a period of months, after eight years of attritional warfare, with unquestioned personal authority and time to deliberate. That is a slow-crisis decision. Using it to price a fast-crisis risk imports an assumption about the stress regime that does not hold. Different distribution.

4.2 Threat rigidity

The mechanism converting institutional damage into escalation risk is well characterized in the organizational literature and it is not simply that people become emotional.

The relevant finding is *threat rigidity*: under acute threat, individuals and organizations restrict information processing and revert to dominant, over-learned responses. Time horizons compress. Working memory available for counterfactual reasoning collapses. Attribution of hostile intent rises. Control is centralized precisely when distributed judgment would be more useful. This pattern was formalized by Staw, Sandelands, and Dutton in 1981 and has been replicated across individual, group, and organizational levels.

The strategic implication is specific. Decision-makers under acute threat do not reason freshly from first principles. They execute whatever script is most deeply installed.

The literature is not unidirectional, and the paper's argument is stronger for saying so. Acute threat can also induce hyper-caution, decision paralysis, and extreme centralization of control. A cadre that has watched its immediate superiors killed has powerful personal reasons to avoid the action that invited the strike. A reader could reasonably object that decapitation produces timidity rather than trigger-pulling.

The objection is correct and does not rescue the situation, for a reason specific to Configuration B: **caution at the center is not stabilizing when authority sits at the periphery**. A leadership paralysed by fear of its own destruction is, functionally, a leadership that cannot issue a stand-down order — and thirty-two provincial commands operating on pre-existing orders do not require one to act. Hyper-caution and fragmentation are not opposing forces

that cancel. They are complementary failure modes: the first removes the brake, the second leaves the accelerator distributed. The stabilizing case requires caution and authority to be located in the same place, which is precisely what has been dismantled.

This reframes the role of ideology in a way that matters. The standard claim is that religious ideology *motivates* — that leaders seek apocalypse. Section 3 gives good reasons to reject that. The stronger claim, and the one this paper adopts, is that ideology functions as a **default operating system**: not what is chosen, but what remains when deliberative capacity degrades. That mechanism does not require anyone to be an apocalyptic actor. It requires only that they be under sufficient load.

4.3 The historical record of leaders under existential attack

Section 3’s optimistic case rests on Iranian behavior under pressure. It is worth asking how leaders in general behave when the threat becomes existential rather than merely severe. The record is not encouraging.

Table 1: Leadership behavior under existential threat: illustrative cases (not a sample)

Case	Observed behavior
Iraq, 1991	Held chemical weapons in reserve and did not employ them despite regime-threatening defeat. Deterrence held.
Japan, August 1945	After two nuclear detonations, a faction of the military attempted a palace coup to prevent surrender. Rational updating did not occur among those holding operational authority.
Germany, March 1945	The Nero Decree ordered the destruction of Germany’s own remaining infrastructure. Implementation was blocked by subordinate refusal, not by reconsideration.
Syria, 2013–2018	Repeatedly escalated to chemical weapons employment as regime survival came into question.
Libya, 2011	Fought to personal destruction rather than accept a negotiated exit that was available.

These cases are selected, not sampled, and the ratio should not be read as a base rate. They were chosen because they are well documented and memorable, and memorability correlates with escalation — regimes that faced existential threat and quietly capitulated, negotiated, or dissolved are systematically underrepresented in the historical record because nobody writes about them. The table is illustrative of a mechanism, not evidence of its frequency. Noting this is required by consistency: Section 4.1 rejects the optimistic case for drawing inferences from an unrepresentative population, and the same standard applies here.

With that qualification, one case supports deterrence and four support escalation. The Japanese instance is the most damaging to the optimistic reading precisely because it is the strongest possible test: even after two nuclear weapons had been used against the homeland, a faction with operational authority attempted to seize the palace and prevent surrender. Breaking the deadlock required an intervention from outside the normal decision structure.

At the Dinner Table

The comforting story is that when things get bad enough, people come to their senses. The historical record says something less comfortable: when things get bad enough, some people come to their senses and others reach for the thing they have been preparing to reach for their entire careers. Which type is holding the authority at that moment is mostly luck.

4.4 What restraint under weakness does and does not predict

There is a further limitation on Section 3's evidence, and it is the strongest objection to this paper's own optimistic passages.

Every year of demonstrated Iranian caution occurred while Iran was militarily and economically weak. Restraint under vulnerability is weak evidence about behavior under nuclear cover. A state that has been strategically defensive throughout its existence because it could be destroyed may behave differently once it believes it cannot be.

This paper does not resolve that objection. It weights the behavioral record heavily because that record includes cases where the regime chose survival at extraordinary ideological cost — but the extrapolation from weakness to strength is genuinely uncertain, and any analysis claiming otherwise is overselling its evidence.

The optimistic case is built on a base rate derived from a population that has been substantially removed, applied to a stress regime it was never measured under, and extrapolated across a capability transition it has never been tested through. Each of those three problems is independently serious. Together they mean the reassuring answer cannot bear the weight commonly placed on it.

5 The Sequencing Argument

This section states the paper's central claim.

5.1 Two configurations

Nuclear acquisition is not a single event with a single risk profile. It occurs against an institutional background, and that background does more work than the acquisition itself.

Configuration A — consolidate, then acquire. The state possesses clear and uncontested succession, subordinated military commands reporting to an identifiable civilian or clerical authority, an intact senior officer corps with accumulated crisis experience, and an arsenal sufficiently dispersed and hardened that it does not fear losing it in a first strike.

Configuration B — acquire during collapse. Succession is contested or was resolved under duress. Senior command has been repeatedly decapitated. Operational authority sits with officers of short tenure. Doctrine authorizes independent action when higher command is unreachable. The arsenal is small, poorly dispersed, and known by its owners to be vulnerable.

The historical cases that generate the reassuring pattern — China after 1964, India and Pakistan after 1998, and arguably North Korea — are all Configuration A. Each acquired weapons from a position of institutional stability, and each subsequently exhibited the pattern

of conventional caution combined with increased sub-conventional aggression that Glenn Snyder named the stability–instability paradox. No state has acquired nuclear weapons in Configuration B, which is why the empirical literature offers so little guidance here.

5.2 The four failure modes

The pathways to nuclear use that should concern an analyst are not the deliberate ones. Four mechanisms account for most of the risk, and each is amplified by Configuration B.

Use-it-or-lose-it. A small, poorly dispersed arsenal facing an adversary with a declared preventive-strike doctrine creates a standing incentive to launch during any severe crisis rather than risk losing the force on the ground. The instability lies not in intent but in *vulnerability*. Counterintuitively, a more survivable Iranian arsenal would be stabilizing.

Pre-delegation. A regime that fears decapitation may distribute launch authority downward to guarantee retaliation. This solves the credibility problem and creates a far worse one: multiple parties with release authority, operating on degraded communications, under wartime conditions, without the political context available at the center.

Warning-time collapse. The Israel–Iran dyad has flight times measured in single-digit minutes. Every serious nuclear near-miss on record — the 1979 NORAD training-tape incident, Stanislav Petrov’s 1983 judgment call, Able Archer, the 1995 Norwegian sounding-rocket alert — was survived because someone had minutes to hours in which to think. That margin does not exist here.

Cascade triangulation. Saudi Arabia has stated it will match an Iranian capability. Every stable nuclear relationship in history has been bilateral. Nobody has solved three-body deterrence, and the Middle Eastern version would run multiple simultaneous dyads with no verification architecture, no hotlines, and no shared crisis vocabulary.

None of the four failure modes requires anyone to decide to destroy anyone. They are properties of a system, not choices by an actor. This is why the debate over Iranian intentions — however rigorously conducted — addresses the smaller part of the problem.

5.3 Why fragmentation amplifies each

The four mechanisms are not independent of institutional coherence. Each is worsened by it in a specific way.

The claim is therefore not that Configuration B is somewhat worse. It is that Configuration B degrades every one of the four principal safety mechanisms simultaneously, and that the mechanisms are correlated rather than independent — a crisis that triggers one is likely to trigger others.

At the Dinner Table

There is a version of a nuclear Iran that the world could probably live with, unhappily, for decades — the way it has lived with a nuclear Pakistan. Getting there requires Iran to be stable first and armed second. What is being manufactured right now is the reverse

Table 2: Failure modes under each configuration

Mechanism	Configuration A	Configuration B
Use-it-or-lose-it	Mitigated by dispersal and hardening achieved before crisis.	Maximal: arsenal built under fire, no time to harden.
Pre-delegation	Unnecessary; center is survivable and communications hold.	Near-mandatory; center has already been struck repeatedly.
Warning-time collapse	Partially offset by trained staff and rehearsed procedure.	Compounded: short-tenure officers, no rehearsal, no procedure.
Cascade triangulation	Slow: neighbors respond to a stable fact.	Fast: neighbors respond to a chaotic one and hedge earlier.

order.

6 The Evidence: Iranian Command Attrition, 2025–2026

The preceding sections are analytical. This one is empirical, and it is the reason the argument is not merely theoretical.

6.1 The attrition record

Iranian senior leadership has been struck in at least three distinct waves between June 2025 and April 2026.

Table 3: Selected senior Iranian leadership losses, June 2025 – April 2026

Date	Individual	Position
13 Jun 2025	Hossein Salami	Commander-in-Chief, IRGC
13 Jun 2025	Amir Ali Hajizadeh	Commander, IRGC Aerospace Force
28 Feb 2026	Ali Khamenei	Supreme Leader
28 Feb 2026	Mohammad Pakpour	Commander-in-Chief, IRGC
17 Mar 2026	Ali Larijani	Secretary, Supreme National Security Council
6 Apr 2026	Majid Khademi	Director, IRGC Intelligence Organization

Reporting places the total at more than forty senior officials killed since the February 2026 strikes began. The Larijani case carries additional weight: he was publicly identified as the primary interlocutor for a ten-point negotiating framework, and his removal was described as resetting the clock on a potential settlement. The decapitation campaign has therefore eliminated not only command capacity but the specific individuals through whom de-escalation would have been transacted.

6.2 Quantifying the memory loss

The strategic significance is best captured by in-post tenure, which proxies for accumulated crisis judgment in a specific role.

As of 1 June 2025, the three senior IRGC commands examined here were held by officers with the following in-post tenures: Commander-in-Chief, approximately 74 months; Aerospace Force, approximately 188 months; Ground Forces, approximately 194 months. Accumulated in-post experience across the three positions:

$$74 + 188 + 194 = 456 \text{ months} \approx 38 \text{ years.}$$

As of 23 July 2026, the same three commands are held by officers with in-post tenures of approximately 4.8, 13.3, and 13.1 months respectively:

$$4.8 + 13.3 + 13.1 = 31.2 \text{ months} \approx 2.6 \text{ years.}$$

The reduction is

$$1 - \frac{31.2}{456} = 0.932,$$

a decline of roughly 93 percent, or a factor of approximately 14.6, in under fourteen months.

That figure alone would be open to a fair objection: the incoming officers are not novices, so why should a short tenure imply degraded judgment? The answer is visible only when the second metric is placed beside the first.

Total IRGC service is approximately unchanged. The three predecessors had served roughly 45, 44, and 46 years respectively; the three current incumbents have served roughly 46, 43, and 40 years. Accumulated total service across the three commands therefore moved from about 135 years to about 129 years:

$$1 - \frac{129}{135} = 0.044,$$

a decline of roughly 4 percent.

Measure (three senior IRGC com- mands)	Jun 2025	Jul 2026
Accumulated total IRGC service	≈ 135 yr	≈ 129 yr (−4%)
Accumulated in-post tenure	≈ 38 yr	≈ 2.6 yr (−93%)

The divergence is the finding. Military experience is essentially intact; role-specific tenure has been very nearly erased. The two measures moving in opposite directions isolates what the decapitation campaign actually destroyed, and it is not competence. It is the accumulated understanding of what *this particular command* does in a crisis — built through repetition in the seat, and not transferable from a different seat however distinguished the service in it.

Assumptions Named

Service-year figures are approximate to within roughly two years, reconstructed from publicly reported entry dates and career histories; in-post tenures are accurate to within about one month. The claim defended here is narrower than "Iran's military is now run

by novices" — that claim is false, and the total-service figures above establish that it is false. The claim is that role-specific crisis judgment has been destroyed while general military competence has not. Readers who reject the premise that role-specific judgment matters independently of total service should reject this section; the rest of the argument does not depend on the magnitude of the effect, only on its direction.

The Commander-in-Chief position illustrates the acceleration directly. In the 74 months to June 2025 it changed hands once. In the 13 months since, it has changed hands twice.




Figure 1: Accumulated in-post tenure across three senior IRGC commands (Commander-in-Chief, Aerospace Force, Ground Forces). The decline is approximately 93 percent in fourteen months.

6.3 The distributed-command architecture

The second element of the evidence is doctrinal rather than personnel-based, and it is the more consequential of the two.

The IRGC Ground Forces established thirty-two independent provincial commands during the 2000s, structured explicitly to continue operating after a decapitation strike. Each maintains its own chain of command and pre-existing orders. This was a deliberate design response to precisely the campaign now being conducted against it, and by the standard of institutional resilience it appears to be working.

By the standard of nuclear stability, it is close to a worst case. A force architecture optimized to act without reference to a center is, by construction, an architecture in which the center cannot reliably restrain action. Resilience against decapitation and controllability under crisis are in direct tension, and Iran resolved that tension in favor of resilience two decades ago, for reasons that were sound at the time.

6.4 The succession problem

The third element is authority itself. Reporting indicates that the Assembly of Experts selected Mojtaba Khamenei as Supreme Leader within days of his father's death, under pressure from the IRGC, with the decision withheld from public announcement and with contested adherence to constitutional procedure.

Whatever the merits of the selection, the resulting authority is of a different character

than the one it replaced. Ali Khamenei had thirty-six years of accumulated standing. His successor was installed at speed by the security services during active bombardment. The most consequential public assessment of the situation was that there is no clear answer to who decides when the missiles stop.

Deterrence requires an address. It requires an identifiable party who can be threatened, who can be reassured, who can receive a signal and act on it, and who can order a stand-down that will be obeyed. What has been produced over fourteen months is a system in which that address is genuinely unclear — not as rhetoric, but as an operational fact acknowledged by observers across the analytical spectrum.

7 The Selection Filter

Sections 4 and 6 establish that the decision-making population has changed and that authority has fragmented. This section specifies the mechanism by which those facts translate into elevated risk.

7.1 An antecedent: what organizational network analysis already measured

The argument this paper makes qualitatively — that decapitation destroys role-specific judgment while leaving general competence intact — has a formal antecedent in a research program that has been running for four decades, and naming it sharpens the claim.

Beginning in the late 1980s, the corporate anthropologist Karen Stephenson developed what became organizational network analysis (ONA): the mapping of an organization's informal structure of trust and consultation, as distinct from its formal chart. Her central empirical finding is that if one maps who actually consults and trusts whom, three mathematically defined roles recur across every organization. *Hubs* are high-centrality connectors, the shortest route to needed information. *Gatekeepers* are high-betweenness nodes: they sit on the only path linking otherwise-unconnected parts of the network, so their removal does not subtract a person but severs a connection nothing else spans. *Pulsetakers* are the nodes positioned to monitor the health of the whole. The framework was developed in part on defense organizations; Stephenson worked with U.S. government data on the corrosive effects of bureaucratic structure on institutional function.

Stated in these terms, this paper's thesis is a claim about the **targeted removal of high-betweenness nodes**. A decapitation campaign does not sample uniformly (Section 7.1); it removes the locatable and the senior, which is a description of hubs and gatekeepers. General competence survives because the rank-and-file and the technical hubs survive. Coherence collapses because the betweenness nodes — the people who constituted the only path between the center and the periphery — are precisely who is killed. In-post tenure, the metric of Section 6.2, is a field-observable proxy for a quantity ONA measures directly: the betweenness centrality of the individuals removed.

The framework has already crossed from corporate application into the analysis of clandestine and adversary networks. Recent work in terrorism studies uses betweenness centrality to locate “hinge figures” — brokers connecting otherwise-separated cells — in exactly Stephenson's sense. This paper's application to a decapitated state military is a step along an established line, not an improvisation.

One further connection is worth naming, because it reframes a structure introduced earlier. The IRGC's thirty-two provincial commands (Section 6.3), built to keep operating when the center is removed, are in network terms a deliberately *low-betweenness-dependence architecture*: engineered so that no single gatekeeper's removal severs the network. That design is why the structure survives decapitation — and, by the same property, why it cannot be centrally restrained. Resilience against decapitation and controllability under crisis are the same tradeoff viewed from two sides, and the IRGC resolved it in favor of resilience two decades ago. A network built to have no indispensable gatekeeper is a network with no address at which a stand-down order can be received and enforced.

At the Dinner Table

An organization is not its org chart. It is the handful of people through whom everything actually passes — and if you remove them, the chart still looks intact while the thing stops working. Someone spent forty years learning to measure exactly those people. This paper is measuring what happens when an adversary spends a year killing them.

7.2 Decapitation is not random removal

A targeted-killing campaign does not sample uniformly from a leadership population. It removes those who are locatable, and it promotes those who survive. Survivorship under sustained targeting selects for specific traits: operational security discipline, willingness to accept personal risk, suspicion of communications, and — critically — the absence of the visible political profile that comes with being a designated negotiator.

The Larijani case makes this concrete. The individual identified as the primary channel for a negotiated settlement was killed. Whatever the intent, the effect on the surviving population is directional: it removes the people who talk to the other side and retains the people who do not.

Compounding this, the surviving cohort has personal grievance that its predecessors did not. The officers now in post watched their superiors killed and in many cases narrowly avoided the same. Threat rigidity operates more strongly on those with direct personal exposure to the threat.

7.3 Ideology as fallback, not motive

This paper rejected the apocalyptic-motive thesis in Section 3 and does not reintroduce it here. The mechanism is different and does not require the thesis.

Under threat rigidity, deliberative capacity degrades and behavior reverts to over-learned defaults. The question then becomes empirical: what is the default?

The honest answer cuts both ways. Forty-seven years of IRGC institutional practice encode asymmetric warfare, deniability, absorb-and-survive, and preservation of the revolution as the terminal value. That is not eschatology. It is daily rehearsed doctrine, and it is arguably encoded more deeply than millenarianism because it is what the institution actually does rather than what it says on commemorative occasions. This is a genuine argument for stability and it should not be discarded.

But institutional scripts live in people. Three rounds of decapitation degrade the script at the same rate they degrade the cadre. An officer with two decades of internalized escalation-management judgment has been replaced by one with months in post — who retains the

ideological frame, which is taught and public, while lacking the tacit operational judgment, which is learned and personal.

At the Dinner Table

Doctrine is what an organization writes down. Judgment is what its senior people know without having to look it up. Bombs destroy the second much faster than the first, and the second is what keeps you out of trouble at three in the morning.

7.4 The strongest objection, and the limit of this paper's claim

The most serious objection to the entire argument is that institutions are more robust than this account allows. Formal structures, written doctrine, standard operating procedures, and organizational culture persist through personnel turnover — that is much of what institutions are for. Iran's system was explicitly designed by Khomeini to outlast any individual, and the 1989 succession, resolved in days despite the designated successor lacking the constitutionally required clerical credentials, demonstrated that the procedure works under pressure.

This objection has real force and this paper does not dismiss it. The response is a distinction: institutional design confers *continuity of function* but not *continuity of judgment*. The system can continue to issue orders, pay soldiers, and fire missiles. What is not preserved by design is the specific capacity that produced 1988 — a leadership able to recognize that the strategic situation had become unwinnable and to absorb the ideological cost of saying so.

That is not a procedure. It is judgment, exercised by a person with the standing to make it stick, and it is exactly what three rounds of decapitation remove.

The risk is not that Iran's leadership believes dangerous things. It is that there may no longer be a coherent leadership to hold beliefs at all, and that ideology's operative function under these conditions is as a selection filter and a cognitive fallback rather than as a motive. A fanatical regime can at least be deterred as a unit. A fragmented one cannot be deterred as anything.

7.5 A test in another domain

The argument to this point is inferential: command coherence is asserted to have degraded, but the degradation is measured through personnel records rather than observed directly. Iranian internal command arrangements are not published. That is a real weakness, and it is worth asking whether any domain exists in which the same organization's coherence can be observed rather than inferred.

One does. Iranian offensive cyber operations against U.S. critical infrastructure are documented in dated, public U.S. government reporting, attributed to a named element of the same institution, with indicators timestamped to the month. The evidence is imperfect and, as shown below, partly confounded — but it is independent of everything in Section 6, and it does not all point the same way.

The source. On 7 April 2026 the FBI, CISA, NSA, EPA, DOE, Cyber Command's Cyber National Mission Force, and Treasury jointly published advisory AA26-097A, documenting exploitation of internet-exposed programmable logic controllers across U.S. government facilities, water

and wastewater systems, and the energy sector. The advisory attributes the activity to Iranian-affiliated actors and links it to prior campaigns by a group affiliated with the IRGC Cyber Electronic Command. It was substantially updated on 22 July 2026. The authoring agencies assess that this targeting has escalated in connection with hostilities involving Iran, the United States, and Israel.

Infrastructure lifetime collapses in the same window as command tenure. The advisory publishes indicators of compromise with the month in which each address was first and last associated with the actors. Treated as a timeline rather than a blocklist, it yields a measure of how long the organization can sustain a piece of operational infrastructure.

Table 4: Actor infrastructure lifetime, derived from AA26-097A indicator tables

Infrastructure first observed	<i>n</i>	Mean	Median	Max
Before February 2026	9	13.0 mo	15 mo	15 mo
February 2026 onward	12	1.5 mo	2 mo	2 mo

The decline in mean sustained lifetime is

$$1 - \frac{1.5}{13.0} = 0.885,$$

approximately 88 percent, a factor of 8.7. One address block ran continuously from January 2025 through March 2026; nothing established after February 2026 has survived beyond two months.

Set beside Section 6.2 — a 93 percent decline in command tenure, a factor of 14.6, over the same window — two independent measures of the same institution move in the same direction at comparable magnitude, one of them derived entirely from a public government document.

Assumptions Named

This correlation is real and its causal interpretation is ambiguous. February and March 2026 contain both the decapitation events and the run-up to the first public advisory, so infrastructure burned on discovery is indistinguishable here from infrastructure churned by organizational disruption. Faster rotation is also the expected response to intensified hunting, which would indicate competence rather than degradation. The published fields record when the FBI ceased observing an address, not necessarily when the actors ceased using it. Sample sizes are 9 and 12. The observation is offered as a candidate indicator worth tracking forward, not as evidence for the thesis.

Two further features of the timeline are worth noting. Monthly coverage is unbroken across the entire period, including March and April 2026 — the campaign did not pause through any of the three decapitation waves, nor around the killing of the official identified as the primary negotiating channel. And the authoring agencies characterise the expansion beyond a single vendor as opportunistic, which describes victim selection rather than command direction, but is at least consistent with technical opportunity driving target sets rather than strategic timing.

The July update is an intent change, not a scope expansion. The more consequential development is qualitative. The April advisory described disruption through project-file interaction and

manipulation of operator displays, producing operational disruption and financial loss. The July update describes something different in kind: a project file that preserved normal downstream function while adding logic overriding the instruction sets responsible for maintaining safe operating parameters, and modifications that disabled shutdown and alarm logic so that systems could enter unsafe conditions without alerting operators.

Normal function preserved, safety envelope defeated, operator blinded. That combination is not harassment and it is not signalling. It is the tradecraft of pre-positioning for a physical-consequence event, held against a decision not yet taken — and it is documented in municipal water and energy systems.

Pre-delegation, implemented in code. Section 5.2.2 argued that a regime fearing decapitation may distribute release authority downward, solving a credibility problem and creating a worse one. The advisory documents that mechanism already operating in a lower-consequence domain.

Describing an earlier campaign beginning in November 2023, the authoring agencies note that malicious ladder logic deployed to victim devices *continues to be observed*. Two and a half years later, code placed by an earlier campaign remains resident on critical infrastructure — outliving the operation, the operators, and in several documented cases the commanders who authorised it.

That is an action that executes without a live decision-maker, requires no surviving communication path, and cannot be recalled by a centre that no longer exists. It is the abstract failure mode of Section 5.2.2 in concrete, verified form.

What cuts the other way. Honesty requires the counter-evidence, and it is substantial. Tradecraft *improved* between April and July 2026: exfiltration of device project files was added, manipulation of reusable code modules is more sophisticated than the earlier work, and the target set expanded across three manufacturers with protocol-specific knowledge for each. An organisation that is learning requires functioning feedback loops, which require coherent command. Five months after the Supreme Leader was killed, this element of the IRGC is measurably better at its work than it was before.

That is a real problem for any strong reading of this paper's thesis, and it should temper one.

It does not dissolve the argument, for a reason that must be stated as a limit rather than a rebuttal. Cyber units are among the most survivable elements of any military: small, geographically distributed, not co-located with senior leadership, and low-signature. Their coherence is weak evidence about the coherence of strategic weapons release authority, which sits elsewhere under different rules. **Autonomous cyber action has never implied delegated nuclear release in any state holding both, and this paper does not claim that it does here.** The domain is offered as a partially observable proxy with a known inferential gap, not as a substitute for evidence about nuclear command.

The one domain where this institution's coherence can be observed rather than inferred returns a split verdict: infrastructure durability has collapsed on a curve resembling command tenure, the campaign has never paused through three decapitations, standing implants demonstrably outlive the commands that placed them — and the unit is simultaneously getting better at its job. A reader should take from this that the thesis is

testable and partly tested, not that it is confirmed.

8 Estimates, Indicators, and Falsification

8.1 Estimates

A first pass at these figures, made before the command-attribution evidence of Section 6 was assembled, placed nuclear use via any pathway at 10 to 20 percent over a consolidated ten-to-twenty-year horizon, with deliberate Iranian first strike at 2 to 5 percent. Incorporating the sequencing argument and the 2026 attrition record moves those figures upward and compresses the horizon. Both columns are shown rather than only the conclusion, because the direction and magnitude of the movement are themselves informative: it is the attrition evidence, not the theory, that does the work.

Table 5: Estimates. All figures are conditional on Iranian nuclear acquisition occurring under Configuration B. They are not unconditional probabilities of nuclear use.

Outcome	First pass	Stated here
Serious crisis involving nuclear alerting, no use	60–70%	55–65%
Nuclear use via any pathway	10–20%	20–30%
Deliberate first strike to destroy Israel	2–5%	category degrades

The conditioning is load-bearing and must not be dropped. This paper does not assess the probability of Iranian acquisition. That is a separate analytical problem requiring enrichment-stockpile, weaponization-timeline, and inspection-access evidence that is outside its scope. The consequence is that every figure above is a conditional probability, and a reader who carries "20–30 percent" away from this paper as an unconditional risk of nuclear use will have overstated it by a large and unknown factor.

To make the arithmetic explicit, let A denote acquisition under Configuration B within the horizon and U denote nuclear use via any pathway. The paper estimates $P(U | A) \in [0.20, 0.30]$. The unconditional figure is

$$P(U) = P(U | A) \cdot P(A).$$

Supplying an illustrative range of $P(A) \in [0.25, 0.50]$ — a range this paper does not defend and offers only to demonstrate the sensitivity — gives

$$P(U) \in [0.25 \times 0.20, 0.50 \times 0.30] = [0.05, 0.15],$$

or roughly 5 to 15 percent rather than 20 to 30. The conditional and unconditional figures differ by a factor of about three at the midpoint, and readers should substitute their own $P(A)$ rather than adopt the illustrative one.

The third row carries the substantive finding. Under fragmented authority, the distinction between deliberate and inadvertent use ceases to be analytically useful. A provincial commander acting on pre-existing orders, under bombardment, believing the regime is ending, operating from an ideological frame because deliberative capacity has degraded, produces an event that neither category describes. The prior taxonomy assumed a coherent decision-maker

choosing among options. That assumption may not hold, and taxonomies that depend on it should be retired rather than re-weighted.

8.2 Decision-cycle arithmetic

The warning-time argument in Section 5.2.3 can be made quantitative, with assumptions stated.

Tehran to Tel Aviv is approximately 1,580 kilometres. A medium-range ballistic missile flown at that range has a total flight time on the order of eleven minutes. Against this, the canonical Cold War case — Moscow to Washington, roughly 7,800 kilometres by intercontinental trajectory — allowed approximately thirty minutes.

The ratio is

$$\frac{11}{30} \approx 0.37,$$

so the Israel–Iran dyad operates on roughly one-third of the Cold War’s decision budget. But the operative figure is not flight time; it is the residual after detection and assessment. Allowing approximately one minute for detection and three to five for assessment and verification, the deliberation window is on the order of

$$11 - 1 - 4 \approx 6 \text{ minutes},$$

before transmission and execution consume further margin. Effective deliberation is therefore plausibly three to five minutes.

Assumptions Named

Flight-time figures are order-of-magnitude estimates derived from range and typical ballistic trajectories, not from system-specific performance data, none of which is used here. Detection and assessment intervals are illustrative rather than measured. The argument does not depend on the precise figures — it depends on the ratio, which is robust to substantial error in the inputs.

For calibration: Stanislav Petrov’s 1983 judgment, the single most consequential individual decision in nuclear history, was made with more time than this dyad structurally allows.

8.3 Four indicators

The thesis implies specific, partially observable indicators. These are proposed for instrumentation in preference to the rhetorical signals that dominate open-source monitoring.

Indicator 1 — Cadre turnover velocity. Mean in-post tenure across senior military and security commands, tracked as a rolling figure. Rising tenure indicates consolidation and is stabilizing. Continued decline indicates that institutional judgment is being destroyed faster than it can be rebuilt. Currently at roughly 10 months against a pre-2025 baseline near 150.

Indicator 2 — Command re-subordination. Evidence that provincial and service commands are being brought back under a single identifiable authority, versus evidence that autonomy is being formalized. Any indication of pre-delegation of release authority is the loudest alarm in the system and should be treated as a discontinuity rather than a data point.

Indicator 3 — Decision-cycle compression. Whether any communication channel — direct, brokered, or tacit — exists between the parties, and whether warning time is being extended

or reduced by force posture. The existence of a channel at all is the single most valuable stabilizing asset available, and it is currently absent.

Indicator 4 — Infrastructure churn and campaign modulation. Sustained lifetime of actor-controlled operational infrastructure, and whether campaign intensity tracks diplomatic events or proceeds independently of them. Falling durability alongside unmodulated escalation is consistent with execution of standing intent; recovering durability alongside intensity that responds to negotiation windows indicates a functioning centre. Section 7.4 establishes a baseline of roughly 13 months' mean infrastructure lifetime before February 2026 against roughly 1.5 months after, with the confounds stated there.

This fourth indicator carries a practical advantage the first three lack. Indicators 1 through 3 require visibility into Iranian internal arrangements that open sources do not reliably provide. Indicator 4 is published: CISA and its partner agencies issue dated advisories with time-stamped indicators, and revisions arrive on a cadence of months. It is the one measure here that a research desk can track in near-real time without privileged access, and for that reason it is the one most worth instrumenting first.

Note the counterintuitive direction of Indicator 2's sub-component: *rising Iranian arsenal survivability is stabilizing*, because it reduces use-it-or-lose-it pressure. Analysts reflexively code hardening and dispersal as threatening. Under this framework they are the opposite.

8.4 What would falsify this thesis

The argument is constructed to be wrong in identifiable ways. It should be abandoned or substantially revised if:

1. A single, publicly legible authority over strategic forces emerges and demonstrates control — most convincingly through a stand-down order visibly obeyed across dispersed commands during an active crisis.
2. Mean in-post tenure across senior commands recovers above roughly 24 months and holds for four consecutive quarters, indicating that consolidation is outpacing attrition.
3. Provincial and service commands are demonstrably re-subordinated, with autonomy reduced rather than formalized.
4. A functioning deconfliction channel is established through any intermediary both parties accept, extending effective decision time.
5. Iran achieves demonstrable arsenal survivability prior to any crisis, converting Configuration B toward Configuration A.

Conversely, the thesis is strengthened by evidence of pre-delegation, by further decapitation events, and by the formalization of independent provincial authority.

9 Market Implications: The Unpriced Layer

9.1 What the current premium prices

As of 23 July 2026 Brent trades near \$98.50, having risen approximately five percent on confirmed attacks against two Saudi tankers enforcing the Houthi blockade declared on 20 July. The World Bank's 2026 baseline of approximately \$86 now reads as conservative.

The structural situation is unusual. The Strait of Hormuz has been contested or closed since late February. Saudi Arabia responded by rerouting crude through the Bab el-Mandeb corridor, which has now itself become effectively impassable for Saudi shipments. Full closure

of Bab el-Mandeb would remove approximately seven percent of global supply, principally by stranding Saudi exports. Both the primary and the secondary exit routes for Persian Gulf crude are simultaneously under threat.

The critical structural point is **stranded spare capacity**. In a conventional supply scare, OPEC+ spare capacity absorbs the shock. That capacity exists. It sits physically behind the chokepoints that are closed. The buffer cannot reach the market, which is why strategic reserves are being asked to perform a function that spare capacity normally performs — and reserves are a finite drawdown, not a renewable one.

9.2 Scenario distribution and why its mean is uninformative

Over the five months to December 2026, with assumptions as stated:

Table 6: Brent scenario distribution, August–December 2026

Scenario	Prob.	Range (midpoint)
Partial stabilization: transit-toll regime, Chinese pressure, escorted Bab el-Mandeb	35%	\$88–95 (91.5)
Sustained attrition: both straits contested, unresolved	35%	\$100–115 (107.5)
Systematic closure or mass-casualty escalation	22%	\$130–160 (145)
Sharp de-escalation: ceasefire, Hormuz re-opens	8%	\$75–85 (80)

The probability-weighted expectation is

$$E[P] = 0.35(91.5) + 0.35(107.5) + 0.22(145) + 0.08(80) = \$107.95.$$

The standard deviation of the same distribution is

$$\sigma = \sqrt{\sum_i p_i (x_i - \bar{x})^2} = \sqrt{459.3} \approx \$21,$$

giving a coefficient of variation of approximately 20 percent, with roughly 22 percent of the mass at or above \$130.

9.3 Sensitivity: does the dispersion claim survive the inputs?

The scenario probabilities above are judgment, and reporting a standard deviation computed from them risks lending a qualitative estimate the appearance of quantitative precision. That objection is fair on presentation. It is answerable on substance, because the claim being made is structural — that dispersion is large relative to the mean — and structural claims can be tested against input error.

Across every perturbation tested — including the uniform case, which is what an analyst claiming no knowledge whatever would assign — the coefficient of variation remains between 16.6 and 23.2 percent, and the expected value moves within a \$16 band. The specific figures are judgment. The conclusion drawn from them is not sensitive to that judgment: under any plausible weighting, and under the assumption of no knowledge at all, dispersion remains roughly a fifth of the mean.

Table 7: Dispersion under alternative scenario weightings (midpoints held constant)

Weighting	$E[P]$	σ	CV
Base case (as stated above)	\$107.95	\$21.43	19.9%
+10pp to stabilization	\$102.60	\$17.90	17.4%
+10pp to escalation	\$113.30	\$23.26	20.5%
+10pp to attrition	\$106.88	\$19.39	18.1%
Escalation weight halved	\$103.14	\$17.13	16.6%
Escalation weight doubled	\$118.74	\$24.27	20.4%
Uniform — maximum ignorance	\$106.00	\$24.54	23.2%

Assumptions Named

This sensitivity test perturbs the probabilities while holding the four scenario midpoints constant. It therefore tests robustness to error in the weighting, not to error in scenario definition or in the price ranges assigned to each. A reader who believes the escalation branch would clear \$180 rather than \$145 is changing an input this test does not vary, and the dispersion figure would rise rather than fall.

The expected value is \$108 and it is close to useless. The distribution is bimodal: either brokered stabilization pulls the price toward \$90, or an escalation event carries it past \$130. Very little probability mass occupies the middle that the mean describes. Positioning to the expectation is positioning to an outcome the model says is unlikely to occur.

9.4 The layer that is not priced at all

Everything above prices *supply interruption*. It is the market's native competence: barrels, routes, days of cover, rerouting cost, insurance rates. The premium is real and roughly appropriate to that risk.

Command failure is a different variable and it is not in the price.

The reason is structural rather than negligent. Value-at-risk and conditional-VaR frameworks are estimated on historical return distributions. A command-and-control failure in a newly nuclear state has no historical instances, and therefore contributes zero to any variance-covariance estimate however long the lookback. It is not underweighted. It is structurally absent, in the same way that a weaponized biological release is absent from the same models.

This is the same analytical failure in a different domain: risks that are non-recurrent and institutionally novel cannot be captured by estimators built on recurrence. The correct response is not to attempt to price them precisely, which is not possible, but to hold convexity against them and to instrument the indicators that would signal a regime change in the underlying distribution.

At the Dinner Table

The market is very good at asking how many barrels are missing. It has no way of asking whether the people with the missiles still answer the phone. The second question is the one that determines whether the first question matters.

9.5 A second transmission channel that bypasses the oil price entirely

Everything above routes Iranian escalation to Western portfolios through one variable: the price of crude. That is the channel the market watches and the channel this paper has modelled. It is not the only one.

Advisory AA26-097A documents Iranian-affiliated actors producing operational disruption and financial loss at U.S. critical infrastructure operators — water and wastewater utilities, energy sector assets, and government facilities — with no chokepoint involved and no barrel displaced. The July 2026 update establishes that safety and alarm logic has been defeated at victim sites while normal operation continued.

This is a distinct exposure with distinct characteristics. It reaches domestic assets directly rather than through a commodity price. It concentrates in regulated utilities and municipal issuers rather than in energy futures. Its loss distribution is driven by remediation cost, regulatory response, service interruption, and — in the tail — liability arising from a safety-system defeat that produces physical harm. None of that correlates cleanly with Brent, which means a portfolio hedged against the scenarios in Section 9.2 is not thereby hedged against this at all.

The point is not to size the exposure, which this paper cannot do. It is that an analyst modelling Iranian escalation risk exclusively through energy prices has specified the problem too narrowly, and the government has published the evidence that it is too narrow.

9.6 Actor calculus, briefly

Two corrections to the common framing are worth stating, because they bear on which stabilization branch is plausible.

Russia and China are not aligned on this. Russia is a net hydrocarbon exporter that benefits directly from elevated prices, from Western attention diverted to the Gulf, and from European fiscal and political bandwidth consumed by energy costs. Its preference is managed instability: hot enough to bleed, not hot enough to collapse demand.

China is the world's largest crude importer and the principal buyer of Iranian exports. Every week the chokepoints remain contested, China pays more and absorbs supply-chain disruption. It wants Iran preserved as a strategic partner and the straits *open*. It therefore holds both the strongest lever over Tehran and the strongest incentive to use it.

The divergence is the exploitable seam. If any actor brokers the stabilization branch, it is most likely China, and it will price the service in diplomatic capital. Analysts modeling the two as a monolithic bloc will misread the most probable path to de-escalation.

10 Conclusion

The argument of this paper can be compressed to a single reordering.

The public debate asks whether Iran acquires a nuclear weapon and, if so, whether its leadership is rational enough to be deterred. Both questions are reasonable and neither is the

operative one. The evidence assembled here suggests that the Islamic Republic has been, across forty-seven years, an instrumentally rational and survival-oriented actor, and that the doctrinal case for apocalyptic motivation does not survive contact with the sources. It also suggests that this reassuring finding has been substantially undermined — not because anyone's beliefs changed, but because the institution that generated the evidence has been dismantled.

What remains is a system in which authority is contested, senior command has lost roughly ninety-three percent of its accumulated in-post experience in fourteen months, ground forces operate under doctrine authorizing independent action when the center is unreachable, the surviving cadre is selected by attrition for hardness over caution, and the people through whom de-escalation would have been transacted have been removed. Into that system, a nuclear capability may arrive.

The sequencing is the finding. A nuclear Iran that consolidates first and arms second is a chronic, expensive, bounded problem of the kind the international system has absorbed before. A nuclear Iran that arms during institutional collapse is a different distribution entirely, and it is the one current conditions are manufacturing.

This finding cuts across the usual policy alignments. Those who favor continued military pressure must account for the possibility that the pressure is producing the more dangerous of the two configurations. Those who favor restraint must account for the fact that restraint, on Pape's own framing, delivers a nuclear Iran on a compressed timeline. Neither camp gets a clean answer from this analysis, and analysis that hands every camp a clean answer is usually not analysis.

Two things follow that are not consolation.

The first is that the failure modes identified here are, individually, the most tractable part of the problem. Deterrence architecture — deconfliction channels, extended warning time, accepted intermediaries, launch-detection transparency — is what converted the Cold War from a coin flip into a survivable standoff, and it was built by people who worked out in unpleasant detail exactly how the failures occur. None of that infrastructure exists in the Israel–Iran dyad. Its absence, not the arsenals, is the variable most amenable to change.

The second is that this infrastructure has historically been built only after a near-miss. The Moscow–Washington hotline followed the Cuban Missile Crisis; the Incidents at Sea Agreement followed years of dangerous encounters. Systems learn from surviving their own worst moment. The danger of the configuration described here is that its worst moment may not leave a learning period behind it.

The hundred-year test

The standing test applied to foundational decisions in this research programme is whether someone examining the choice in one hundred or one thousand years would conclude it was the right thing to do.

Applied here, the test does not resolve the strategic question — reasonable people will weigh the counterfactual risks of pressure and restraint differently, and this paper does not adjudicate between them. But it resolves an analytical one. A future examiner would ask whether the people who could see the command-fragmentation problem forming said so plainly while the decisions were still open, or whether they confined themselves to the more comfortable debate about intentions and capabilities. That question has an unambiguous answer, and it is available to anyone willing to write the uncomfortable version down.

Stop asking whether Iran gets a weapon. Start asking whether anyone will be able to give a stand-down order and have it obeyed. The first question is being answered by physics and enrichment cascades on a timeline no analyst controls. The second is being answered right now, by decisions that are still open, and it is the question on which the outcome actually turns.

At the Dinner Table

Every argument about this comes down to whether the people in charge can be reasoned with. That is the wrong worry. The right worry is whether there is anyone in charge to reason with — and that is a question about org charts and chains of command, not about theology.

A Command Attrition Detail

In-post tenure figures used in Section 6.2, with sources as noted in the References.

Table 8: Senior IRGC command tenure, before and after June 2025

Command	Incumbent	Tenure (mo.)	Terminated
Commander-in-Chief	H. Salami	≈ 74	Jun 2025, strike
Commander-in-Chief	M. Pakpour	≈ 8.5	Feb 2026, strike
Commander-in-Chief	A. Vahidi	≈ 4.8	incumbent
Aerospace Force	A. A. Hajizadeh	≈ 188	Jun 2025, strike
Aerospace Force	M. Mousavi	≈ 13.3	incumbent
Ground Forces	M. Pakpour	≈ 194	promoted Jun 2025
Ground Forces	M. Karami	≈ 13.1	incumbent

Accumulated in-post tenure, three commands, 1 June 2025: $74 + 188 + 194 = 456$ months.
Accumulated in-post tenure, same commands, 23 July 2026: $4.8 + 13.3 + 13.1 = 31.2$ months.
Reduction: $1 - (31.2/456) = 93.2\%$; ratio $456/31.2 \approx 14.6$.

Note. Tenure figures are computed from publicly reported appointment and termination dates and are approximate to within roughly one month. The measure is in-post tenure only; all current incumbents are career officers with multi-decade total service.

B Source Confidence

Claims in this paper fall into four confidence bands, marked here rather than hedged inline.

References

Theory and scholarship

Table 9: Confidence classification of principal claims

Band	Content
Established	Historical events prior to 2025: the 1988 ceasefire, the 1983 Hojjatieh dissolution, classical Twelver source content, Cold War near-miss incidents, threat-rigidity findings, power-transition and stability–instability theory.
Well-reported	2025–2026 events: the leadership losses in Appendix A, the February 2026 strikes, the Larijani killing, the Houthi blockade declaration, current Brent pricing. Multiple independent outlets; details may be revised.
Single-source	The Mojtaba Khamenei succession, reported by Iran International as exclusive and described as withheld from public announcement. Treated as probable but not confirmed; the argument does not depend on the identity of the successor, only on the contested character of the process.
Official assessment	The attribution in Section 7.4 of the cyber activity to Iranian-affiliated actors and to an IRGC-affiliated element, and the assessment that targeting escalated in connection with hostilities. This is the joint published judgement of seven U.S. government agencies — the strongest provenance available short of independent verification, and not the same thing as it. The infrastructure-lifetime figures are the author’s computation from the published indicator tables, with confounds stated in Section 7.4.
Author’s thesis	The sequencing argument, the selection-filter mechanism, ideology-as-fallback, the four indicators, all probability estimates, and the unpriced-layer claim. These are analysis, not reporting, and should be evaluated as such.

- Arjomand, S. A. *The Shadow of God and the Hidden Imam*. University of Chicago Press, 1984.
- Amanat, A. *Apocalyptic Islam and Iranian Shi'ism*. I.B. Tauris, 2009.
- Blair, B. G. *The Logic of Accidental Nuclear War*. Brookings Institution, 1993.
- Feaver, P. D. *Guarding the Guardians: Civilian Control of Nuclear Weapons in the United States*. Cornell University Press, 1992.
- Gilpin, R. *War and Change in World Politics*. Cambridge University Press, 1981.
- Narang, V. *Nuclear Strategy in the Modern Era: Regional Powers and International Conflict*. Princeton University Press, 2014.
- Organski, A. F. K., and J. Kugler. *The War Ledger*. University of Chicago Press, 1980.
- Pape, R. A. *Bombing to Win: Air Power and Coercion in War*. Cornell University Press, 1996.
- Pape, R. A. *Dying to Win: The Strategic Logic of Suicide Terrorism*. Random House, 2005.
- Sachedina, A. *Islamic Messianism: The Idea of the Mahdi in Twelver Shi'ism*. SUNY Press, 1981.
- Sagan, S. D. *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons*. Princeton University Press, 1993.
- Sagan, S. D., and K. N. Waltz. *The Spread of Nuclear Weapons: A Debate Renewed*. W. W. Norton.
- Stephenson, K. "The Quantum Theory of Trust." *Strategy+Business*, 2002; and associated work on organizational network analysis, Netform, 1990s–2000s.
- Snyder, G. H. "The Balance of Power and the Balance of Terror." In P. Seabury, ed., *The Balance of Power*. Chandler, 1965.
- Staw, B. M., L. E. Sandelands, and J. E. Dutton. "Threat-Rigidity Effects in Organizational Behavior: A Multilevel Analysis." *Administrative Science Quarterly* 26, no. 4 (1981): 501–524.

Contemporary analysis and reporting

- Atlantic Council. "The Houthis just announced a blockade on Saudi Arabia. What does it mean for the global economy?" July 2026.
- Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, National Security Agency, Environmental Protection Agency, Department of Energy, U.S. Cyber Command Cyber National Mission Force, and Department of the Treasury. *Iranian-Affiliated Cyber Actors Exploit Programmable Logic Controllers Across US Critical Infrastructure*. Joint Cybersecurity Advisory AA26-097A. Published 7 April 2026; last revised 22 July 2026.
- Council on Foreign Relations. "Another Hormuz? What to Know About the Houthi Threat to the Red Sea." 20 July 2026.
- Iran International. "A wartime succession in Iran: why the IRGC backed Mojtaba Khamenei." 4 March 2026.
- Lewis, B. "August 22." *The Wall Street Journal*, 8 August 2006.
- Pape, R. A. *The Escalation Trap*. Substack, 2026, and associated interviews (*Diary of a CEO; Triggernometry; Amanpour and Company*, 25 March 2026).
- Time*. "The New Leaders Calling the Shots in Iran." 6 May 2026.
- Yildiz, G. "Khamenei Is Dead — Iran's Missiles Keep Flying Without Him." *Forbes*, 28 February 2026.
- Market data: Brent crude settlement and intraday pricing, 15–23 July 2026, via Bloomberg, CNBC, and Trading Economics. World Bank *Commodity Markets Outlook* 2026 baseline.

This paper is research and commentary. It is not investment advice and does not constitute a recommendation to buy or sell any security or commodity.