

THE UNIVERSE BUTTERFLY (UB) PROGRAM

Aggregated-Asteroid Rotating Habitats:
A Scientific Feasibility Assessment and
Fifty-Year Implementation Roadmap

WHITE PAPER · VERSION 1.0

Kailong Cai

Universe Butterfly Program · universebutterfly.com

July 2026

DOI: 10.5281/zenodo.21190731

Released under Creative Commons Attribution 4.0 (CC BY 4.0). Copy, adapt, build, and commercialize freely — with attribution to the Universe Butterfly (UB) Program, K. Cai, 2025-.

Abstract

The Universe Butterfly (UB) Program proposes a staged, fifty-year pathway to the first kilometer-scale rotating space habitat built almost entirely from asteroid material. The architecture rests on four physical constraints that force its design uniquely: artificial gravity must come from rotation, not mass; aggregation of bodies is docking physics, not collision physics; decades of orbital lead time substitute for orders of magnitude of propulsive energy; and all operations must remain in verifiably non-Earth-threatening staging orbits. The Program's capability stack — target characterization, long-horizon trajectory design, micro-thrust orbit modification, non-cooperative rendezvous and capture, and precision controlled delivery — is monetized in four successive markets of increasing heliocentric distance: orbital-debris capture and recycling; asteroid-resource wholesaling in cislunar space; impact-assisted volatile delivery in support of lunar and Mars development; and a pre-positioned deep-space rapid-response reserve for planetary defense against the "large-and-fast" threat class. No element of the architecture violates known physics; every element has a published proof of principle or a flight mission in progress as of 2026. The binding constraints are scale, capital, calendar time, and organization. The Program's governing doctrine, the Time-Asymmetry Principle, holds that motivation is an instantaneous variable while capability is an integral one: planetary-scale capability must be built before the motivation for it arrives, because it cannot be built after. This paper presents the physical foundations, answers the three strongest anticipated objections quantitatively, surveys the state of the art, details the habitat construction architectures and the four-engine economic structure, and proposes a five-phase roadmap to full operation of UB Station One in 2076, with explicit decision gates, an honest treatment of schedule probability, and a governance framework engineered as verifiable physics rather than promises.

1. Introduction and Program Definition

1.1 The proposition

Humanity's multiplanetary programs concentrate on two destinations — the Moon and Mars — both of which are planetary surfaces at the bottom of gravity wells, both of which offer partial gravity of unknown biological adequacy, and both of which remain permanently inside the same star system and therefore inside the same envelope of systemic risk as Earth. A third settlement architecture has existed in the literature for a century: the free-flying rotating habitat, from Bernal's 1929 hollowed-sphere colonies [1] through O'Neill's 1970s cylinders [2] to modern engineering treatments [3-5]. Its historical objection has always been mass: a habitat providing Earth-normal gravity and Earth-equivalent radiation shielding requires on the order of 10^8 - 10^{11} tonnes of material, unlaunchable from Earth at any foreseeable cost.

The UB Program's proposition is that this objection dissolved quietly over the last decade. The mass is not on Earth; it is in the near-Earth asteroid population — more than 38,000 cataloged objects, a number about to grow severalfold under new surveys [6,7] — and the technologies to characterize, move, join, and reshape that mass have each been demonstrated in principle

by planetary-defense missions, sample-return missions, commercial satellite-servicing operations, and published structural analyses. What does not yet exist is the integrated program that assembles these capabilities into a construction industry. UB is a design for that program.

Program definition. UB will: (i) select volatile- and metal-bearing near-Earth asteroids in the 10–500 m class from the post-2027 survey catalogs; (ii) modify their orbits over multi-year horizons using gravity-assist cascades and continuous micro-thrust, delivering them to long-term-stable staging orbits (Sun–Earth L4/L5; lunar distant retrograde orbit, DRO); (iii) aggregate them by controlled contact at centimeter-per-second relative velocities with mechanical anchoring; (iv) construct, by one of three candidate architectures (Section 5), a rotating habitat of approximately kilometer scale providing 1 g interior gravity, 5 t/m² of native-regolith radiation shielding, and a closed-loop ecology for a permanent population of order 1,000; and (v) finance the fifty-year effort through four sequential markets (Section 6), each of which is independently profitable and each of which flight-qualifies the capabilities of the next. Target date for full operation of UB Station One: 2076.

1.2 What this paper is and is not

This is a feasibility assessment and roadmap, not a mission design. All quantitative statements are order-of-magnitude analyses using standard parameters, with derivations collected in Appendix A; their purpose is to locate the boundary between the physically excluded, the merely difficult, and the near-term practical. Mission-grade design is identified throughout as Phase 0 and Phase 1 output. The paper is written to be attacked: Section 3 states the three strongest objections to the Program in their strongest forms and answers them with numbers, because a program of this ambition earns credibility only by hosting its own opposition.

1.3 On the name: sensitive dependence as an engineering lever

The Program is named for the butterfly effect — Lorenz's demonstration that in nonlinear dynamical systems, arbitrarily small differences in initial conditions grow exponentially into system-scale divergences [43]. The name is not ornament; it is a statement of the Program's core mechanism. The gravitational N-body problem is exactly such a system: near planetary close encounters, neighboring trajectories diverge exponentially — which is why the long-term prediction of certain asteroid orbits is hard, and why the *steering* of those same orbits is, counterintuitively, cheap. A millimeter-per-second impulse applied years before a gravitational-keyhole passage is amplified by planetary gravity into kilometer-per-second-scale changes in the subsequent orbit (Section 2.3). Chaos, the forecaster's adversary, is UB's amplifier of action.

This mechanism answers a question larger than habitat construction: *by what path does a civilization whose power is measured in gigawatts obtain outcomes measured in planets?* On the Kardashev scale [44], commanding planetary-scale masses and energies is the mark of a Type I civilization — a status centuries away by any extrapolation of raw energy growth. Sensitive dependence collapses that gap for a specific and precious class of problems: wherever a system is chaotic and foresight is long, small forces applied at the sensitive point purchase planetary-scale results. Orbital mechanics is the largest such system within human reach. The butterfly is therefore both the Program's method — decades of foresight, grams of force, planets of consequence — and its thesis about how a young species grows: not by

waiting to become large, but by learning where to be small.

2. Physical Foundations: Four Constraints That Force the Design

The UB architecture is not one choice among many. It is the unique solution compatible with four physical facts. A reader who internalizes these four constraints can reconstruct the entire Program from first principles.

2.1 Constraint I — Gravity must come from rotation, not mass

For a body to retain a breathable atmosphere over geological time, its escape velocity must exceed the thermal velocity of atmospheric molecules by roughly a factor of six; Earth (11.2 km/s) qualifies, Mars (5.0 km/s) marginally, the Moon (2.4 km/s) does not. The total mass of the main asteroid belt is $\approx 2.4 \times 10^{21}$ kg — about 3% of the Moon's mass and 0.04% of Earth's, with Ceres alone accounting for roughly a third [8]. If every belt object were miraculously fused into a single sphere of density 3,000 kg/m³, the result would have radius ≈ 576 km, surface gravity ≈ 0.05 g, and escape velocity ≈ 0.75 km/s (Appendix A.1) — a body that can hold neither an atmosphere nor human skeletal health. This is not an engineering limit; the solar system inside 3 AU simply does not stock the raw material for a naturally habitable world. Jupiter saw to that 4.5 billion years ago.

The consequence is definitional for UB: **habitability is manufactured by rotation and enclosure, not accreted by mass.** Interior gravity follows $\omega^2 R = g$. At the classical vestibular comfort limit of 2 rpm, 1 g requires $R \geq 224$ m — fixing the "kilometer-class" scale of a full station — while more recent human-factors work suggests adaptation up to ~ 4 rpm may be acceptable [9], which would admit 1 g at $R \approx 56$ m and materially de-risks the 100-m-class prototypes of Phase 2. Atmosphere is retained by a pressure hull, not by gravity; radiation protection is provided by meters of regolith, not by a magnetosphere. Once these substitutions are accepted, the required mass drops from the 10^{23} kg of a "real planet" to the 10^8 - 10^{11} kg of a large engineered structure — from physically excluded to logistically hard.

2.2 Constraint II — Aggregation is docking, not collision

Whether two bodies merge or shatter on contact is governed by the ratio of their relative velocity to their mutual escape velocity. For a homogeneous body of density ρ and radius r , $v_{\text{esc}} = r \cdot \sqrt{(8\pi G \rho / 3)}$; at $\rho = 2,500$ kg/m³ this gives $v_{\text{esc}} \approx 1.2 \times 10^{-3} \text{ s}^{-1} \times r$ — about **0.6 m/s for a 1-km body and 0.06 m/s for a 100-m body** (Appendix A.2). Natural collisions in the main belt occur at ~ 5 km/s, four orders of magnitude above the merger window; the specific kinetic energy of such impacts ($\sim 10^7$ J/kg for comparable masses) exceeds the catastrophic-disruption threshold Q^*_D of km-scale bodies ($\sim 10^2$ - 10^3 J/kg [10]) by four orders of magnitude. Impact is a disassembly tool. Nature nonetheless provides existence proofs of the merger regime: comet 67P/Churyumov-Gerasimenko is a contact binary formed by a gentle ancient collision [11], and a large fraction of small asteroids are rubble piles held together by self-gravity and weak cohesion — objects that themselves accreted at low speed.

The engineering consequence: **UB's "aggregation" is rendezvous-and-docking technology scaled up.** Multi-year micro-thrust brings the orbits of two bodies into near-identity; terminal approach occurs at cm/s; contact is followed by mechanical consolidation (tethers, nets, anchors, sintered joints). The core operation — autonomous approach, inspection, and capture of a non-cooperative, possibly tumbling object — has already been performed twice in geostationary orbit by commercial servicing vehicles [12] and is being industrialized by the debris-removal sector (Section 6.1). Nothing in the aggregation step is new physics; all of it is new scale.

2.3 Constraint III — Lead time substitutes for energy

Direct transport of asteroids scales linearly in momentum with mass. Moving a 1-km body ($\sim 1.3 \times 10^{12}$ kg) through the ~ 5 km/s of a belt-to-1-AU transfer requires $\sim 10^{16}$ kg-m/s of momentum and, with electric propulsion at specific impulse 3,000 s, on the order of 10^{11} kg of propellant and $\sim 10^{20}$ J of energy — comparable to a year of global electricity generation (Appendix A.4). Whole-body transport of km-class objects is excluded. Three mechanisms, however, trade **calendar time** against this energy, and together they are the enabling physics of the Program:

1. **Gravitational keyholes.** Many near-Earth asteroids make repeated planetary flybys. A millimeter-per-second perturbation applied at a precisely computed point before a flyby is amplified by the planet's gravity into a large change in the subsequent orbit — the planetary-defense "keyhole" phenomenon [13], run in reverse as a steering tool.
2. **Flyby cascades.** Venus-Earth-Earth gravity-assist sequences, the standard fuel-free workhorse of interplanetary navigation, apply equally to a slowly guided asteroid, circularizing and phasing its orbit over successive encounters at near-zero propellant cost.
3. **Integrated micro-thrust.** Solar-electric tugs, ion-beam shepherding, laser ablation, and even albedo modification of the Yarkovsky thermal force [14] deliver small accelerations that, integrated over 20–30 years, accumulate to the tens-to-hundreds of m/s that well-chosen transfers actually require.

The phrase "computing the orbits decades in advance" in the Program's founding question is therefore not rhetorical decoration; **it is the necessary condition for the energy budget to close.** It also fixes the Program's first core asset: not rockets, but a long-horizon trajectory-planning and target-characterization capability. This constraint is, further, the mechanical root of the Time-Asymmetry Principle developed in Section 3.3.

2.4 Constraint IV — Staging at L4/L5 and DRO, never in Earth-crossing orbits

An uncontrolled reentry of a 100-m body carries the energy of hundreds of Hiroshima-scale detonations, and any capability to move asteroids is intrinsically dual-use — the "deflection dilemma" articulated by Sagan and Ostro [15]. UB therefore adopts as non-negotiable design law: all aggregation and construction occurs at Sun-Earth L4/L5 (provably long-term stable — Earth's natural Trojan asteroids 2010 TK7 and 2020 XL5 reside there [16] — at 1 AU from Earth with ~ 8 -minute one-way light time, on orbits that never intersect Earth's) or, for early small-scale work, in lunar distant retrograde orbit (DRO, $\sim 3.8 \times 10^5$ km, the staging point selected by NASA's Asteroid Redirect Mission studies [17]). Every UB asset obeys a

two-maneuver separation rule: its trajectory is designed such that placing it on any Earth-threatening course would require at least two large, months-long-observable maneuvers, independently verifiable by any nation's tracking assets. Under the Outer Space Treaty (Art. VI-VII) and the 1972 Liability Convention, launching states bear absolute liability for surface damage [18]; UB's safety case is designed to be a property of orbital mechanics that any third party can check, not a promise that any party must trust. Section 9 develops the full governance architecture.

3. Anticipated Objections: The Three Hard Questions

A proposal of this scope earns a hearing only by stating its strongest objections in their strongest forms — steel men, not straw men — and answering them quantitatively.

3.1 "Why not launch from Earth? Fully reusable heavy lift will make launch cheap."

The objection, at full strength. Starship-class reuse is projected to reduce launch costs toward \$100-200/kg. Earth has unlimited industrial capacity and every material known. The International Space Station was built by launch and assembly; scale that up rather than herding rocks across the solar system.

Answer, in three layers.

Layer 1 — the mass ledger. The mass floor of a 1-g habitat is set by radiation shielding: reduction of galactic cosmic radiation to terrestrial-equivalent dose requires roughly 5 t/m² of areal density — about 2.5 m of regolith at 2,000 kg/m³ [19]. Over the ≈56 km² of interior area computed by Miklavčič et al. for a spin-deployed rubble-pile habitat [4], the shielding alone is ≈2.8×10¹¹ kg — **280 million tonnes** (Appendix A.3). At \$150/kg — an aggressive projection — launch fees alone are ≈\$42 trillion, requiring on the order of two million heavy-lift flights, before a single structural member, tank, or tonne of water is counted. A single C/S-type asteroid of 500-650 m diameter is that mass, already in orbit. For missions denominated in hundreds of megatonnes, Earth launch fails on arithmetic, independent of how inexpensive rockets become.

Layer 2 — the energy ledger. Lifting 1 kg from Earth's surface to escape requires a minimum of 62.5 MJ (GM_E/R_E); chemical rockets, after gravity losses and staging, expend roughly an order of magnitude more than this in propellant energy, and the rocket equation's exponential penalty holds payload fractions to single-digit percentages. Moving 1 kg already in heliocentric orbit through the tens-to-hundreds of m/s of a well-chosen transfer, with electric propulsion at I_{sp} = 3,000 s, costs two to three orders of magnitude less energy per kilogram (Appendix A.4). This is not an engineering gap that closes with progress; it is the depth of the gravity well. Earth is the most expensive quarry in the solar system for the sole reason that we live at the bottom of it.

Layer 3 — the crossover and the division of labor. Earth launch is not always the loser. For ISS-class missions (hundreds of tonnes), it wins outright. The cost crossover falls near 10⁴-10⁵ tonnes of total delivered mass: below it, launch is cheaper; above it, a reusable electric-tug

fleet of order \$3-5B, amortized across many retrievals, plus in-situ material, dominates decisively. UB therefore adopts a division-of-labor law: **launch carats, mine tonnes**. Chips, precision machinery, biological seed stock, and people — complexity valued at up to millions of dollars per kilogram — fly on rockets. Structure, shielding, water, and propellant — mass valued at dollars per kilogram — come from the asteroid. Heavy-lift reuse is not refuted by this analysis; it is assigned its correct role as the trunk line for complexity and crews, the very capability that makes UB's equipment cheap to deploy. Launch makes *getting up* cheap. Only asteroids make *having mass up there* cheap.

3.2 "O'Neill's original plan used lunar material and a mass driver. Why bypass the Moon?"

The objection, at full strength. The classical 1970s designs [2] baseline a lunar electromagnetic mass driver flinging regolith to a catcher at L5. The Moon is three days away, is the object of funded national programs, and its material properties are far better characterized than any asteroid's.

Answer. The lunar route is UB's sibling, not its rival — but four structural facts establish the independent necessity of the asteroid route. **(i) The Moon is still a gravity well.** Every kilogram exported pays a 2.4 km/s tax, collectible only after a gigawatt-class power and industrial base exists on the surface — a bootstrap problem the mass driver presupposes rather than solves. An asteroid's well is measured in centimeters per second, and the *entire body* can be moved. **(ii) The Moon is volatile-poor.** Lunar regolith is severely depleted in carbon, nitrogen, and hydrogen — precisely the closed-ecology shopping list — with polar ice as a localized, contested exception. C-type asteroids carry that list natively: the Ryugu and Bennu samples returned hydrated phyllosilicates, carbonates, phosphates, evaporite salts, amino acids, and nucleobases [20,21]. O'Neill himself conceded that colony carbon and nitrogen would need asteroidal supply. **(iii) Δv economics.** From low Earth orbit, the cheapest near-Earth asteroids are reachable for ≈ 4 km/s — *less* than the ≈ 6 km/s of a soft lunar landing; by the energy metric, hundreds of asteroids are closer than the Moon's surface, a counterintuitive fact institutionalized in NASA's NHATS accessibility catalog [22]. **(iv) Tenure.** Lunar territory and polar ice are already the contested board of national base programs; near-Earth asteroids number in the tens of thousands, unclaimed, and are the only celestial objects a new program can acquire whole. The correct relationship is commercial: the Moon is UB's first *customer* — for made-to-order volatile delivery (Section 6.3) — not its quarry.

3.3 "Technology improves every year. Why not wait?" — The Time-Asymmetry Principle

The objection, at full strength. In twenty years AI, robotics, and launch will each be an order of magnitude better. Everything built now will look like a detour. The rational strategy is patience.

Answer — and the Program's first doctrine. The objection is correct for every task that technology accelerates, and exactly backwards for every task that consumes calendar time regardless of technology. Three classes of work belong to the second category, and UB's value is concentrated entirely within them:

1. **Characterization inventory.** Taking a specific asteroid from detection to full knowledge of size, composition, spin state, and internal structure requires years of ground observation plus at least one proximity mission — a per-target clock of years to a decade.
2. **Orbital pre-positioning.** Delivering a body from its natural orbit to a usable station (L4/L5, DRO, resonant orbits) consumes multi-year low-thrust transfers gated by planetary geometry; a missed Venus window is not negotiable with better software.
3. **Flight-verified confidence.** The engineering reliability of contact, anchoring, and controlled delivery is fed only by real missions, accumulated in real years.

These clocks share one property: **they measure the passage of the calendar, not the state of the art.** No conceivable AI advances the date of a gravity assist. Motivation, by contrast, is an instantaneous variable: a single comet warning, a single impact event, any systemic shock can align global will overnight — as 1939 aligned it for radar and fission, as 2020 aligned it for vaccines. Humanity survived those episodes because the crises granted reaction windows measured in years. Celestial mechanics issues no such guarantee (Section 6.4 documents a threat class with warning times measured in months). The correct temporal strategy is therefore not "wait for technology" but **"let technology improve while starting, today, every task whose duration is set by the calendar — since those tasks run no faster for having been started late."**

Stated as doctrine: **motivation can appear overnight; capability cannot; therefore capability must exist before motivation does.** Insurance cannot be bought after the fire has started. Every commercial engine in Section 6 is a monetized application of this principle: clearing orbits before the cascade, staging ore before the demand, delivering volatiles before the settlement wave, and banking momentum before the comet.

4. State of the Art, 2026: The First Half of the Road Is Being Paved

The Program does not begin from zero. The decade 2022–2031 is delivering, at global public and private expense, flight validation of every precursor capability UB requires.

Orbit modification, demonstrated and being calibrated. NASA's DART impactor (579 kg at 6.14 km/s) shortened the orbital period of the 150-m moonlet Dimorphos by ~33 minutes in September 2022 — humanity's first measured alteration of a celestial trajectory — with a momentum-enhancement factor $\beta \approx 3.6$ from impact ejecta [23]; the impact also measurably reshaped the rubble-pile target, a cautionary datum for UB's insistence on gentle methods. ESA's Hera spacecraft, launched October 2024, **arrives at Didymos in November 2026** with two CubeSats to measure Dimorphos's mass and interior, converting a one-off demonstration into a predictive engineering model [24]. China's first planetary-defense mission — an observer-plus-impactor pair — is planned for launch in 2027 with a kinetic impact on the ~30-m Aten asteroid 2015 XF261 in April 2029, followed by 6–12 months of on-site effect assessment [25]; notably, the target's orbit ($a \approx 0.99$ AU, near 1:1 Earth resonance) lies in exactly the orbital family UB designates for reserve nodes. In April 2029 the 340-m asteroid Apophis passes Earth at ~31,600 km; ESA's Ramses and NASA's OSIRIS-APEX will observe

planetary tidal forces resculpting a large body in real time — a natural experiment in celestial-body modification [26].

Surface operations and materials knowledge. Laboratory sample inventories now exist for three asteroids (Itokawa, Ryugu, Bennu). Bennu's 121.6 g returned hydrated phyllosilicates, carbonates, phosphates, evaporite salts recording an ancient brine, fourteen protein-forming amino acids, and all five nucleobases [21]; Ryugu's samples yielded comparable organics including uracil [20]. OSIRIS-REx's sampling arm nearly sank into Bennu's unexpectedly fluid-like surface — a discovery that rewrote rubble-pile geotechnics and directly informs every anchoring design [27]. China's Tianwen-2, launched 29 May 2025, completed its arrival maneuvers at the quasi-satellite 2016 HO3 by June 2026 (confirmed by independent Doppler tracking) and will attempt the world's first anchor-and-attach asteroid sampling — the direct ancestor of UB construction operations — returning samples in late 2027 before continuing to main-belt comet 311P [28]. JAXA's Hayabusa2 extension reaches the ~30-m fast rotator 1998 KY26 in 2031, the first field survey of a "transportable-class" object. NASA's Psyche arrives at its metal asteroid in August 2029.

Cataloging. The Vera C. Rubin Observatory began its ten-year survey in 2025 and is expected to multiply the known asteroid census several-fold; NASA's NEO Surveyor infrared telescope (in integration and test as of mid-2026, launch no earlier than September 2027, to Sun-Earth L1) will find two-thirds of near-Earth objects larger than 140 m within five years and, critically for target selection, measures sizes and thermal properties directly [6,7].

Commercial precursors. In asteroid resources: AstroForge's DeepSpace-2 is slated for a Q4 2026 launch toward a near-Earth asteroid rendezvous; Karman+ targets a 2027 hydrated-regolith excavation mission; TransAstra continues NIAC-lineage optical-mining development [29]. In the operations that matter most to aggregation — non-cooperative proximity work — Northrop Grumman's MEV vehicles have twice docked with and extended uncontrolled GEO satellites [12], and Astroscale's ADRAS-J performed the first commercial close inspection of a large derelict rocket stage in 2024, with removal missions and ESA's ClearSpace-1 in the pipeline [30]. Regulatory demand is codified: the U.S. FCC's five-year post-mission disposal rule is in force and has produced its first enforcement action.

Interstellar visitors — the observational basis for Section 6.4. 1I/Oumuamua (2017), 2I/Borisov (2019), and 3I/ATLAS (2025) establish the flux of high-velocity, short-warning objects as measured fact rather than conjecture; Hera passed through 3I/ATLAS's tail in October 2025, a footnote on how close these encounters run. Long-period comet C/2013 A1 (Siding Spring) was discovered only 22 months before its 2014 Mars flyby [31].

Implication. By 2031 humanity will possess: two calibrated kinetic-deflection experiments, four in-situ surveys of distinct small-body types, three laboratory sample suites, two next-generation survey systems, and a commercialized non-cooperative-capture industry. Phase 0 of UB is being funded by the world.

5. Habitat Construction: Architectures, Materials, and Closure

5.1 Three candidate construction architectures

UB carries three architectures to a competitive down-select in Phase 2, on the principle that the choice should be made by subscale flight data, not by white paper.

Architecture A — Spin-bag deployment (baseline; Miklavčič et al., 2022 [4]). A high-strength fiber mesh bag (carbon-nanotube-class tensile members) encloses a rubble-pile asteroid of a few hundred meters. The body is spun up; centrifugal force splays the loose material outward against the bag, which shapes it into a rotating cylinder whose wall — several meters of native rock — is simultaneously the structure's ballast and its radiation shield. The published analysis of a Bennu-class (~490 m) progenitor yields on the order of 56 km² of interior area (comparable to Manhattan) under artificial gravity at a large fraction of 1 g. The architecture's elegance is that it cooperates with, rather than fights, the rubble pile's incohesion, and that essentially none of the structural mass is launched.

Architecture B — Hollowed monolith (Maindl et al., 2019 [5]). Applicable to mechanically competent (non-rubble) rocky bodies: the interior is excavated and the shell spun. Published stress analysis supports ~0.3 g without failure for shells of a few hundred meters' radius; reaching 1 g requires interior tensile liners to carry hoop stress. Candidate-dependent; contingent on finding monolithic targets, which sample-return geotechnics suggest are the minority.

Architecture C — External truss ring. The asteroid serves as counterweight, shield stockpile, and feedstock; an independent engineered ring rotates around or beside it. Structurally the most conservative and materially the least efficient; retained as the low-risk fallback and as the probable form of early crewed prototypes.

A design option common to all three deserves note: the radiation shield need not co-rotate. A static (non-rotating or slowly counter-rotating) shield shell decoupled from the pressure hull — as in the Kalpana One lineage [3] — removes the shield's weight from the rotating structure's hoop-stress budget, reducing required hull strength to that needed for atmospheric pressure and interior payload alone (Appendix A.6 quantifies: a 225-m, 1-g hull carrying atmosphere plus decoupled shielding requires of order 10⁻¹ m of steel-equivalent section — heavy, but manufacturable from asteroidal Fe-Ni by the in-orbit metallurgy line of Section 6.1).

5.2 The materials inventory: the asteroid as a stocked warehouse

Laboratory ground truth from three sample suites establishes the shopping list. C-type bodies supply: water (hydrated phyllosilicates at ~5-10% by mass; a 100-m body at 10% water content holds ~10⁴ tonnes — initial water inventory for a thousand-person ecology plus thousands of tonnes of electrolyzed propellant); carbon and nitrogen (organics, carbonates — the two elements lunar regolith critically lacks); phosphorus (phosphates — the fertility bottleneck of any agriculture); sulfur and salts. S- and M-type bodies supply silicate aggregate and iron-nickel metal for structural fabrication — with platinum-group metals as the M-type's cash crop. UB's material self-sufficiency target: **>99% of structural and shielding mass never launched.** Earth ships machines, chips, seeds, and people.

5.3 Ecological closure: the honest hard problem

Closure — the fraction of water, air, and nutrient mass recycled rather than resupplied — is the least technologically mature element of the entire Program and is treated as such. Current state of the art: ISS environmental control has demonstrated ~98% water-loop recovery [32]; ground analogs (BIOS-3; China's Lunar Palace 365, a 370-day, 4-crew closed test at ~98% system closure [33]; ESA's MELiSSA program) have validated multi-month bioregenerative loops at small scale. What has never been demonstrated is multi-year, multi-generation (crop and human) closure at >95% with population-scale agriculture — and, per Section 8's dual-speed model, **this validation runs in real biological time and is not accelerable by computation**. The roadmap therefore reverse-engineers its crewed milestones from this floor: Phase 2's prototype hosts its first crew by ~2043 precisely so that Decision Gate 2 (~2048) can adjudicate on ≥ 5 years of real closure and human-physiology data.

6. The Four-Engine Architecture: Monetizing One Capability Stack at Four Distances

UB's capability stack is single: target characterization → long-horizon trajectory design → micro-thrust orbit modification → non-cooperative rendezvous, capture, and consolidation → precision controlled delivery. It is sold into four markets of increasing heliocentric distance, each of which is independently profitable, each of which flight-qualifies the next, and each of which is a monetized application of the Time-Asymmetry Principle (Section 3.3).

6.1 Engine 1 — Orbital debris capture and recycling (2020s-2030s): the training ground

Removing a derelict upper stage decomposes operationally into: long-range approach, non-cooperative inspection, de-tumbling, capture, rigid attachment, controlled orbit change. **This is, item for item, the operational decomposition of aggregating an asteroid**, differing only in target mass (tonnes versus kilotonnes) and range (hundreds of kilometers versus tens of millions). Low Earth orbit is the most forgiving classroom — benign lighting, seconds of light delay, abort-to-safe options — and it has paying customers now: tracked objects exceed 36,000, derelict large bodies number in the thousands, Kessler-cascade risk is a sovereign-level concern, and disposal mandates are in force with enforcement precedent. Commercial capture is arriving (Section 4); UB's differentiated entry is **recycling metallurgy**: derelict stages are stockpiles of aerospace-grade aluminum and titanium whose in-orbit replacement value, at launch cost, is thousands of dollars per kilogram, and in-orbit additive manufacturing has years of ISS flight heritage [34]. The strategic outputs of Engine 1 are cash flow, flight hours, and a robotic workforce holding verified licenses in the exact operations Phase 1 requires.

6.2 Engine 2 — Asteroid-resource wholesaling (2030s): the ore parking lot

The wholesaler model: relocate entire 10-m-class volatile- or metal-rich bodies to a cislunar depot (DRO) where extraction companies work at their doorstep — the commercial resurrection of the ARM architecture. The engineering baseline is established: the 2012 Keck

Institute study concluded that a ~40-kW solar-electric vehicle could return a small (~7-m-class, several-hundred-tonne) asteroid to lunar orbit within about a decade for ≈\$2.6B with near-existing technology [17]. Water is the anchor commodity (delivered-to-orbit water from asteroids undercuts Earth-launched water by an order of magnitude at scale), platinum-group metals the cash crop, and lunar programs plus orbital refueling the demand side.

6.3 Engine 3 — Impact-assisted terraforming support, IAT (2040s onward): the volatile logistics company

The scientific opening. The definitive inventory assessment of Mars terraforming (Jakosky & Edwards, *Nature Astronomy*, 2018 [35]) concluded that Mars lacks accessible CO₂ to warm itself by in-situ means — settled science that closes the classical greenhouse route and thereby *opens* the exogenous one. The canonical exogenous scheme is Zubrin & McKay (1993) [36]: steer ammonia- and water-rich small bodies onto Mars, each ~10¹³ kg object contributing of order a few degrees of warming with nitrogen (Mars's other silent deficit) as co-benefit; systematic modification requires tens of such deliveries across decades. UB modernizes the logistics: target selection from the post-Rubin/NEO Surveyor catalogs (near-Earth and main-belt hydrated bodies at far lower Δv than the outer-solar-system objects of the 1993 paper, individually smaller but deliverable at higher cadence); impact siting at polar caps or uninhabited antipodes; sequencing strictly pre-settlement or at exclusion distance from bases; dust-injection anti-greenhouse transients budgeted into delivery cadence; and full deference to COSPAR planetary-protection process pending the outcome of Mars life-detection science. Stated honestly: IAT converts Mars terraforming from *materially impossible to a century-scale logistics problem* — an upgrade, not a shortcut.

The lunar product line makes no atmosphere (2.4 km/s escape velocity retains none) and sells two concrete services. *Made-to-order volatile delivery*: controlled emplacement of 100-m-class hydrated bodies (~10⁵ t water-equivalent) into designated polar cold traps or engineered receiving craters, creating ore bodies whose grade and coordinates are specified by the customer. The controlling physics is retention: the minimum lunar impact speed equals escape velocity (2.4 km/s), and impact flash-vaporization losses are severe — published simulations span retention fractions from a few percent to a few tens of percent depending on velocity, angle, and trap temperature [37] — so **delivery efficiency is governed by exactly the capability UB exists to master: driving terminal relative velocity toward zero** (propulsive terminal braking, "hard-landing delivery," rather than ballistic impact). *Impact earthmoving*: controlled impacts as megascale excavation (shield pits, lava-tube access) — a far-term option contingent on mature debris-safety modeling for lunar orbital assets.

The narrative foundation is mainstream science. The Moon is the product of a giant impact; a substantial fraction of Earth's ocean water arrived on volatile-bearing small bodies (whose composition the Bennu and Ryugu samples now display in the laboratory); Chicxulub cleared the ecological stage for mammals. **Impact is the solar system's native mechanism for distributing volatiles and resetting ecologies. UB domesticates the mechanism — from random catastrophe to scheduled freight.** The one-sentence version for lunar and Mars programs: settlement is bottlenecked on water, nitrogen, atmosphere, and schedule — all four of which are flying around the solar system with return addresses — and UB is the logistics company that delivers them by coordinate, by tonnage, and by timetable.

6.4 Engine 4 — The Deep-Space Rapid-Response Reserve (permanent public good): banking momentum before the comet

The concept, stated precisely. UB's planetary-defense role is not merely "the technology is dual-purpose." It is a specific system proposal: **pre-characterized, propulsion-equipped small bodies of 10^3 - 10^5 tonnes, stationed at three to five computed orbital nodes, held as humanity's standing response to the threat class that launch-on-need cannot reach.**

The threat class is real and observationally established. The current defense architecture — survey, catalog, launch on need — covers threats that are *known and slow*: cataloged near-Earth asteroids whose impact risks are identified decades out, for which DART-class missions suffice. Outside the catalog lies a demonstrated class that is *large and fast*: long-period comets and interstellar objects, arriving from deep space at high velocity with warning measured in months. Calibration by observation: comet Siding Spring was discovered 22 months before its Mars encounter [31]; three interstellar objects in under a decade (2017, 2019, 2025) establish the flux as fact.

Why neither Earth launch nor the nuclear fallback closes. The deflection Δv required to convert an impact into a miss scales as $k \cdot R_{\oplus} / t_w$ (k a geometry factor of order 1–3): ≈ 0.2 - 0.6 m/s at one year's warning, double at six months (Appendix A.5). Two sizing cases:

- *Case A — 300-m rocky body ($\approx 3.5 \times 10^{10}$ kg), 12 months' warning.* Required momentum change: 0.7 - 2×10^{10} kg·m/s — roughly 550–1,600 times the momentum DART actually transferred to Dimorphos, or 2,000–6,000 times DART's own launch momentum. At closing speed 10 km/s and $\beta = 3$, the impactor must mass 230–700 t, delivered to a deep-space intercept: multiple heavy-lift launches plus months-to-years of transfer, on windows dictated by planetary geometry. The calendar rarely closes.
- *Case B — km-class comet nucleus ($\rho \approx 600$ kg/m³, $\approx 3 \times 10^{11}$ kg), 6 months' warning.* Required momentum: ~ 1 - 4×10^{11} kg·m/s. At closing 20 km/s and a conservative $\beta = 2$ for porous ice, the impactor masses **of order 10^3 - 10^4 tonnes** — hundreds of fully laden heavy-lift flights direct-injected to deep space. Earth launch is excluded outright. The nuclear standoff option — the current planning fallback for exactly this case — suffers uncertain energy coupling into high-porosity nuclei, fragmentation risk that can convert one bullet into a shotgun blast of fragments, and the treaty wall of the Partial Test Ban and Outer Space Treaty [38]. A pre-positioned 10^4 -tonne slug, already characterized, already fueled, resolves the case in a single engagement and requires no treaty waiver.

Conclusion: for the large-and-fast class, a deep-space mass reserve is the only known response that closes simultaneously in physics, on the calendar, and in law.

System design and the double ledger. Three to five nodes — Sun–Earth L4/L5, lunar DRO, and Earth 1:1-resonant quasi-satellite orbits (the family of 2016 HO3, whose materials properties Tianwen-2 is presently establishing) — each holding several characterized bodies of 10^3 - 10^5 t, resident electric tugs, and locally produced propellant (electrolyzed C-type water). In peacetime the nodes *are* UB's aggregation yards and ore depots: **the ammunition ledger and the ark-material ledger are two account names on the same assets.** Defense-readiness contracts fund the standing inventory; habitat construction draws it down; future arks (Section 7.3) inherit the surplus. This is the most elegant reuse in the Program's

economics.

The honest governance clause. A reserve squares the Sagan-Ostro dilemma: a system that can push a body toward a threat can, in principle, push one toward Earth. UB's answer is to make safety a verifiable property of physics rather than a promise: the two-maneuver separation rule of Section 2.4 applied to every reserve body (any Earth-threatening trajectory requires ≥ 2 large maneuvers with months of observable lead, checkable by any nation's tracking network); real-time public publication of all reserve orbits; international observers in the command chain; and no reserve body carrying single-burn propellant margin sufficient to reach an Earth-crossing course. Section 9 places this package within the Program's overall legal architecture.

6.5 The four engines under one doctrine

Engine 1 clears the orbits before the cascade; Engine 2 stages the ore before the demand; Engine 3 delivers the volatiles before the settlement wave; Engine 4 banks the momentum before the comet. **UB sells no heroics at the moment of crisis. UB sells what is already in orbit when the crisis arrives.** The Program never presumes an early awakening of collective human will; it presumes commerce and single-jurisdiction regulation, and spends the years before any awakening compressing the response time after it.

7. Motivation: Three Layers Under One Doctrine

Megaprojects die of motivation failure more often than of technical failure. Apollo died of détente; the Asteroid Redirect Mission died of an election. A fifty-year program must therefore carry a motivation structure that does not depend on any single political cycle. UB carries three layers, addressed to three audiences, all governed by the Time-Asymmetry Principle.

7.1 Layer 1 — Biology: the 1 g question no surface settlement can answer

Thirty years of station data document human deterioration in microgravity: bone-density loss, muscular atrophy, ocular and cardiovascular remodeling [39]. Human data at 0.17 g (Moon) and 0.38 g (Mars) — for long-duration health, for child development, and above all for **reproduction** — is zero. If partial gravity proves inadequate for gestation and development, then surface settlements are permanently rotating-crew outposts, and "multiplanetary species" dead-ends in biology regardless of propulsion progress. The rotating habitat is the only architecture that supplies exact 1 g — and, uniquely, adjustable gravity by radial station, enabling the controlled-gradient experiments that will actually answer the partial-gravity question. **UB is not the competitor of the lunar and Mars programs; it is their biological insurance policy.** This layer rests entirely on published aerospace medicine and leads every academic and agency-facing presentation of the Program.

7.2 Layer 2 — Anti-fragile economics

Section 6 in full. The structural point bears restating: **UB never asks anyone to fund a fifty-year vision.** Each five-year period sells a standalone product — debris recycling, ore

wholesale, volatile delivery, defense readiness — and each product is separately survivable through a change of government, a funding winter, or a founder's absence. This is engineered immunity to the failure mode that killed ARM: grand vision, no interim product, canceled at the next transition.

7.3 Layer 3 — The Ark: from multiplanetary to multistellar

The layer no other settlement architecture can claim begins with a question the others avoid: *what can the Moon and Mars not save us from?*

Answer: risks at the scale of the solar system itself, which all three worlds share. The quantifiable inventory: solar luminosity rises ~10% per billion years, closing Earth's habitability window within roughly 1 Gyr [40]; the star Gliese 710 transits the Oort cloud in ~1.3 Myr at ~0.06–0.17 light-years, with ensuing comet showers indifferent to which inner-system body they strike [41]; nearby gamma-ray bursts and supernovae are likewise system-wide events. Beyond the quantifiable inventory lie systemic unknowns that are unquantifiable in principle yet logically inadmissible to exclude. The first law of risk management is not to predict every catastrophe but to refuse a single common exposure — and "the solar system" is itself a single exposure.

Now the Program's decisive observation: **among all known settlement architectures, only the asteroid-built rotating habitat is structurally isomorphic to a generation ship.** This is engineering fact, not metaphor. A closed ecology, 1 g, meters of rock shielding, a thousand-person society, and zero dependence on any planetary gravity well — this is, line for line, the hull specification of the worldship literature, from Bernal (1929) [1] through the British Interplanetary Society's Daedalus/Icarus lineage and its generation-ship successors [42]. The sole missing subsystem is propulsion — addable at L5 on a century timescale, with cometary ice (deliberately excluded from habitat structure for its mechanical weakness, Section 5) re-entering the architecture as the propellant reserve, a configuration anticipated by the Enzmann starship studies. Planetary settlements are *more baskets*; **UB is the basket that can move.** Multiplanetary secures humanity for centuries; the mobile habitat secures it for millennia. And the Time-Asymmetry Principle delivers this layer's closing argument: if some morning a telescope produces the reason to need an ark, *that* is the most expensive possible morning to begin selecting asteroids.

8. Implementation Roadmap, 2026–2076

8.1 The dual-speed model

Two corrections govern the schedule, one decelerating, one accelerating.

Organizational realism (the decelerating correction). The Program presumes no global treaty consensus, because history shows that treaties redistributing advantage stall indefinitely absent existential pressure. The regulatory architecture is three-layered: (i) a **primary national track** — Outer Space Treaty Article VI requires national authorization and supervision, implemented in the United States through launch licensing, spectrum licensing,

and the mission-authorization framework now being completed in legislation; all critical UB assets and entities sit under this single jurisdiction; (ii) a **coalition instrument** on the Artemis Accords model — the demonstrated template by which 50+ states aligned on norms with zero treaty renegotiation — extending recognition to the reserve and its transparency regime; (iii) **unilateral transparency** — reserve orbits published in real time, verification open to every nation's tracking assets — so that safety legitimacy derives from independently checkable physics rather than from signatures. The cost of this correction is deferred internationalization; the benefit is the removal of the least controllable wait-state from the critical path.

AI acceleration, credited by compressibility (the accelerating correction). Evidence from 2022–2026 (code generation, robotics foundation models, autonomous operations) supports crediting 3–10× compression to: engineering design iteration (simulation-driven substitution for physical test), mission software and verification, operations staffing, and — decisively for Phase 3 — **autonomous swarm construction under 8-minute light delay**, which earlier planning treated as the binding risk and which the current AI trajectory converts into a tailwind. No credit is taken against three incompressible floors: **low-thrust transfer times** (a NEA-to-DRO transfer remains 2–4 years; flyby windows are set by planetary geometry — AI optimizes the route, not the trip); **ecological closure validation** (crop generations and biogeochemical loops run in real time; closure data cannot be simulated into existence); **multi-year human physiology** (the 1-g health and reproduction record requires real people and real years). **AI makes the human-paced work fast; it does not make time-paced work fast. The shape of the schedule is set by the second category; the margin is released by the first.**

8.2 Five phases and two gates

Phase 0 — Foundation (2026–2031). Program incorporation and IP architecture; first peer-reviewed publications; first competitive government concept award (NIAC-class); Engine-1 commercial entry (debris-recycling metallurgy); target-selection pipeline stood up on the Rubin/NEO Surveyor data stream. Capital scale: 10^5 – 10^6 USD. External milestones arriving free of charge: Hera (2026), Tianwen-2 samples (2027), China's kinetic impact and the Apophis encounter (2029).

Phase 1 — 10-meter demonstration (2031–2036). Humanity's first controlled aggregation: two bodies of 1–10 m class (or one body and an instrumented counterweight) brought to contact in lunar DRO at cm/s, anchored, and instrumented — the direct test of contact mechanics and consolidation methods; megawatt-class solar-electric endurance qualification; first commercial ore delivery (a 10-m-class hydrated body to the DRO parking lot); first reserve node stood up under its commercial name. Reference cost: \$2–4B (Keck-baseline extrapolation), sourced from a national flagship line plus commercial joint venture, licensed through the single national track. **Decision Gate 1 (~2035): contact mechanics agree with models; electric propulsion demonstrates $\geq 10^4$ -hour endurance; autonomous assembly passes an uncrewed subscale test.** Pass: 2076 remains on the board. Fail: 5–8-year slip.

Phase 2 — 100-meter prototype (2036–2049). Competitive subscale fly-off among Architectures A/B/C, then full-scale conversion of a selected ~100-m near-Earth body (quasi-satellite family; materials ground truth by then established by Tianwen-2); **first crew**

aboard ~2043 — a date reverse-engineered from the biological floor, since Gate 2 requires ≥ 5 years of crewed data; partial-gravity research decks (0.3-1 g by radial station); first IAT contract executed (lunar polar volatile delivery into the Artemis-era demand window); reserve expanded to three nodes and moved onto standing defense procurement — Engine 4's first independent revenue. **Decision Gate 2 (~2048): multi-year crewed 1-g health data nominal; material closure $\geq 95\%$.**

Phase 3 — Kilometer-scale aggregation (2049-2064). At Sun-Earth L4/L5: a several-hundred-meter core body converted to the rotating nucleus; tens of 10-100-m bodies arriving by flyby cascade, docked and consolidated in sequence; spin-up staged toward 1 g; pressurization and habitation by segment. The protagonists are autonomous robotic swarms and the in-orbit metallurgical line — assets tested in Phase 1, apprenticed in Phase 2, and deployed at scale here, where the AI tailwind pays its full dividend. Capital scale: 10^{11} USD across the phase, borne substantially by the by-then-mature four-engine economy.

Phase 4 — Habitation and closure (2064-2076). Population ramp to $\sim 1,000$; full multi-year ecological validation, including agricultural generations and — a coordinate in human history — the first child born under engineered 1 g; legal status of the artificial body confirmed through the national-track-plus-coalition instrument. **Target: full operation of UB Station One, 2076.**

8.3 Critical path, probability, and the floor sequence

The three external preconditions: a launch-cost revolution (mainstream expectation for the 2030s; not UB-controlled); an autonomy revolution (reclassified from risk to tailwind per Section 8.1; partially UB-accelerated through Phase 0-1 research); a capital mechanism spanning political cycles (the four engines plus long-horizon vehicles). Critical-path analysis is unambiguous: **the 10-m controlled-aggregation demonstration is the longest pole in the entire Program** — everything before it can proceed in parallel; everything after it consumes its data. Advancing its start from the mid-2030s toward the turn of the decade is the highest-leverage action available in the 2020s, and is where the Program's early advocacy, publication, and proposal effort concentrates.

Stated with candor: under the dual-speed model, completion of the full kilometer-class station by 2076 is an aggressive target whose probability is substantially set at Gate 1; the neutral case slides toward 2085-2095. The **floor sequence**, however, holds under neutral assumptions: controlled aggregation demonstrated mid-2030s; a crewed 100-m prototype in the mid-2040s; the kilometer-class main structure closing in the 2060s under optimistic assumptions. The Program's promise is therefore precise: the transition of the aggregated-asteroid habitat from literature to hardware occurs within the working lifetime of people reading this paper — and the probability of the 2076 date is, from this point forward, a function of decisions rather than of physics.

9. Governance, Law, and Safety by Design

Four commitments, stated as design law rather than aspiration. **(1) Two-maneuver separation, universally applied:** no UB asset, commercial or reserve, ever occupies a trajectory from which a single burn produces an Earth-threatening course; the ≥ 2 -maneuver, months-of-observable-lead property is verifiable by any nation's tracking network. **(2) Unilateral transparency:** all asset orbits published in real time; observer access to reserve command chains offered internationally; the safety case is physics anyone can check, not promises anyone must trust. **(3) Jurisdictional clarity:** all critical assets under a single national authorization regime (OST Art. VI), with liability posture per the 1972 Convention accepted rather than contested; coalition recognition sought on the Artemis Accords template rather than through treaty renegotiation. **(4) Planetary protection deference:** Engine 3's Mars line proceeds only within COSPAR process and subordinate to the outcome of Mars life-detection science; the legal identity of the aggregated artificial body — a genuinely novel object in space law — is developed proactively with the academic space-law community rather than left to litigation. The deflection dilemma is real and is doubled by a reserve; UB's position is that the answer to dual-use is not abstinence — the large-and-fast threat class exists whether or not the reserve does — but verifiability.

10. Program Structure, Open Model, and Invitation

UB is published open. All Program documents are released under **CC BY 4.0**: copy, adapt, build, and commercialize freely — with attribution to the *Universe Butterfly (UB) Program, K. Cai, 2025-*. The Program's theory of progress is that a correct architecture, published early with its numbers shown, recruits its own builders. Specific invitations: small-body dynamicists (contact mechanics and consolidation modeling — the Gate 1 problem); ISRU and metallurgical engineers (the Engine 1→5.2 processing line); autonomy and swarm-robotics researchers (the Phase 3 enabler); closed-ecology programs (the Section 5.3 floor); space-law scholars (the artificial-body question); mission designers (the 10-m demonstration concept, which the Program is developing toward agency concept-study solicitations); and agencies and companies for whom any single engine is a business today.

11. Conclusions

1. A kilometer-scale, 1-g, closed-ecology rotating habitat built from asteroid material violates no physics; its design is uniquely forced by four constraints (rotation for gravity, docking-speed aggregation, lead-time-for-energy, and non-Earth-threatening staging).
2. The mass argument against free-flying habitats inverts once the material source is asteroidal: launch carats, mine tonnes.
3. Every precursor capability is flight-demonstrated or in flight as of 2026; the decade to 2031 delivers the remainder at global public expense.
4. Four sequential markets — debris recycling, ore wholesale, volatile delivery, defense reserve — finance the Program in five-year self-sufficient increments, each an application

of the Time-Asymmetry Principle: capability must precede motivation, because it cannot follow it.

5. The deep-space rapid-response reserve is the only response to the large-and-fast threat class that closes in physics, on the calendar, and in law — and its assets are the ark's raw material under a second ledger name.
6. The 2076 target for UB Station One is aggressive and gate-dependent; the floor sequence (aggregation demo mid-2030s, crewed prototype mid-2040s) holds under neutral assumptions. The variable separating the two outcomes is the start date of the 10-m demonstration — a decision, not a law of nature.

The butterfly effect is this Program's namesake and its method: millimeter-per-second thrusts that rewrite a body's destiny over twenty years; small first moves that rewrite a species' options over fifty. The butterfly will leave the garden.

References

- [1] Bernal, J.D. (1929). *The World, the Flesh and the Devil*. Kegan Paul.
- [2] O'Neill, G.K. (1976). *The High Frontier: Human Colonies in Space*. William Morrow; O'Neill, G.K. (1974). "The Colonization of Space." *Physics Today* 27(9).
- [3] Globus, A., et al. "The Kalpana One orbital space settlement." *NSS Space Settlement Journal* and AIAA papers (2007 rev.).
- [4] Miklavčič, P.M., et al. (2022). "Habitat Benu: Design concepts for spinning habitats constructed from rubble pile near-Earth asteroids." *Frontiers in Astronomy and Space Sciences* 9:645363.
- [5] Maindl, T.I., Miksch, R., Loibnegger, B. (2019). "Stability of a rotating asteroid housing a space station." *Frontiers in Astronomy and Space Sciences* 6:37.
- [6] Vera C. Rubin Observatory / LSST solar-system science projections (Ivezić et al., 2019, *ApJ* 873:111).
- [7] NASA NEO Surveyor mission documentation, JPL/UCLA (mission page, 2026; launch NET Sept 2027).
- [8] DeMeo, F., Carry, B. (2013). "The taxonomic distribution of asteroids." *Icarus*; Krasinsky et al. (2002) belt mass estimates.
- [9] Globus, A., Hall, T. (2017). "Space settlement population rotation tolerance." *NSS Space Settlement Journal*.
- [10] Benz, W., Asphaug, E. (1999). "Catastrophic disruptions revisited." *Icarus* 142:5-20.
- [11] Massironi, M., et al. (2015). "Two independent and primitive envelopes of the bilobate nucleus of comet 67P." *Nature* 526:402.
- [12] Northrop Grumman SpaceLogistics, MEV-1/MEV-2 mission reports (2020, 2021).
- [13] Chodas, P., Yeomans, D., keyhole analyses, JPL CNEOS technical literature.
- [14] Vokrouhlický, D., Bottke, W.F., et al. (2015). "The Yarkovsky and YORP effects." In *Asteroids IV*, U. Arizona Press.
- [15] Sagan, C., Ostro, S.J. (1994). "Dangers of asteroid deflection." *Nature* 368:501.
- [16] Santana-Ros, T., et al. (2022). "Orbital stability analysis of Earth Trojan 2020 XL5." *Nature Communications* 13:447.
- [17] Brophy, J., et al. (2012). *Asteroid Retrieval Feasibility Study*. Keck Institute for Space Studies, Caltech/JPL.
- [18] Outer Space Treaty (1967), Arts. VI-VII; Convention on International Liability for Damage Caused by Space Objects (1972).
- [19] NASA radiation shielding studies; Simonsen, L.C., et al., regolith shielding analyses (NASA TP series).

- [20] Naraoka, H., et al. (2023). "Soluble organic molecules in samples of Ryugu." *Science* 379; Oba, Y., et al. (2023). "Uracil in Ryugu samples." *Nature Communications* 14:1292.
- [21] Laretta, D.S., et al. (2024–2025). Bennu sample analyses. *Nature / Nature Astronomy* (evaporites, amino acids, nucleobases).
- [22] NASA Near-Earth Object Human Space Flight Accessible Targets Study (NHATS), CNEOS.
- [23] Cheng, A.F., et al. (2023). "Momentum transfer from the DART mission kinetic impact on asteroid Dimorphos." *Nature* 616:457; Daly, R.T., et al. (2023). *Nature* 616:443.
- [24] ESA Hera mission documentation (2024–2026); arrival November 2026.
- [25] Zou, Y., et al. (2024). "Scientific objectives and payloads of the asteroid defense test mission." *Journal of Deep Space Exploration* 11(2):169–176 (in Chinese); CNSA/Deep Space Exploration Laboratory announcements (2023–2025).
- [26] ESA Ramses mission approval (2025); NASA OSIRIS-APEX extended mission documentation.
- [27] Laretta, D.S., et al. (2022). "Spacecraft sample collection and subsurface excavation of asteroid Bennu." *Science* 377:285 (surface mechanics).
- [28] CNSA Tianwen-2 mission announcements (2025–2026); independent Doppler confirmation of arrival maneuvers, June 2026.
- [29] Company disclosures: AstroForge (DeepSpace-2, Q4 2026); Karman+ (2027 mission); TransAstra NIAC reports.
- [30] Astroscale ADRAS-J mission results (2024); ESA ClearSpace-1 program documentation; U.S. FCC 5-year disposal rule (2022) and first enforcement action (2023).
- [31] Farnocchia, D., et al., C/2013 A1 (Siding Spring) discovery-to-encounter timeline, CNEOS; 1I/Oumuamua, 2I/Borisov, 3I/ATLAS discovery literature (2017, 2019, 2025).
- [32] NASA ISS ECLSS reports: ~98% water recovery achieved (2023).
- [33] Fu, Y., et al., Lunar Palace 365 experiment reports, Beihang University (2018–2020); MELiSSA program, ESA.
- [34] Redwire (Made In Space) ISS additive-manufacturing flight heritage reports.
- [35] Jakosky, B.M., Edwards, C.S. (2018). "Inventory of CO₂ available for terraforming Mars." *Nature Astronomy* 2:634–639.
- [36] Zubrin, R., McKay, C. (1993). "Technological requirements for terraforming Mars." AIAA-93-2005.
- [37] Lunar volatile delivery/retention modeling literature (e.g., Ong, L., et al. 2010, *Icarus*; Prem, P., et al. cold-trap capture simulations).
- [38] National Academies (2022–2023) planetary defense strategy reviews; NNSA/LLNL nuclear deflection coupling studies.
- [39] NASA Human Research Program evidence reports; Twins Study, Garrett-Bakelman et al. (2019), *Science* 364.
- [40] Kasting, J.F. (1988); Wolf, E.T., Toon, O.B. (2015) — solar luminosity evolution and Earth's habitability lifetime.
- [41] Berski, F., Dybczyński, P.A. (2016). "Gliese 710 will pass the Sun even closer." *Astronomy & Astrophysics* 595:L10 (and Gaia-era updates).
- [42] British Interplanetary Society, *JBIS* worldship literature: Project Daedalus (1978), Project Icarus, and generation-ship special issues; Enzmann starship studies.
- [43] Lorenz, E.N. (1963). "Deterministic nonperiodic flow." *Journal of the Atmospheric Sciences* 20:130–141; Lorenz, E.N. (1972). "Predictability: Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?" Address, AAAS 139th Meeting.
- [44] Kardashev, N.S. (1964). "Transmission of information by extraterrestrial civilizations." *Soviet Astronomy* 8:217–221.

Appendix A — Order-of-Magnitude Derivations

A.1 The aggregate-belt body. $M_{\text{belt}} \approx 2.4 \times 10^{21} \text{ kg}$ at $\rho = 3,000 \text{ kg/m}^3 \rightarrow V = 8 \times 10^{17} \text{ m}^3 \rightarrow R = (3V/4\pi)^{1/3} \approx 576 \text{ km}$. $g = GM/R^2 \approx 0.48 \text{ m/s}^2 \approx 0.05 \text{ g}$. $v_{\text{esc}} = \sqrt{2GM/R} \approx 0.75 \text{ km/s}$ — versus the $\sim 6 \times$ thermal-velocity retention criterion, hopeless for N_2/O_2 at habitable temperatures.

A.2 The merger window. $v_{\text{esc}} = \sqrt{2Gm/r} = r\sqrt{8\pi G\rho/3}$. At $\rho = 2,500 \text{ kg/m}^3$: $v_{\text{esc}} \approx 1.18 \times 10^{-3} \text{ s}^{-1} \times r \rightarrow 0.59 \text{ m/s}$ ($r = 500 \text{ m}$), 0.059 m/s ($r = 50 \text{ m}$). Belt collision speeds $\sim 5 \text{ km/s}$ carry specific KE $\sim 1.25 \times 10^7 \text{ J/kg}$ (comparable masses) versus $Q^*_D \sim 10^2\text{--}10^3 \text{ J/kg}$ for km-scale bodies [10]: disruption by four orders of magnitude.

A.3 Shielding versus launch. Areal density 5 t/m^2 ($\approx 2.5 \text{ m}$ regolith at $2,000 \text{ kg/m}^3$) over $56 \text{ km}^2 \rightarrow 2.8 \times 10^{11} \text{ kg}$. At $\$150/\text{kg}$: $\$4.2 \times 10^{13}$. At 150 t per heavy-lift flight: $\approx 1.9 \times 10^6$ flights. Equivalent single asteroid: $2.8 \times 10^{11} \text{ kg}$ at $\rho = 2,000\text{--}2,500 \rightarrow$ diameter $\approx 600\text{--}650 \text{ m}$.

A.4 Transport energetics. Whole 1-km body ($1.3 \times 10^{12} \text{ kg}$) through $\Delta v = 5 \text{ km/s}$: $p = 6.5 \times 10^{15} \text{ kg}\cdot\text{m/s}$; EP propellant at $I_{\text{sp}} = 3,000 \text{ s}$ ($v_e = 29.4 \text{ km/s}$): $m_p = m(e^{\Delta v/v_e} - 1) \approx 2.4 \times 10^{11} \text{ kg}$; jet energy $\approx \frac{1}{2}m_p v_e^2 \approx 10^{20} \text{ J} \approx$ one year of world electricity. Contrast a selected NEA-to-L5 transfer at $\Delta v = 100 \text{ m/s}$ on a 10^6-kg (10-m-class) body: $p = 10^8 \text{ kg}\cdot\text{m/s}$, $m_p \approx 3.4 \times 10^3 \text{ kg}$ — a single tug's budget. Earth-surface reference: $GM_E/R_E = 62.5 \text{ MJ/kg}$ minimum; chemical practice $\sim 10 \times$ with payload fractions of a few percent.

A.5 Deflection sizing. Miss-generation $\Delta v \approx k \cdot R\Theta/t_w$, $k \approx 1\text{--}3$. $t_w = 1 \text{ yr}$: $0.20\text{--}0.61 \text{ m/s}$; 6 months: double. Case A: $m = (4/3)\pi(150)^3 \cdot 2,500 \approx 3.5 \times 10^{10} \text{ kg} \rightarrow \Delta p = 0.7\text{--}2.1 \times 10^{10} \text{ kg}\cdot\text{m/s} = 550\text{--}1,600 \times$ DART's transferred momentum ($\beta \cdot m \cdot v \approx 1.3 \times 10^7$) or $2,000\text{--}6,000 \times$ DART's launch momentum (3.6×10^6). Impactor at $v_{\text{rel}} = 10 \text{ km/s}$, $\beta = 3$: $230\text{--}700 \text{ t}$. Case B: comet, $\rho = 600$, $r = 500 \text{ m} \rightarrow m \approx 3.1 \times 10^{11} \text{ kg}$; $t_w = 6 \text{ mo} \rightarrow \Delta p \approx 1.2\text{--}3.8 \times 10^{11}$; at $v_{\text{rel}} = 20 \text{ km/s}$, $\beta = 2$: $3 \times 10^3\text{--}10^4 \text{ t}$.

A.6 Spin structural loads. 1 g at 2 rpm : $R = g/\omega^2 = 224 \text{ m}$; rim speed $v = \omega R \approx 47 \text{ m/s}$. Self-supporting hull hoop stress $\sigma = \rho_{\text{hull}} v^2 \approx 5.5 \text{ MPa}$ at $\rho = 2,500$ — trivial for steel or basalt fiber. Loaded hull (thin-wall): hoop tension $T = P_{\text{eff}} \cdot R$ with $P_{\text{eff}} =$ atmosphere ($\sim 100 \text{ kPa}$) + rotating shield weight ($5 \text{ t/m}^2 \times 9.8 \approx 49 \text{ kPa}$) + payload; at $P_{\text{eff}} \approx 150 \text{ kPa}$, $T \approx 3.4 \times 10^7 \text{ N/m} \rightarrow \approx 0.14 \text{ m}$ steel-equivalent at 250 MPa allowable. Decoupling the shield (static shell) removes the 49 kPa term and roughly a third of the section — the design driver behind the Kalpana-lineage option noted in §5.1. At 4 rpm tolerance [9], 1 g at $R = 56 \text{ m}$: prototype-scale hulls shrink accordingly.

Universe Butterfly (UB) Program — White Paper v1.0, July 2026. Suggested citation: Cai, K. (2026). "The Universe Butterfly (UB) Program: Aggregated-Asteroid Rotating Habitats — A Scientific Feasibility Assessment and Fifty-Year Implementation Roadmap." White Paper v1.0, universebutterfly.com. Licensed CC BY 4.0. DOI: 10.5281/zenodo.21190731.