Deterrence Elasticity: Building Scalable Strategic Capacity in the Simultaneity Era

by

Joseph H. Lyons

Introduction

The defining strategic challenge of the emerging multipolar nuclear era is not "one more threat," but simultaneity—two peers pressuring the United States at once. The 2026 National Defense Strategy makes this explicit, warning that it is "only prudent" to prepare for adversaries acting together, or opportunistically across multiple theaters. In that world, deterrence credibility depends on more than having a modern nuclear "Triad" on paper. It depends on whether the United States can generate and sustain usable capability quickly enough to deny adversaries – including adversaries who *collaborate* in aggression, or who engage in it opportunistically once the U.S. forces are occupied elsewhere – a window in which temporary force advantages can be converted into coercive leverage.

This essay proposes a simple organizing concept for that problem: deterrence elasticity. Elasticity is the ability to expand usable, survivable strategic capacity on *decision-relevant timelines* – measured in single-digit years rather than decades. Elasticity is not a substitute for modernization. Modernization replaces ageing systems, whereas elasticity adds margin in terms of the systems one already has.

Here, "fast enough" does not mean responding to a crisis next month. It means adding usable capacity inside the period when the strategic environment can deteriorate meaningfully but before long-cycle modernization programs mature. For the purposes of this essay, that means roughly five to eight years: long enough for production-rate changes, sustainment expansion, and industrial ramp decisions to affect force availability, but short enough that new submarine hulls, new warhead infrastructure, or wholly new strategic systems are unlikely to change the balance.

The 2025 National Security Strategy frames strategy as the hard work of connecting objectives to available means, rather than assembling a laundry list of aspirations. Deterrence elasticity forces that connection by asking a blunt question: *What strategic deterrent capacity can the United States actually scale fast enough for this to matter as the threat environment evolves in this era of simultaneity?*

The Simultaneity Problem and the Case for Deterrence Elasticity

The traditional posture debate often proceeds as if time is available: that crises can be treated as sequential rather than overlapping, forces can be shifted after the initial shock, and industrial or nuclear-enterprise constraints can be worked through later if the threat worsens. The simultaneity problem erodes those assumptions because overlapping pressures compress the time available to reposition forces, recover lost readiness, or expand production before coercive leverage hardens. In practice, the problem is not that the United States can never "catch up." It is that catch-up may occur too late to matter if two peers create overlapping demands before new capacity can be fielded.

Elasticity is the posture attribute that determines whether that margin can be created. Here, margin means more than a larger inventory on paper: it includes excess deployable capacity, production headroom, readiness headroom, and reserve or upload flexibility that can be converted into usable force before a crisis hardens into a new status quo. A force can be survivable and still be inelastic; it can be

modern and still be slow – not slow to respond once employed, but slow to field additional usable capacity once the strategic environment deteriorates. In the simultaneity era, those shortcomings matter because routine bottlenecks in factories, shipyards, depots, and the nuclear enterprise can become strategic vulnerabilities when there is no longer enough time to recover lost margin.

A practical way to think about elasticity is a timeline test. Some posture adjustments can plausibly affect usable capacity in *five to eight years* – for example, increasing production rates on an existing aircraft program, expanding depot and sustainment capacity, or growing the training pipeline that turns inventory into operational availability. Other adjustments remain measured in decades even when the policy choice is clear: building additional ballistic-missile submarine hulls, creating new nuclear production infrastructure at scale, or fielding major new strategic systems that require long testing and certification cycles. As the Strategic Posture Commission put it in 2023, decisions "need to be made now" to address threats "arising during the 2027–2035 timeframe," and the United States and its allies must be prepared to "deter and defeat both adversaries simultaneously." That makes elasticity a near-term requirement rather than a distant aspiration.

Inventory Illusions: Total Inventory Does Not Equal Capacity

Deterrence elasticity is measured in usable capacity – what can be generated, sustained, and employed when leaders need options under time pressure. Nominal inventories can be useful for budgeting and long-range planning, but they can obscure the practical question simultaneity forces: what portion of a force is actually available for operational use, and what portion is mission-capable when it matters.

The bomber force provides a clean example. A recent Mitchell Institute policy paper by Mark A. Gunzinger and Heather Penney breaks the bomber fleet from total aircraft inventory down to primary mission aircraft, defined as the aircraft available for operational assignment after subtracting training/test aircraft and backup/attrition reserves, and finally mission-capable aircraft after

applying readiness rates. In that accounting, the Air Force's 2025 bomber inventory falls from 141 aircraft on the books to 96 primary mission aircraft, and then to roughly 47 mission-capable bombers available at any one time, as shown in **Figure 1** below. The implication is not that bombers are uniquely "unreliable." It is that thin inventories leave little margin to absorb routine maintenance cycles, training demands, and reserves without collapsing usable capacity. That current availability picture is not itself the elasticity window; it is the baseline condition from which any meaningful increase in usable capacity must begin.

This inventory illusion is what simultaneity punishes. When effective inventory is already thin, "surge capacity" has to be built deliberately – by scaling inventory and readiness together and by being honest about what share of the force is actually relevant to the mission set a given crisis demands.

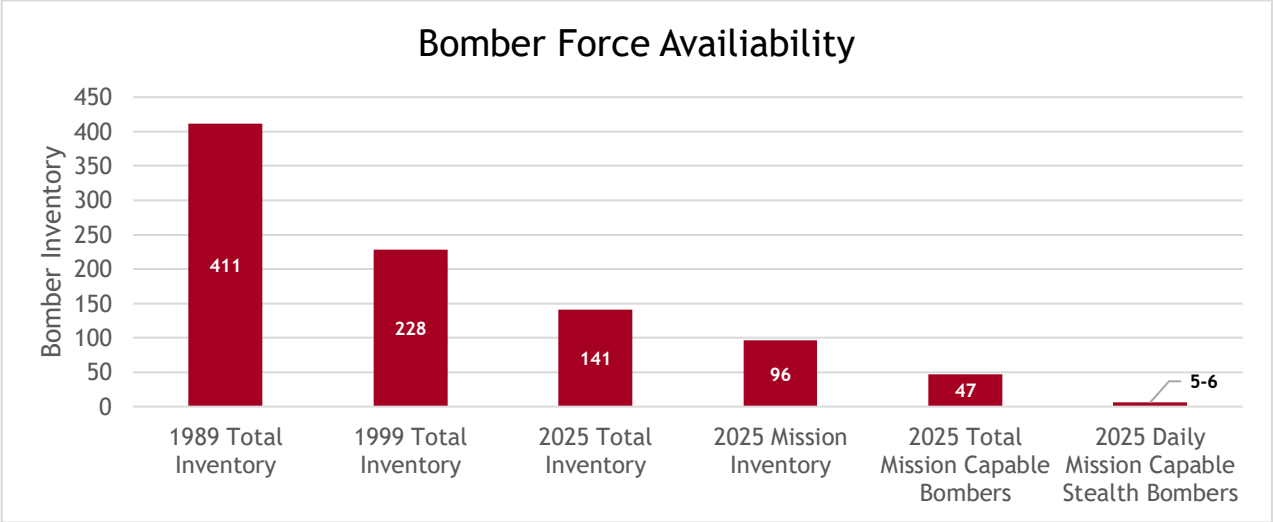

Figure 1. Bomber Force Availability.[1]

Applying the Elasticity Test Across the Triad

If simultaneity makes time the binding constraint, elasticity should be judged by a simple test: *what can be scaled in five to eight years, and what cannot?* This is not a comment on the intrinsic value of each

leg of the Triad. It is a recognition that different legs translate funding and policy decisions into usable capability on very different timelines.

Not all elasticity is the same. For the purposes of this essay, additive elasticity means the ability to field more usable strategic capacity on relevant timelines. Protective elasticity means preserving existing capacity by preventing schedule slips, transition turbulence, or enterprise bottlenecks from reducing it. That distinction matters because the Triad's legs do not contribute in the same way: some can add capacity in the near term, while others contribute mainly by protecting the force from losing margin during modernization.

Two cross-cutting constraints dominate before any platform-specific debate begins. The first is industrial capacity: the specialized workforce, facilities, and supplier networks needed to build and sustain strategic systems do not surge on demand. The second is the nuclear weapons enterprise itself. Even if leaders wanted to expand deployed capacity rapidly, the ability to produce and certify key components – particularly plutonium pits – sets a hard ceiling on how quickly the stockpile can grow or be refreshed. The Government Accountability Office (GAO) has repeatedly highlighted that the U.S. National Nuclear Security Administration still lacks a comprehensive schedule and cost estimate for the production of plutonium "pits" – the physical core of an implosion-type nuclear weapon – which is a strategic constraint in an era that prizes speed and margin.[2] In practice, even the planned 80-pit-per-year enterprise is already largely committed to the modernization program on the books – 30 war reserve pits per year at the Los Alamos National Laboratory and at least 50 at the Savannah River Plant – leaving little room to create excess margin quickly.[3]

Elasticity, however, is not limited to delivery platforms. Upload capacity and reserve stockpile management can also provide a form of strategic margin by increasing the number of warheads available for deployment on existing systems. That is a different kind of elasticity from the one emphasized here. Upload flexibility can change the payload of an already fielded force, but it does not solve the problem

of how many usable platforms the United States can generate, sustain, and signal in the simultaneity window. This essay therefore focuses primarily on delivery-system and force-generation elasticity – the part of the problem most directly shaped by industrial timelines, readiness, and operationally available capacity.

Against that backdrop, the land-based leg offers an instructive case. Sentinel is indispensable for replacing a rapidly aging system, but its near-term contribution is primarily one of protective elasticity rather than additive elasticity. In the 2027–2035 window, that means preserving continuity and transition stability: keeping the force credible during a complex modernization period and preventing schedule turbulence from becoming an operational vulnerability. GAO's work on the U.S. Air Force's "Sentinel" Intercontinental Ballistic Missile (ICBM) program emphasizes transition and schedule risk, including gaps in the Air Force's plans to manage the Minuteman-to-Sentinel transition and the risks that follow from delayed infrastructure and testing milestones.[4] In practical terms, the elasticity question for the intercontinental ballistic missile force is less about adding missiles quickly and more about ensuring the modernization transition does not create a period of reduced confidence, reduced readiness, or prolonged dependence on an increasingly strained legacy system.

The sea-based leg is even less "surgeable" in the near term. The planned *Columbia*-class nuclear-powered ballistic missile submarine (SSBN) is the most survivable element of the U.S. nuclear Triad, but submarine hull production operates on decades-long timelines, and schedule disruptions have strategic consequences because they directly affect at-sea presence. Here again, the relevant contribution is protective elasticity rather than numerical expansion. GAO has highlighted that the Columbia program lacks essential schedule insight because the shipbuilder has not conducted schedule risk analysis, and it notes persistent construction challenges.[5] The implication is not simply cost and delay; it is that the margin for error is thin. In a simultaneity era, protecting the schedule – and the industrial base that supports it – is therefore an elasticity move in the

protective sense, because it preserves the one part of the force that cannot be recreated quickly.

Taken together, Sentinel and *Columbia* illustrate that some elements of deterrence elasticity are protective rather than additive. Their near-term contribution is not the rapid creation of new deployed capacity; it is the prevention of avoidable losses in credibility, readiness, and schedule integrity while modernization is underway. That is still strategically important, because it preserves margin the United States would otherwise lose. But it also clarifies where additive near-term elasticity must come from. If the demand signal is additional usable capacity on single-digit timelines, the clearest lever is the strategic force element that can scale through production rate and readiness investments rather than new hull construction or new warhead infrastructure.

Bomber Elasticity as the Near-Term Pressure Valve

If the question is additive elasticity – what can increase usable capacity on decision-relevant timelines – the bomber leg is the clearest near-term lever. Here 'near-term' means changes that can plausibly affect usable capacity before 2035, not changes that could fully transform the force before a sudden crisis next month. Unlike new submarine hulls or major warhead infrastructure, bombers can add usable capacity by changing two variables that respond on single-digit timelines – production rate and readiness/sustainment capacity – and by shifting the mission emphasis of a dual-capable force. The Mitchell Institute's "effective inventory" breakdown shows why this matters: when training/test aircraft and mission-capable rates are accounted for, usable bomber capacity is far smaller than the top-line inventory suggests.[6]

Why Bombers are Elastic

The production-rate lever is the central difference. The Air Force's stated inventory objective for the B-21 "Raider" dual-capable (*i.e.*, conventional or nuclear payload) strategic bomber is a minimum

of 100 aircraft.[7] Production capacity can be expanded by investing in facilities, suppliers, and workforce—an industrial scaling problem that is difficult but generally faster than adding new strategic submarine hulls or new ICBM silos. Recent public reporting suggests the Air Force and Northrop Grumman are actively negotiating a production expansion agreement for the B-21 enabled by congressional funding, with company leadership indicating a deal could arrive on the order of months rather than years.[8]

The second lever is readiness. In a small force, readiness is not merely a maintenance problem; it is a margin problem. When inventories are thin, routine depot cycles and training demands consume a larger share of the available fleet, making it harder to preserve both forward options and nuclear mission assurance. The solution is not "buy more aircraft and hope readiness improves." It is to pair force growth with sustainment capacity: parts, facilities, workforce, and modernization paths that keep mission-capable rates from becoming the hidden ceiling on usable capacity.

Bomber elasticity also derives from dual-capability. Unlike strategic systems tied more narrowly to a single mission, bombers can shift the emphasis of the force between conventional strike and nuclear signaling as conditions change. That flexibility is especially valuable in a simultaneity environment because it gives leaders more ways to allocate scarce capacity across theaters without immediately forcing an all-or-nothing choice. But it also creates arbitration pressure: the same aircraft, crews, and support infrastructure may be pulled toward competing conventional and nuclear demands at the same time. Bomber elasticity, then, is not just about adding airframes. It is also about how a dual-capable force is apportioned, sustained, and signaled under stress.

B-21 Acquisition Bands and Elastic Deterrence Alternatives

The most useful way to discuss potential B-21 fleet sizes is to treat them as *decision bands* tied to assumptions, not as a single procurement target. In the simultaneity era, the question is not

whether 100, 145, 200, or 288 is the "right number" in the abstract. It is what each band buys in terms of usable options, what assumptions it relies on, and whether it creates a hedge that can absorb surprise.

1) *Band 1: 100 as a baseline floor.* The Air Force's B-21 fact sheet describes an inventory objective of a minimum of 100 aircraft.[9] That floor matters because it anchors program cost and basing plans, but it is not the same as an elastic deterrence posture. Congressionally mandated force design work by the Center for Strategic and Budgetary Assessments (CSBA) warned that a planned force of 100 B-21s could fall short of penetrating strike capacity needed for even a single major high-end conflict, and it tied meaningful capacity to ramping production toward 10–20 aircraft per year by the late 2020s.[10] In other words, 100 is a starting point—not a hedge.

2) *Band 2: 145 as a hedge step under simultaneity.* The figure of 145 B-21s has entered mainstream policy discourse through STRATCOM testimony and subsequent analysis.[11] It is best understood as a hedge step: a force large enough to improve operational flexibility and resilience compared to the minimum program objective, but still bounded by the reality that "usable capacity" depends on readiness, enablers, and how many aircraft must be withheld for nuclear mission assurance.

3) *Band 3: 200 as a minimum for sustained penetrating capacity.* The Mitchell Institute goes beyond "bomber relevance" arguments and quantifies the implication of a thin effective inventory: if the force must deny sanctuary, sustain long-range strike demand, and contribute to homeland defense and nuclear deterrence simultaneously, the bomber inventory must be able to

generate roughly two to three times today's long-range strike sortie capacity. Their recommendation is a bomber inventory of 300 or more aircraft, including at least 200 penetrating B-21s, and they argue that 145 is "not enough" and "not fast enough."[12] Whether one agrees with their campaign assumptions, the key analytic contribution is the link between effective inventory and sustained sortie generation under concurrency.

4) *Band 4: 288 as a high-band elasticity hedge.* The upper band should be treated carefully, but it is not an ad hoc number. CSBA's congress-required future force inventory recommended 383 total bombers, including 288 B-21s (total aircraft inventory) and 206 B-21 primary mission aircraft.[13] This is best read as an elasticity hedge case: a force sized to sustain strategic deterrence, defend the homeland, and be prepared to defeat major aggression by China and Russia, rather than planning to handle those demands in sequence. It does not mean the Air Force should announce "288" as a procurement target tomorrow. This band is best interpreted as an upper-bound hedge case for simultaneity planning, not a near-term program objective. It means that credible simultaneity planning pulls the analysis toward a larger, more resilient penetrating bomber force than the minimum program objective.

Elasticity Enablers and the "Readiness Critique"

Bomber elasticity is not simply a function of airframes. The limiting factor can become tankers, basing, munitions stockpiles, maintenance capacity, and the ability to generate sorties at scale. That is where the readiness critique arises: if mission-capable rates are low, why buy more aircraft? The correct answer is that low readiness in a thin force is often an indicator of insufficient margin and sustainment

capacity, not proof that the platform category is a poor investment. A small fleet cannot absorb depot and modernization downtime without translating routine friction into strategic scarcity. The Mitchell inventory breakdown makes that point implicitly: the effective force collapses because the total inventory is too small to sustain both routine demands and crisis requirements at the same time.[14]

The implication for policy is straightforward. If bombers are the near-term pressure valve, elasticity has to be funded as a package—and that package carries tradeoffs. Increasing B-21 production without expanding tanker capacity, basing and infrastructure resilience, munitions stocks, maintenance throughput, and training capacity simply shifts the bottleneck from the factory to the flightline. Likewise, bridging choices (like extending legacy aircraft) can preserve margin, but they impose opportunity costs in sustainment funding and workforce bandwidth that compete with modernization. Owning those tradeoffs up front is the point of the elasticity lens: it forces planners to connect added capacity to the enabling ecosystem required to make that capacity usable.

That package can be debated on cost and priorities, but it is at least measurable on a timeline relevant to the simultaneity problem.

Bridging Capacity: Don't Retire What You Cannot Replace

Elasticity also requires bridging choices that preserve penetrating capacity while the new force grows. Other legacy capabilities—including retained sea-based capacity or selected legacy warhead options—can also provide forms of hedge or bridging elasticity, but the focus here is scalable strategic capacity across the Triad, with the bomber force as the principal case study. Both CSBA and the Mitchell Institute make versions of the same argument: do not accelerate B-2 retirement if it creates a temporary "bathtub" in penetrating capacity before the B-21 force is sizable and fully operational. CSBA recommended sustaining and modernizing the B-2 until approximately 2040.[15] Mitchell recommended refraining from retiring B-2s until B-21 inventories surpass 100 aircraft and are fully

operational in the 2030s.[16] Bridging is not a substitute for growth, but it is an elasticity move because it preserves options during the exact window when simultaneity risk is rising and modernization concurrency is most stressed.

A Decision Rule for Bomber Elasticity

The bomber leg becomes the near-term pressure valve when policy choices treat elasticity as the objective rather than treating the program of record as the ceiling. The decision rule is simple: *prioritize the options that add usable, survivable capacity on single-digit timelines while protecting the schedule integrity of the systems that cannot surge.* In practical terms, that means treating the B-21 not only as a replacement program but as an elasticity lever that can expand the effective inventory – if production, readiness, and enablers are scaled together.

Conclusion

Deterrence in the simultaneity era is less about declaring new inventory targets than about restoring and expanding margin – usable capacity that can be generated, sustained, and signaled under time pressure. Deterrence elasticity offers a practical way to judge posture options against that reality. The key question is not simply what the United States intends to modernize over decades, but what it can scale within the window in which strategic conditions can change.

Through that lens, the Triad's legs contribute in different ways. Sentinel and Columbia remain essential to long-term survivability and strategic stability, but their near-term contribution is primarily protective elasticity – risk reduction through schedule protection, transition stability, and preservation of enterprise capacity. The bomber force is distinct because it offers the clearest form of additive near-term elasticity, while also allowing leaders to shift the emphasis of a dual-capable force as conditions change, provided that enablers and readiness are funded as part of the package.

The decision rule is straightforward: prioritize investments that measurably increase usable, survivable capacity on decision-relevant timelines while protecting the schedule integrity of systems that cannot be expanded quickly. If the United States adopts that standard, debates about force structure become less about advocating for a platform and more about matching strategy to means in the environment the 2026 National Defense Strategy describes: one where simultaneity is the default condition, not the exception.

* * *

About the Author

Maj. Joseph H. Lyons *is an active-duty United States Air Force officer and combat-experienced B-52H instructor pilot. A career bomber aviator, he has supported the full range of global strike operations. He is currently a Department of the Air Force Fellow at the Defense Threat Reduction Agency and a doctoral candidate at Missouri State University's School of Defense and Strategic Studies. His doctoral and 2025-2026 Fellows Program research examines how the United States should adapt its nuclear posture to credibly deter multiple nuclear-armed adversaries – concurrent deterrence in a multipolar era.*

The views expressed in this article are the author's and do not necessarily reflect the official policy or position of the Department of the Air Force, the Defense Threat Reduction Agency, the Department of Defense, or the U.S. Government.

Notes

1 *Visualization adapted by the author from* Col Mark A. Gunzinger & Heather R. Penney, "Strategic Attack: Maintaining the Air Force's Capacity to Deny Enemy Sanctuaries," Mitchell Institute Policy Paper, Vol. 64 (February 2026), Figure 3 & accompanying discussion, https://www.mitchellaerospacepower.org/app/uploads/2026/02/Strategic_Attack_Denying_Sanctuary_Policy_Paper_64.pdf.

2 U.S. Government Accountability Office, *Nuclear Weapons: NNSA Does Not Have a Comprehensive Schedule or Cost Estimate for Pit Production Capability*, GAO-23-104661 (GAO, 2023), https://www.gao.gov/assets/gao-23-104661.pdf.

3 U.S. Government Accountability Office, *Nuclear Security Enterprise: Assessments of NNSA Major Projects*, GAO-26-107777 (GAO, February 2026), 15–16, 38, & 55, https://www.gao.gov/assets/gao-26-107777.pdf

4 U.S. Government Accountability Office, *ICBM Modernization: Air Force Actions Needed to Expeditiously Address Critical Risks to Sentinel Transition*, GAO-25-108466 (GAO, 2025), https://www.gao.gov/assets/gao-25-108466.pdf.

5 U.S. Government Accountability Office, *Columbia Class Submarine: Program Lacks Essential Schedule Insight amid Continuing Construction Challenges*, GAO-23-106292 (GAO, 2023), https://www.gao.gov/assets/gao-23-106292.pdf.

6 Gunzinger & Penney, "Strategic Attack," Figure 3.

[7] U.S. Air Force, "B-21 Raider," fact sheet, updated March 21, 2025, inventory objective ("minimum of 100 aircraft"), https://www.af.mil/About-Us/Fact-Sheets/Display/Article/2682973/b-21-raider/.

[8] Thomas Novelly, "Northrop CEO: deal to accelerate B-21 production could arrive in months," *Defense One*, January 28, 2026, https://www.defenseone.com/business/2026/01/northrop-grumman-ceo-deal-accelerate-b-21-production-could-arrive-months/411021/.

[9] U.S. Air Force, "B-21 Raider."

[10] Mark Gunzinger et al., *An Air Force for an Era of Great Power Competition* (Center for Strategic and Budgetary Assessments, 2019), xi, https://csbaonline.org/uploads/documents/CSBA_AFAIS_Report_v9.pdf.

[11] Jennifer DiMascio & Anya L. Fink, *U.S. Strategic Bombers*, IF12945 (Congressional Research Service, updated January 6, 2026), https://www.congress.gov/crs-product/IF12945.

[12] Gunzinger & Penney, "Strategic Attack," 17 & 21.

[13] Gunzinger et al., *An Air Force for an Era of Great Power Competition*, ix (Table 2).

[14] Gunzinger & Penney, "Strategic Attack," Figure 3.

[15] Gunzinger et al., *An Air Force for an Era of Great Power Competition*, xi.

[16] Gunzinger & Penney, "Strategic Attack," 21.