


ARGUS-PD/POD-7K Programme Capstone

Operational Plausibility, Build Sequence, Cost Exposure, Governed Control and Strategic Space Advantage

Christopher I. V. Farmer

Standalone strategic programme study • Independent preprint

JULY 2026



proof by year 8 • direct mode \$18.1B • P80 \$27.1B • governed strategic advantage

Abstract

The first three ARGUS-PD/POD-7K papers establish a controlled theoretical programme: a complete moving-mass boundary and projectile family; a fixed seven-kilometre integrated architecture; and a coupled feasibility model with uncertainty propagation and falsification criteria [1–3]. They do not establish that a full system can presently be built, afforded, deployed or authorised. This programme capstone addresses that gap without converting a research concept into an operational counterspace plan. It asks what a government–industry consortium would have to build, in what order, at what indicative cost, under which control regime, and with what credible strategic outcome.

The analysis reaches four principal findings. First, eight years is a defensible horizon for a decision-grade research programme, not for assuming an operational seven-kilometre spacecraft. By year eight, an unusually well-funded programme could plausibly deliver validated component models, representative pulse and moving-interface experiments, an integrated ground-based demonstrator at a deliberately reduced scale, and a non-energetic orbital assembly-and-control pathfinder. A first supervised civil capability is more credibly placed in a conditional 12–16-year window. The full reference architecture remains unschedulable until mass, pulse-power, thermal, structural, projectile-survival, in-space assembly and legal authorisation gates close.

Second, commercial batteries can be useful buffer assets but do not solve the pulse-power problem. The 514.286 GJ reference nameplate equals approximately 36.5 current Tesla 4-hour Megapacks by stored energy, but their combined continuous output is only about 35.7 MW against a 30.6 TW mean admitted-power screen during the 14 ms reference launch: a power-rate gap of about 8.57×10^5 . Sandia's 22 MJ Z facility is an important pulsed-power analogue, yet the ARGUS nameplate is about 2.34×10^4 times larger by stored energy [7–9]. This is not evidence of impossibility; it is evidence that batteries, intermediate stores, switches, transmission, cells, structure and controls must be qualified as distinct tiers.

Third, a line-by-line re-audit reconciles my own earlier estimate with programme risk. The mutually exclusive direct work packages have a \$18.1 billion most-likely sum. After sampling their three-point ranges and a separate integration-and-reserve factor, the constant June-2026-dollar screen produces a P50 of \$24.3 billion and a P80 of \$27.1 billion for the eight-year proof programme. The risk-loaded mathematical input envelope is \$10.1–48.45 billion. These are research-planning assumptions, not quotations or an acquisition baseline; NASA and GAO practice requires a complete work-breakdown structure, supplier data, schedule logic and joint cost-and-schedule confidence before commitment [10–13]. Full deployment is deliberately not given a false point estimate at this maturity.

Fourth, the strategic hypothesis requires precise but direct language. The competition is already public: the United States seeks space superiority; China describes a journey towards a strong space presence; the European Union seeks resilience, strategic autonomy and the role of a global space power; the United

Kingdom seeks to preserve strategic advantage; NATO treats space as an operational domain; and India is expanding sovereign and commercial space capacity [24,28,33–37]. United States Space Force doctrine defines space superiority as freedom to operate at a chosen time and place without prohibitive interference while denying the same to an adversary [24]. ARGUS alone cannot establish that complete condition. If its technical gates close, however, the programme could confer substantial strategic space advantage through pulsed-power competence, launch and in-space-assembly capacity, space-domain awareness, materials science, logistics, standards, alliance leadership and control of a high-value planetary-defence service. Those benefits are strongest when the capability is transparent, civil-led, independently supervised and internationally consulted. An overt national “supremacy machine” would increase counterspace ambiguity, vulnerability, balancing behaviour and legal risk, potentially destroying the governance advantage it sought to create.

The investable proposition is therefore a stage-gated national research programme with commercially valuable spillovers and explicit stop rules. Funding buys verified knowledge and enabling infrastructure at every gate; continued scale-up is earned by evidence.

Keywords: programme architecture; planetary defence; pulsed power; in-space assembly; technology readiness; cost uncertainty; space governance; strategic advantage; dual-use systems.

1. Executive proposition

Decision proposition. *Fund an 18-month formulation tranche that creates the independent authority, complete work-breakdown structure, digital thread, test plan and supplier evidence needed for Gate G0. Do not authorise full-scale procurement, orbital energetic operation or a claim of deployability. Preserve the right to narrow, pause or terminate every non-closing branch.*

This is a proposal for controlled capability creation rather than a promise that the current reference system will work. Its value is asymmetric. If the central mechanism eventually closes, the sponsor acquires a strategically important civil-space infrastructure capability. If it does not, the sponsor still obtains reusable assets in pulsed power, dynamic materials, hypervelocity diagnostics, autonomous assembly, precision metrology, high-assurance control and planetary-defence modelling.

The initial authorisation should have a ceiling of approximately \$1.0–1.8 billion over 18 months. It should procure evidence, not a production design. Release of later tranches should require:

1. experimentally measured performance against pre-registered thresholds;
2. an independent technical and safety review;
3. a refreshed cost-and-schedule confidence analysis;
4. legal authority and continuing-supervision arrangements; and

5. public or internationally shareable evidence sufficient to distinguish a civil research programme from covert counterspace preparation.

The programme's most credible strategic message is simple: a sponsor intends to lead the industrial and governance architecture for planetary defence and very-large space infrastructure. It is not asking others to trust an unverifiable claim of benign intent.

2. What the existing research establishes

2.1 Controlled inheritance from Papers I–III

The capstone treats the prior papers as controlled inputs rather than marketing collateral.

Prior work	What it establishes	What remains unestablished
Paper I [1]	complete 0.300 kg moving-mass boundary; projectile family; energy–momentum Pareto trade; material and flight-environment variables	survivable moving article; selected flight material; release accuracy
Paper II [2]	fixed 7,000 m architecture; ideal force-area-field bounds; 6,000 physical module-bay hypothesis; energy, thermal, structure and governance ledgers	representative cell; physical module closure; structure and cost baseline
Paper III [3]	coupled mass-area-energy iteration; covariance; sensitivity; evidence levels; quantitative stop rules	integrated hardware; validated target response; mission authorisation
Launcher study [4]	variable-energy modes; physical module interpretation; rear counter-impulse arithmetic; dust, atmosphere and application envelopes	buildable launcher design; trajectory survival; operational mission plan

For the reference moving assembly,

$$E_k = \frac{1}{2} m v^2 = 150 \text{ GJ}, p = m v = 300 \text{ kNs}, \quad (1)$$

and, for constant acceleration over $L = 7,000 \text{ m}$,

$$t_L = \frac{2L}{v} = 14 \text{ ms}, \dot{F} = \frac{E_k}{L} = 21.429 \text{ MN}. \quad (2)$$

At an assumed electrical efficiency of 35% and a 20% installed reserve, the reference admitted energy is 428.571 GJ and nameplate energy is 514.286 GJ [1–4]. These are conservation and bookkeeping results. They are not hardware specifications validated by test.

2.2 The principal physical blockers

The programme is dominated by coupled blockers, not by a shortage of conceptual diagrams.

- **Moving interface.** No experiment demonstrates that the complete moving article can tolerate the acceleration, electromagnetic coupling, release transient and thermal environment across the reference envelope.
- **Pulsed power.** Bulk energy, pulse power, switching, transmission, synchronisation and cell coupling are separate requirements. Satisfying one does not satisfy the others.
- **Structure and reaction.** Momentum must leave the launcher-platform system. A rear valve or counter-mass can export opposing momentum, but it does not erase transient loads. For the lower-speed T-DART case already analysed, 3 MN s on a 2,000 t platform gives 1.5 m s^{-1} uncompensated motion; ideal matching with 4.5 km s^{-1} expelled mass requires 666.7 kg per event [3,4].
- **Flight survival.** The flight path is not empty. Dust, residual gas, impact plasma and charge exchange are trajectory-dependent hazards. The published fracture calculation is intentionally a lower-bound crack-energy screen, not a ballistic limit. A 50 mm cross-section also sweeps about 20.3 kg of atmospheric column in a vertical Earth passage; an intact atmospheric mode is therefore excluded [1,3,4].
- **Materials.** Boron-doped diamond remains a functional-material candidate, not a selected monolithic structure. CVD/PCD diamond, diamond-SiC, toughened SiC, boron carbide, tungsten, carbon-rich sacrificial layers and graded hybrids require matched dynamic tests [1,4].
- **Assembly and control.** A seven-kilometre flexible structure is far beyond ordinary spacecraft integration. Environmental magnetic torque can assist slow momentum management in LEO but does not provide the required rapid reference slew [6].

These blockers determine the build order.

3. Technology and market reality

3.1 Batteries are the buffer tier

Tesla currently lists a 4-hour Megapack at 3,916 kWh and 979 kW, or approximately 14.10 GJ of stored energy [9]. A simple energy-equivalent count is

$$N_{MP} = \frac{514.286 \text{ GJ}}{14.098 \text{ GJ}} = 36.48. \quad (3)$$

Using 36.48 equivalent units for comparison, not procurement, gives 35.7 MW combined output. The mean admitted-power screen is

$$\dot{P}_{admit} = \frac{428.571 \text{ GJ}}{0.014 \text{ s}} = 30.612 \text{ TW}, \frac{\dot{P}_{admit}}{P_{MP, fleet}} \approx 8.57 \times 10^5. \quad (4)$$

The conclusion is not that batteries are useless. Commercial grid batteries can recharge, buffer and stabilise a slow primary-energy tier. They cannot directly serve as the millisecond pulse tier. The power architecture must be hierarchical: grid or generation; electrochemical buffer; intermediate pulsed store; switching and pulse compression; local module; launcher cell. Each boundary requires a measured efficiency, thermal ledger, fault-containment case and replacement strategy.

3.2 Public pulsed-power comparators

Sandia reports that Z stores up to 22 MJ and can deliver approximately 80 TW to an experiment on a roughly 100 ns scale [7]. Its Marx banks, intermediate water capacitors, pulse-forming lines and triggered switches illustrate why stored energy and delivered pulse are different engineering layers [8]. Sandia has also discussed a possible 130 MJ stored-energy facility for future magnetically driven high-energy-density research [30].

The ARGUS reference nameplate is

$$R_z = \frac{514.286 \text{ GJ}}{22 \text{ MJ}} = 2.3377 \times 10^4 \quad (5)$$

times Z by stored energy. The comparison is not like-for-like in pulse shape, load or topology. It nevertheless eliminates the idea that an existing capacitor catalogue can be multiplied directly into a complete solution. The right question is whether a progressively scaled cell-and-string architecture can preserve efficiency, timing, insulation, force transfer and fault containment as energy grows.

3.3 Launch and in-space assembly

SpaceX describes Starship as designed to carry more than 100 metric tonnes to orbit in a fully reusable configuration [15]. That capacity materially changes the logistics of large spacecraft but does not close ARGUS mass. Neither the active launcher nor its energy, thermal, shielding, structure, robotics and consumables have a credible integrated mass estimate.

NASA lists the International Space Station at 419,725 kg with a 94 m truss [16]. That comparison should not be used as a linear cost or mass model, but it exposes the unprecedented length and integration class of a seven-kilometre article. NASA's 2025 in-space servicing, assembly and manufacturing survey documents a growing ecosystem rather than a mature kilometre-class supply chain [17]. SpiderFab and related concepts show that on-orbit fabrication can improve stiffness-to-mass and stowed volume, but they are technology directions, not qualification evidence for a pulsed-power structure [18]. NASA's cancellation of OSAM-1 is an

equally important negative precedent: ambitious assembly programmes can fail on cost, schedule and integration despite strong strategic logic [19].

The build strategy must therefore separate:

1. a non-energetic structural and robotic assembly pathfinder;
2. a ground-based energetic demonstrator;
3. a later orbital energy and thermal demonstration; and
4. any mission-authorised civil operation.

4. Outcome architecture

4.1 What “success at year eight” means

The year-eight outcome is a verified evidence package and integrated demonstrator, not a hidden redefinition of “operational”.

Metric	Year-eight success threshold	Why it matters
critical technologies	independent assessment at approximately TRL 5–6 for the selected subscale architecture	NASA typically expects TRL 6 or higher before flight-system integration [14]
moving interface	repeatable test series with measured force, heating, release and damage distributions	replaces ideal force-density inference
pulse system	representative module and physical string with measured energy, power, timing, efficiency and fault isolation	tests the 6,000-bay architecture
integrated facility	reduced-energy ground demonstrator with closed reaction, thermal and safety ledgers	proves system interaction
materials	down-selected layered projectile architecture validated over reachable test speeds	addresses dust, flaws and interfaces
orbital pathfinder	non-energetic modular structure with assembly, metrology, repair and safe-state control	isolates ISAM risk
governance	enacted authority, licensing/supervision plan, independent authorisation and international consultation protocol	makes civil intent verifiable
cost	refreshed WBS and joint cost-and-schedule confidence for the next phase	prevents commitment by narrative

NASA defines TRL 6 as a high-fidelity model or prototype demonstrated in a relevant environment with critical scaling issues addressed [14]. A subscale ARGUS system can reach that level only for the architecture actually tested. The rating must not be transferred automatically to the full seven-kilometre system.

4.2 What “first deployable capability” means

If Gate G4 closes, years 9–16 can pursue full-scale qualification and a first supervised civil capability. “Deployable” requires:

- a mission class and safe operating envelope;
- independently verified physical and software configuration;
- range, trajectory and casualty-risk analysis held in controlled mission records rather than an open paper;
- launch, registration, licensing and continuing supervision;
- dual authorisation for any irreversible energetic action;
- safe state, abort, inspection and post-event accountability;
- international threat validation and consultation for planetary defence;
- a funded decommissioning and disposal plan.

No single date can be credible before the G4 mass and energy baseline. The central schedule is 12–16 years for a first civil capability. Eight years remains an aggressive earliest deployment decision. Full reference-scale performance may take longer or may be rejected by evidence.

5. Build process and industrial organisation

5.1 Workstream decomposition

The build should be organised around nine accountable workstreams.

Workstream	Build responsibility	Primary gate output
programme and assurance	sponsor authority plus independent review board	controlled requirements, hazards, configuration and decision records
digital engineering	systems integrator with national laboratories and universities	traceable model, test data, uncertainty and verification ledger
projectile materials	materials laboratories and specialised suppliers	ranked layered architectures and dynamic failure statistics
pulsed power	national laboratory–prime contractor team	representative store, switch, transmission and cell chain
structure and reaction	civil/space prime plus structural test centre	transient-load, thermal, vibration and momentum-export closure
ground demonstrators	dedicated test-site authority	progressively integrated strings and safe high-energy testing

Workstream	Build responsibility	Primary gate output
ISAM pathfinder	spacecraft/robotics prime and launch provider	modular assembly, metrology, inspection and repair
flight-science pathfinder	civil space agency and international science partners	non-operational materials/environment experiment
operations and governance	public mission authority	secure control, audit, supervision, consultation and liability arrangements

5.2 Make, buy and partner

The programme should buy commercial maturity where it exists and retain public control where consequences are irreversible.

- **Commercially buy:** conventional generation; Tesla-class or equivalent buffer storage; standard launch services; data-centre hardware; conventional spacecraft buses; cloud-isolated development environments; non-sensitive robotics and metrology.
- **Co-develop:** intermediate pulse storage; switching; superconducting or alternative cell technology; high-rate diagnostics; dynamic material architectures; autonomous large-structure assembly; high-reliability thermal and power distribution.
- **Publicly own or control:** safety case; energetic test facility; independent verification; mission authorisation; controlled trajectory analysis; audit records; civil command authority; legal supervision; incident investigation.

This division gives companies a commercial market for enabling technologies without outsourcing sovereign responsibility.

5.3 Facility sequence


The physical sequence should move from cheap falsification to expensive integration.

1. **Coupons and digital twins.** Calibrate electromagnetic, thermal, fracture, impact and flexible-body models against small experiments.
2. **Representative cell.** Demonstrate measured force and energy transfer into a moving interface with failure containment.
3. **Physical string.** Establish timing, switching, interference, structural reaction and diagnostic observability across multiple cells.
4. **Integrated ground article.** Couple energy buffer, pulse tier, string, vacuum, structure, reaction management, thermal recovery, release and independent safety controls.

5. **Non-energetic orbital structure.** Validate packaging, launch, autonomous assembly, metrology, flexible-body control, repair and disposal without an energetic payload.
6. **Relevant-environment demonstrator.** Scale only after earlier evidence closes; retain an independent authority capable of stopping the programme.
7. **Mission-specific civil deployment.** Proceed only after G4/G5 technical and governance approvals.

The sequence protects the sponsor from a common megaproject failure mode: building an impressive facility before discovering that its core interface cannot meet the required physics.

6. Stage-gated schedule



Eight-year proof programme and conditional deployment path. The first five gates run from formulation to a year-eight deployment decision; years nine to sixteen are conditional on the earlier evidence.

6.1 Gates G0–G4: years 0–8

Gate	Window	Decision evidence	Stop or redirect condition
G0	0–1	credible architecture, authority, hazards, WBS, test plan and independent review	no coherent legal authority, safety boundary or cost basis
G1	1–2	calibrated component models; ranked material, pulse and interface candidates	all candidate branches miss threshold by a predeclared margin
G2	2–4	repeatable representative cell/string data and fault containment	efficiency, damage, timing or force transfer does not scale
G3	4–6	integrated subscale closure of energy, heat, structure, reaction and release	integrated result contradicts component predictions
G4	6–8	relevant-environment ground demonstrator plus	mass, survival, cost, governance or assembly

Gate	Window	Decision evidence	Stop or redirect condition
		non-energetic orbital pathfinder	cannot support deployment formulation

NASA's life-cycle model separates phases with Key Decision Points at which a decision authority determines readiness to proceed [11,12]. The ARGUS gates adopt that discipline but are not represented as NASA approval.

6.2 Gates G5–G6: conditional years 8–16

G5 authorises full-scale qualification only after a credible mass and energy baseline exists. G6 authorises a first civil capability only after mission-specific safety, legal and international review. The distinction is essential. The sponsor may reach year eight with a valuable national laboratory and no rational case for orbital energetic deployment. That is a successful research outcome if the negative result is established before larger sunk costs.

The schedule has real precedent at both ends. DART demonstrated a planetary defence kinetic impact for a projected life-cycle cost of \$330.6 million, but it was a comparatively small spacecraft and mission [21,22]. NEO Surveyor traces from a 2016 concept study through reformulation and a baseline launch no later than June 2028; its current agency baseline commitment is \$1.6 billion [20]. A genuinely unprecedented integrated system should not be scheduled more confidently than those missions.

7. Cost exposure

7.1 Method and limitations

GAO states that reliable estimates should be comprehensive, well documented, accurate and credible [13]. NASA cost practice combines estimating methods with risk and uncertainty; its Joint Cost and Schedule Confidence Level integrates cost, schedule, risk and uncertainty rather than quoting one deterministic number [10,31].

This paper supplies an early screening model:

$$C = F_c \sum_{i=1}^9 \text{Triangular}(C_{i,L}, C_{i,M}, C_{i,H}), F_c \sim \text{Triangular}(1.00, 1.15, 1.50). \quad (6)$$

The nine work packages are defined to be mutually exclusive direct costs. The common factor F_c then represents programme-wide integration, reserve, schedule and supplier-learning uncertainty. Its minimum is unity: it cannot reduce the direct-cost draw. Constant June-2026 dollars exclude future escalation, so inflation is not loaded again into F_c . The factor prevents the false confidence that would result from treating


all workstreams as independent, but it is not a calibrated joint cost-and-schedule model. The reported percentile is

$$C_{p_q} = \inf \{ c : Pr(C \leq c) \geq q \}. \quad (7)$$

The inputs are my own planning assumptions informed by public comparator programmes, not supplier bids. Tesla's current order page publishes power, energy and physical specifications but displays price only after project quantity, site and delivery inputs; it also excludes taxes and installation [9]. High-energy capacitors, switches and representative launcher cells likewise require a defined duty cycle and supplier requests for quotation. The estimate therefore does not multiply a guessed retail battery or capacitor price. It excludes a full operational deployment, financing costs, future escalation beyond constant June-2026 dollars, classified infrastructure, war-reserve stocks and operations after the proof programme.

7.2 Workstream assumptions

Workstream	Low	Mode	High	Scope note
programme, systems engineering and assurance	0.6	1.0	1.8	independent review and control
digital engineering and test infrastructure	0.7	1.2	2.0	models, diagnostics, data and facilities
projectile materials and qualification	0.6	1.2	2.5	multi-material dynamic campaign
pulse power, switching and cells	2.5	4.5	8.0	dominant facility and scale-up risk
structure, vacuum, thermal and reaction	1.2	2.2	4.0	distinct transient-load infrastructure
integrated ground demonstrators	2.0	3.5	6.0	integration excluding cell-development cost
in-space assembly pathfinder	1.0	1.8	3.2	non-energetic orbital assembly
launch and flight-science pathfinder	1.0	1.8	3.2	spacecraft, launch and mission assurance
operations, cyber and governance	0.5	0.9	1.6	secure control and supervision



Illustrative programme cost by workstream. Bars show mode assumptions and whiskers show the low-to-high planning range.

The revised workstream boundaries remove two forms of double loading in the previous draft: pulse-cell development is no longer repeated inside the ground-demonstrator line, and orbital assembly is separated from the civil flight-science mission. Workstream high cases remain visible, but a second unbounded contingency is not hidden inside each mode.

Public analogues bound scale; none is a parametric cost-estimating relationship for ARGUS:

Public analogue	Published cost	Approx. June-2026 USD	Relevance and limit
DART projected lifecycle	\$0.331bn	not converted	lower flight-demonstration comparator; mixed-year lifecycle dollars [21]
NEO Surveyor baseline	\$1.60bn (2022)	\$1.83bn	planetary-defence flight project, not an energetic launcher [20,40]
OSAM-1 rebaseline	\$2.05bn (2022)	\$2.34bn	in-space servicing/assembly scale and overrun warning [38,40]
National Ignition Facility construction	\$3.50bn (2009)	\$5.45bn	unique high-energy research facility with a different mechanism and mission [39,40]

The CPI normalisations use the BLS CPI-U annual averages of 214.537 for 2009 and 292.655 for 2022, and the June 2026 index of 333.952 [40]. They are transparent scale comparisons, not aerospace escalation indices.

7.3 Reconciliation with the earlier \$18.1 billion estimate


The source archive contains an earlier \$18.1 billion figure. It was correctly rejected as a lifecycle cost for that paper’s obsolete 49 km architecture because it omitted material scope. This audit did not reinstate that old cost object. Starting again from the controlled fixed seven-kilometre G0–G4 proof scope, the revised mutually exclusive work packages independently sum to the same \$18.1 billion at their modes. The numerical convergence is useful, but it is not validation of the archived estimate.

The distinction is:

- **\$18.1 billion:** direct most-likely work-package sum in constant June-2026 dollars;
- **\$24.3 billion P50:** median risk-loaded eight-year proof-programme screen;
- **\$27.1 billion P80:** early confidence target for tranche planning; and
- **unpriced:** full operational seven-kilometre deployment, future escalation, financing, classified infrastructure and post-proof operations.

The raw workstream sums are \$10.1 billion low, \$18.1 billion mode and \$32.3 billion high. With the common factor, the explicit risk-loaded mathematical input envelope is \$10.1–48.45 billion. A deterministic 250,000-sample run with the published seed gives:

Statistic	Eight-year proof programme
P10	\$20.8 billion
P50	\$24.3 billion
mean	\$24.5 billion
P80	\$27.1 billion
P90	\$28.6 billion



Illustrative early cost distribution for the eight-year proof programme, with P50 and P80 markers.

The relatively narrower central distribution does not invalidate the much wider mathematical envelope. It reflects the triangular sampling assumptions, which themselves require replacement. A decision authority should manage the programme to a confidence target such as P80 only after a real schedule logic network and correlated risk register exist.

7.4 Funding profile and tranche decisions

For an illustrative P80 programme, annual funding is

$$B_y = f_y C_{P80}, \sum_{y=1}^8 f_y = 1. \tag{8}$$

Programme year	Fraction	P80 annual value
1	3%	\$0.81 billion
2	7%	\$1.89 billion
3	12%	\$3.25 billion
4	16%	\$4.33 billion
5	18%	\$4.87 billion
6	18%	\$4.87 billion
7	15%	\$4.06 billion
8	11%	\$2.98 billion

Those values should not be appropriated as one unconditional block. The programme should use tranche ceilings:

- **Formulation tranche:** 18 months, ceiling of approximately \$1.0–1.8 billion.
- **Laboratory tranche:** to G2, released only after G1 candidate selection.
- **Integration tranche:** to G3, released only after representative data.
- **Relevant-environment tranche:** to G4, released only after integrated closure.

Years 9–16 require a new cost estimate at G4. The system mass and launch architecture are too immature for an honest deployment point estimate. As a stress test, a sponsor should assume exposure comparable to a major national space campaign: NASA's contributions to Artemis were projected at \$93 billion for fiscal years 2012–2025 [29]. That comparison defines scale, not cost.

8. Control potential after deployment

8.1 Control is broader than command

The relevant question is not whether a control room can transmit a command. It is whether public authority can constrain the complete lifecycle: requirements, procurement, configuration, launch, operation, irreversible action, inspection, incident response and disposal.



Governance stack for a high-consequence civil space capability, from public authority to continuing supervision.

Layer	Required control	Evidence
public authority	explicit statutory mandate, budget authority,	published authority and accountable

Layer	Required control	Evidence
	licensing jurisdiction and treaty responsibility	minister/agency
independent authorisation	technical, legal and mission approval separated from routine operations	recorded decision with dissent and conditions
dual authorisation	two-person approval for critical or irreversible operations	NIST recognises dual authorisation as risk reduction for sensitive operations [23]
configuration assurance	signed software/hardware baseline, supply-chain provenance and controlled change	reproducible build, attestation and independent verification
safe state	de-energised, passively stable or otherwise bounded condition after fault or lost communications	demonstrated fault response
telemetry and audit	independent observation, immutable logs and post-event reconstruction	external telemetry and retained evidence
continuing supervision	inspection, licence conditions, incident reporting, sanctions and decommissioning	Article VI-compatible national supervision [5,25]

The architecture should keep detailed command paths, credentials, operational trajectories and target solutions outside the open publication record. Transparency should concern authority, constraints, evidence and accountability—not the disclosure of exploitable operational data.

8.2 International control for planetary defence

Planetary defence is inherently a shared-risk activity. NASA's strategy calls for transparency, multilateral cooperation and international engagement [20,26]. The International Asteroid Warning Network supports detection and characterisation; the Space Mission Planning Advisory Group provides a forum for space agencies to collaborate on mitigation options [27]. An ARGUS mission should therefore require:

1. independent confirmation of the threat;
2. shared target-state and uncertainty review;
3. consultation on the proposed intervention and alternatives;
4. allocation of registration, responsibility and liability;
5. notification, observation and post-event data release; and
6. pre-agreed abort and dispute procedures.

The legal companion paper explains why Outer Space Treaty Article VI imposes state responsibility and supervision but does not itself create a universal licensing or sanctions power [5]. Domestic legislation must supply the actual authority.

9. The strategic-advantage hypothesis

9.1 The race already exists

Space control is not a speculative motive invented for this project. It is a declared national and alliance objective expressed in different political vocabularies.

Actor	Published objective	Competitive implication
United States	space superiority and freedom to operate without prohibitive interference [24,28]	explicit control objective
China	a “new journey towards a strong space presence” and comprehensive space infrastructure [33]	sustained state-led industrial competition
European Union	resilience, protection, threat response, dual-use capability and a stronger role as a global space power [34]	strategic autonomy and collective control
United Kingdom	become a meaningful space actor and preserve strategic advantage with allies [35]	civil-defence integration and sovereign capability
NATO	space as an operational domain essential to deterrence and defence [36]	allied demand for assured services and resilience
India	augment national capability and create a thriving commercial space ecosystem [37]	widening industrial and launch competition

These sources do not show that every actor seeks identical military “supremacy”. They show convergence on the underlying requirements: assured access, protected services, industrial depth, autonomous decision-making, domain awareness and influence over rules. That convergence is a strategic race. A state that waits until another actor controls the launch, construction, data and governance stack must buy access on terms it did not set.

For an established or aspiring space power, non-participation carries its own costs. It risks dependence on another actor’s launch cadence, orbital infrastructure, sensing, technical standards and crisis decisions. The strategic reason to fund ARGUS research is not simply to obtain one machine: it is to buy a credible position in that wider race while there is still time to shape its industrial and legal architecture. A programme that fails its physics gates can be stopped; a state that never builds the underlying competence cannot recover strategic time on demand.

ARGUS is one possible way to accelerate several layers of that stack at once. Its attraction is concentration: one programme forces progress in energy, pulsed power, materials, heavy launch, robotic construction, sensing, high-assurance software and legal control. If the architecture scales modularly, it may produce strategic option value faster and at lower cost than building the same industrial capabilities as unrelated programmes. That is a hypothesis to be tested at G0–G4, not assumed in the financial model.

The programme is not cheap in ordinary terms. Its \$27.1 billion P80 proof screen is, however, below the cited \$93 billion scale of NASA’s Artemis contribution through fiscal year 2025 [29], while purchasing a portfolio of capabilities relevant to civil exploration, national security and commercial infrastructure. The investment argument is one of comparative affordability.

9.2 A testable definition

United States Space Force doctrine defines space superiority as the degree of control that permits operations at a chosen time and place without prohibitive interference while denying the same to an adversary [24]. A 2025 United States policy statement likewise presents space superiority as linked to national economic and security interests [28]. These sources make the strategic hypothesis legitimate to examine. They do not prove that ARGUS would deliver it.

The hypothesis can be decomposed:

Necessary condition for broad superiority	ARGUS contribution	Current assessment
persistent awareness	sensor, tracking, metrology and data-fusion spillovers	credible programme benefit
assured access and logistics	launch demand, modular assembly, servicing and repair	potentially major, dependent on ISAM
resilient operations	distributed modules, repair and alternative modes	conceptual; architecture unvalidated
protected command	high-assurance control and independent telemetry	designable; must be demonstrated
freedom from prohibitive interference	one large structure is observable and potentially vulnerable	not established
ability to deny adversary interference	outside the civil programme and not established by these papers	no evidentiary support
alliance and legal legitimacy	transparent planetary-defence service and shared governance	achievable if designed from inception
industrial depth	pulse power, materials, robotics, energy and launch supply chains	strongest near-term strategic pathway

The strict claim—“ARGUS gives one state military supremacy in space within eight years”—is not supported. The stronger evidence-based claim is:

A sponsor that leads the enabling industrial base, assembly logistics, sensing, standards and governed planetary-defence service could gain durable strategic space advantage and significant agenda-setting power.

9.3 Strategic outcomes ladder

Level	Outcome	Earliest defensible horizon
0	reproducible science and falsification	present publication set
1	national laboratories, skills and supply-chain learning	years 1–4
2	integrated civil-space demonstrator and independent control regime	years 5–8
3	coalition-standard leadership in large-space infrastructure and planetary defence	years 6–12
4	first supervised civil capability with strategic option value	conditional years 12–16
5	broad space superiority	cannot be attributed to ARGUS alone

The programme can contribute to a wider space-superiority posture through industrial capacity, resilience and infrastructure control. It cannot replace launch fleets, communications, navigation, intelligence, alliance access, cyber defence, manoeuvre, distributed architectures or political legitimacy. The ladder exists to stop level 5 being claimed on the strength of level 1.

9.4 Why overt weapon framing destroys value

An openly counterspace ARGUS proposition would have five predictable effects:

1. it would invite pre-emption against a large, observable and infrastructure-dependent object;
2. it would increase incentives for competitors to build asymmetric countermeasures rather than match the system;
3. it would complicate launch licensing, insurance, financing and international partnership;
4. it would undermine the civil evidence-sharing needed for planetary defence; and
5. it would weaken the rule-of-law claim by concentrating irreversible power without shared verification.

The United Nations General Assembly's 2022 resolution on destructive direct-ascent anti-satellite testing reflects international concern over debris-generating counterspace behaviour [32]. ARGUS is not a direct-ascent missile, but the political lesson transfers: strategic legitimacy depends on observable restraint and responsible behaviour.

The most powerful strategic configuration is therefore not unilateral secrecy. It is sovereign technical leadership embedded in a coalition governance structure: the sponsor owns the difficult industrial capability while rules, threat validation and mission authorisation make misuse politically and technically harder.

10. Deployment options

Option	Eight-year plausibility	Strategic value	Principal defect	Decision
ground national research facility	high if funded and gated	pulsed power, materials, skills, deterrent knowledge	not an orbital capability	baseline
non-energetic LEO assembly/control pathfinder	moderate	ISAM, metrology, repair, control and launch cadence	does not prove energetic system	baseline
reduced-energy civil flight-science article	moderate after G3	environment and materials evidence	strict safety and mission limits	conditional
full seven-kilometre LEO energetic system	low within eight years	high option value if it closes	mass, power, vulnerability, dust, control and authorisation	no commitment before G4
cislunar or heliocentric system	not schedulable	less atmosphere/debris exposure for some missions	logistics, communications, thermal and control complexity; geomagnetic method inapplicable	post-G4 study

The geomagnetic companion paper is LEO-specific [6]. It must not be used to support cislunar or heliocentric control. Likewise, the absence of a complete atmosphere beyond low orbit does not mean absence of matter: dust and plasma remain mission variables.

11. Scientific and commercial spillovers

The programme has investment value before full system closure.

- **Pulsed power:** switching, insulation, transmission, diagnostics, fault isolation and high-energy-density research.
- **Energy systems:** orchestration of grid storage, intermediate stores, thermal recovery and high-consequence power controls.
- **Materials:** diamond composites, ceramics, refractory layers, high-strain-rate interfaces and impact diagnostics.
- **Space construction:** autonomous assembly, structural metrology, reconfiguration, servicing and inspection.
- **Space-domain awareness:** object characterisation, uncertainty fusion, conjunction analysis and independent telemetry.
- **Planetary science:** impact physics, equations of state, ejecta characterisation and mitigation modelling.

- **Governance technology:** auditable authorisation, configuration attestation and cross-border mission evidence.

Commercial partners should receive intellectual-property and market rights for ordinary civil applications while the sponsor retains safety, mission-authorisation and controlled-technology rights. A public-private special-purpose programme can coordinate those interests, but it should not transfer treaty responsibility to a corporation.

12. Risk register and falsification

Risk	Leading indicator	Gate response
moving article fails before representative speed	damage and release distributions diverge from models	down-select materials or terminate speed branch
cell/string scaling loses force or efficiency	measured scaling exponent misses preregistered band	re-architect; no length procurement
pulse tier mass becomes prohibitive	energy and power specific performance fail independent lower bounds	reduce energy mode or stop
dust survival cannot close	calibrated impact model predicts unacceptable mission loss	change geometry/material/trajectory or reject mode
structure cannot survive transient loads	strain, vibration or alignment exceed recoverable bounds	shorten, segment or terminate architecture
counter-impulse is operationally excessive	consumables, plume, contamination or timing exceed limit	alternative momentum management or reject mode
ISAM cannot achieve required precision	assembly/metrology error exceeds tolerance	ground-only facility or redesigned modularity
cost confidence degrades	P80 exceeds authorised ceiling	pause and re-scope
legal authority or legitimacy fails	no accountable authorisation and consultation regime	no orbital energetic deployment
strategic posture provokes balancing	partners decline and threat perception rises	strengthen civil coalition structure and transparency

The conditional probability of reaching deployment can be written without inventing numerical confidence:

$$Pr(G_6) = Pr(G_0) \prod_{i=1}^6 Pr(G_i | G_0, \dots, G_{i-1}). \quad (9)$$

The present papers do not supply calibrated values for those conditional probabilities. Publishing a single “chance of success” would therefore be false precision. The purpose of the gates is to learn them.

13. Governance and rule of law

This paper stipulates an underlying strategic concern: that increasing occupation and commercial use of space will reshape global power. Infrastructure owners influence access, standards, insurance, data, dispute resolution and the practical meaning of freedom of operation. The rule of law will not be preserved simply because a leading state has the strongest machine. It will be preserved if technical power is coupled to institutions that make decisions reviewable, attributable and constrained.

For ARGUS, that means:

- public law defining the competent authority;
- separation between operator, mission proposer, authoriser and investigator;
- registration and continuing supervision of non-governmental partners;
- independent scientific validation of a planetary-defence threat;
- documented necessity, proportionality and alternatives analysis;
- consultation under the Outer Space Treaty's due-regard framework;
- liability, compensation and insurance arrangements;
- supply-chain, export-control and cybersecurity governance;
- external observation and post-event publication; and
- decommissioning, disposal and environmental responsibility.

The legal companion article should be read alongside this programme paper [5]. Nothing in a DOI deposit constitutes mission authorisation or legal advice.

14. Investment case

14.1 Why a government should fund formulation

A government is uniquely positioned to value low-probability, high-consequence planetary defence; to build national laboratories; to hold international responsibility; and to sustain a multi-cycle programme. The first tranche is justified if it produces:

- a credible architecture and negative-result capability;
- a national map of relevant suppliers and facilities;
- measured performance of the highest-risk interfaces;
- an independent cost and schedule basis;
- a civil governance model; and
- a portfolio of commercially transferable technologies.

The additional strategic justification is avoiding dependence. Heavy lift, large-structure assembly, high-consequence energy systems and space-domain awareness are likely to shape future access to resources, communications and security. Major powers and alliances are already funding that competition [24,28,33–37]. A formulation tranche buys a sovereign option before a rival or commercial gatekeeper sets the architecture, standards and price of entry.

14.2 Why a major launch-and-energy company should participate

A vertically integrated launch, energy, software and manufacturing company would gain a demanding anchor use case for high-cadence heavy lift, large-scale storage, automated fabrication, robotics, power electronics, digital engineering and in-space logistics. Starship-class payload capability is relevant [15], and Tesla-class storage is relevant as a buffer [9], but neither should be presented as proof that the reference architecture closes.

The commercially sensible entry is a scoped consortium role:

1. launch and orbital logistics architecture;
2. non-energetic structural pathfinder;
3. buffer-energy and supervisory-control demonstration;
4. rapid manufacturing and supply-chain design; and
5. options for later phase participation conditional on evidence.

This lets industry create valuable infrastructure without underwriting an unbounded physics claim.

14.3 The disciplined optimistic conclusion

The concept is not ready for a production commitment. It is ready for a serious national formulation programme. The difference is not rhetorical. Formulation has bounded cost, named outputs, independent review and stop rules.

A well-run programme could reach a decisive result within eight years. If positive, years 9–16 could produce a first supervised civil capability and a strategically important industrial ecosystem. If negative, the sponsor would have bought high-value science and infrastructure while avoiding a much larger failed deployment.

15. Conclusions

The ARGUS-PD/POD-7K research now supports a realistic programme proposition:

1. **Eight years is the proof horizon.** It can plausibly deliver TRL 5–6 evidence for selected subsystems, an integrated reduced-scale ground demonstrator and a non-energetic orbital assembly/control pathfinder.

2. **Twelve to sixteen years is the conditional civil-capability horizon.** It is not guaranteed, and it begins only after the year-eight mass, pulse-power, survivability, assembly, cost and governance gates close.
3. **Commercial batteries help but do not solve the pulse.** The energy count is manageable in tens of Megapacks; the rate gap is approximately 8.57×10^5 . A dedicated hierarchical pulse architecture remains the central engineering problem.
4. **The direct mode is \$18.1 billion and the early P80 proof-programme screen is \$27.1 billion.** The former reconciles my own bottom-up arithmetic; the latter retains integration and uncertainty. Neither is suitable for appropriation as an acquisition baseline.
5. **Strategic advantage is credible; unilateral supremacy is not yet evidenced.** The strongest pathway is leadership in industry, launch, ISAM, sensing, logistics, standards and a governed planetary-defence service.
6. **Control and legitimacy are part of the machine.** Independent authorisation, dual control, telemetry, audit, national supervision and international consultation must be designed from the start.

The correct optimistic finding is not that ARGUS will certainly dominate space. It is that a state or coalition willing to fund and govern the programme could acquire a rare combination of scientific knowledge, industrial capacity, strategic option value and institution-building power. That is enough to justify a formulation tranche. It is not enough to justify anything beyond G0, and this paper does not ask for more.

Data and reproducibility

The accompanying reproducibility package contains `argus_pd_repro/programme_model.py`, `argus_pd_repro/run_programme_analysis.py`, the nine cost-input rows, the 250,000-sample deterministic cost output, annual funding profile, public comparator arithmetic, stage-gate table, figures and unit tests. The model contains no target solutions, firing logic, mission trajectories or live command architecture.

References

1. Farmer, C. I. V., *ARGUS-PD/POD-7K Paper I: Parametric Design-Space Analysis of a Modular Projectile Family and Electromagnetic Launch Interface*, version 1.1.0, 2026, DOI pending.
2. Farmer, C. I. V., *ARGUS-PD/POD-7K Paper II: Integrated Architecture and Feasibility Bounds for a Fixed Seven-Kilometre Electromagnetic Planetary-Defence Research System*, version 1.1.0, 2026, DOI pending.
3. Farmer, C. I. V., *ARGUS-PD/POD-7K Paper III: Coupled Feasibility Boundaries, Uncertainty Propagation and Falsification Criteria*, version 1.1.0, 2026, DOI pending.
4. Farmer, C. I. V., *ARGUS-PD/POD-7K Launcher Science and Application Envelope: Distributed Pulse-Module Architecture, Counter-Impulse Closure, Projectile-Family Physics and Space-Only Research Utility*, version 1.0.0, 2026, DOI pending.
5. Farmer, C. I. V., *Conventional Counterspace Conduct and the Outer Space Treaty: Article VI Supervision, Responsible Behaviour and Nationally Anchored Market Conditionality*, version 1.1.0, 2026, DOI pending.
6. Farmer, C. I. V., *Feasibility Bounds for Geomagnetic Attitude Control of Very Large Flexible Space Structures*, version 1.1.0, 2026, DOI pending.
7. Sandia National Laboratories, “Z Machine—An Engine of Discovery,” reporting up to 22 MJ stored energy and approximately 80 TW delivered power, accessed 29 July 2026, <https://www.sandia.gov/news/publications/hpc-annual-reports/article/z-machine/>.
8. Sandia National Laboratories, “How Does the Z Machine Work?,” accessed 29 July 2026, <https://www.sandia.gov/z-machine/about-z/how-z-works/>.
9. Tesla, “Order Megapack,” 4-hour product specification: 3,916 kWh and 979 kW per unit, accessed 29 July 2026, <https://www.tesla.com/megapack/design>.
10. NASA, *NASA Cost Estimating Handbook*, fourth edition and appendices, accessed 29 July 2026, <https://www.nasa.gov/ocfo/ppc-corner/nasa-cost-estimating-handbook-ceh/>.
11. NASA, *NASA Systems Engineering Handbook*, NASA/SP-2016-6105 Rev. 2, 2016, https://www.nasa.gov/wp-content/uploads/2018/09/nasa_systems_engineering_handbook_0.pdf.
12. NASA, “NASA Program/Project Life Cycle,” including Key Decision Points, accessed 29 July 2026, <https://www.nasa.gov/reference/3-0-nasa-program-project-life-cycle/>.

13. U.S. Government Accountability Office, *Cost Estimating and Assessment Guide: Best Practices for Developing and Managing Program Costs*, GAO-20-195G, 2020, <https://www.gao.gov/assets/gao-20-195g.pdf>.
14. NASA Earth Science Technology Office, “Technology Readiness Levels,” accessed 29 July 2026, <https://esto.nasa.gov/trl/>.
15. SpaceX, “Starship,” design objective of more than 100 metric tonnes to orbit in a fully reusable configuration, accessed 29 July 2026, <https://www.spacex.com/vehicles/starship>.
16. NASA, “International Space Station Facts and Figures,” accessed 29 July 2026, <https://www.nasa.gov/international-space-station/space-station-facts-and-figures/>.
17. Mulvaney, J. et al., *In-Space Servicing, Assembly, and Manufacturing State of Play—2025 Edition*, NASA, 2025, https://ntrs.nasa.gov/api/citations/20250008988/downloads/NASA_ISAM_State_of_Play_2025_Edition.pdf.
18. Hoyt, R. P., *SpiderFab: Process for On-Orbit Construction of Kilometer-Scale Apertures*, NASA NIAC final report NNX12AR13G, 2013, https://www.nasa.gov/wp-content/uploads/2019/03/niac_2012_phasei_spiderfab_hoyt_tagged.pdf.
19. NASA, “On-orbit Servicing, Assembly, and Manufacturing 1 (OSAM-1),” cancellation status, accessed 29 July 2026, <https://www.nasa.gov/mission/on-orbit-servicing-assembly-and-manufacturing-1/>.
20. NASA Office of Inspector General, *NASA’s Implementation and Management of Its Planetary Defense Strategy*, IG-25-006, 2025, <https://oig.nasa.gov/wp-content/uploads/2025/06/final-report-ig-25-006-nasas-implementation-and-management-of-its-planetary-defense-strategy.pdf>.
21. NASA Office of Inspector General, *NASA’s Management of Its Johns Hopkins University Applied Physics Laboratory Portfolio*, IG-22-017, 2022, DART projected life-cycle cost, <https://oig.nasa.gov/docs/IG-22-017.pdf>.
22. Adams, E. et al., *Double Asteroid Redirection Test Final Technical Report*, NASA NTRS 20230015804, 2023, <https://ntrs.nasa.gov/api/citations/20230015804/downloads/DART%20Final%20Technical%20Report.pdf>.
23. National Institute of Standards and Technology, *Enhanced Security Requirements for Protecting Controlled Unclassified Information: A Supplement to NIST SP 800-171*, NIST SP 800-172, dual-authorisation control 3.1.1e, <https://doi.org/10.6028/NIST.SP.800-172>.

24. United States Space Force, *Space Force Doctrine Document 1*, 2025, definition of space superiority, https://www.starcom.spaceforce.mil/Portals/2/Space%20Force%20Doctrine%20Document%201%20FINAL_4Apr25.pdf.
25. United Nations Office for Outer Space Affairs, "Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, including the Moon and Other Celestial Bodies," Articles VI and IX, <https://www.unoosa.org/oosa/en/ourwork/spacelaw/treaties/outerspacetreaty.html>.
26. NASA, *Planetary Defense Strategy and Action Plan*, 2023, https://www.nasa.gov/wp-content/uploads/2023/06/nasa_-planetary_defense_strategy-_final-508.pdf.
27. NASA, "NASA Celebrates First Decade of International Asteroid Warning Network," 2024, <https://science.nasa.gov/directorates/smd/nasa-celebrates-first-decade-of-international-asteroid-warning-network/>.
28. The White House, "Ensuring American Space Superiority," presidential action, 18 December 2025, <https://www.whitehouse.gov/presidential-actions/2025/12/ensuring-american-space-superiority/>.
29. NASA Office of Inspector General, *NASA's Partnerships with International Space Agencies for the Artemis Campaign*, IG-23-004, 2023, <https://oig.nasa.gov/docs/IG-23-004.pdf>.
30. Sandia National Laboratories, "Inertial Confinement Fusion—Pulsed-Power Science," discussion of Z and possible next-step stored-energy facilities, accessed 29 July 2026, <https://www.sandia.gov/pulsed-power/inertial-confinement/>.
31. NASA, *Joint Cost and Schedule Confidence Level Analysis*, NASA Cost Estimating Handbook Appendix J, 2015, https://www.nasa.gov/wp-content/uploads/2020/11/ceh_appj.pdf.
32. United Nations General Assembly, Resolution 77/41, *Destructive Direct-Ascent Anti-Satellite Missile Testing*, adopted 7 December 2022, <https://docs.un.org/A/RES/77/41>.
33. State Council Information Office of the People's Republic of China, *China's Space Program: A 2021 Perspective*, "A New Journey Towards a Strong Space Presence," 2022, https://english.www.gov.cn/archive/whitepaper/202201/28/content_WS61f35b3dc6d09c94e48a467a.html.
34. European Commission, "EU Space Strategy for Security and Defence," including resilience, threat response, dual-use capability and the EU's role as a global space power, accessed 29 July 2026, https://defence-industry-space.ec.europa.eu/eu-space/eu-space-strategy-security-and-defence_en.
35. United Kingdom Ministry of Defence, *Defence Space Strategy: Operationalising the Space Domain*, 2022, <https://www.gov.uk/government/publications/defence-space-strategy-operationalising-the-space-domain/defence-space-strategy-operationalising-the-space-domain>.

36. North Atlantic Treaty Organization, “NATO’s Approach to Space,” updated 8 July 2026, <https://www.nato.int/en/what-we-do/deterrence-and-defence/natos-approach-to-space>.
37. Government of India, *Indian Space Policy—2023*, 2023, https://www.isro.gov.in/media_isro/pdf/IndianSpacePolicy2023.pdf.
38. NASA Office of Inspector General, *NASA’s Efforts to Demonstrate Robotic Servicing of On-Orbit Satellites*, IG-24-002, 2023, <https://oig.nasa.gov/wp-content/uploads/2023/10/ig-24-002.pdf>.
39. U.S. Government Accountability Office, *Nuclear Weapons: Actions Needed to Address Scientific and Technical Challenges and Management Weaknesses at the National Ignition Facility*, GAO-10-488, 2010, <https://www.gao.gov/products/gao-10-488>.
40. U.S. Bureau of Labor Statistics, *Historical Consumer Price Index for All Urban Consumers (CPI-U): U.S. City Average, All Items, December 2025*, <https://www.bls.gov/cpi/tables/supplemental-files/historical-cpi-u-202512.xlsx>; and *Consumer Price Index—June 2026*, <https://www.bls.gov/news.release/pdf/cpi.pdf>.