Terraforming Mars: Mass, forcing, and industrial throughput constraints

Slava G. Turyshev *
Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Drive, Pasadena, California 91109-0899, USA

(Received 27 February 2026; revised 19 April 2026; accepted 18 May 2026; published 28 May 2026)

Terraforming Mars can be evaluated with system-level constraints linking target pressure and composition to volatile inventory, target temperature to radiative control, climate agents to throughput and power, and persistence to collapse, escape, and geochemical sinks. We use transparent order-of-magnitude scalings to compare endogenous CO_2 release, synthetic supergreenhouse gases, CO_2-H_2 collision-induced absorption, engineered aerosols and nanoparticles, orbital reflectors, and regional paraterraforming. Mars requires 3.89×10^{15} kg of atmosphere per millibar of global mean surface pressure, so open-surface human-relevant pressures require 10^{17}–10^{18}-kg-scale gas inventories. Accessible endogenous CO_2 is best treated as an order-of-tens-of-millibar resource: A representative 20 mbar case gives $\lesssim$10 K warming and does not close melt-class climates or breathable composition. Mean surface temperatures of 250–273 K require, in a gray/window-weighted architecture model, effective window-weighted longwave optical depth of order unity to a few, or direct absorbed-solar forcing of order 10^2 W m^{-2}, implying 10^{13}–10^{14}-m^2-scale reflector area. A human-breathable open atmosphere is dominated by O_2 and buffer gas: $p_{O_2} = 21$ kPa and $p_{N_2} = 50$ kPa imply 9.0×10^{17} kg and 1.9×10^{18} kg, respectively, and water electrolysis requires reversible work $\sim 1.3 \times 10^{25}$ J before inefficiencies and sink-management costs. These bounds apply directly to abiotic industrial oxygenation and to the terminal atmospheric inventories that any human-breathable open Mars must contain; biological bootstrapping can change the path to useful habitability, but not the final O_2 and buffer-gas inventory required for an open breathable atmosphere. Under known inventories, regional and covered-area habitability is the physically favored staged path; no surveyed abiotic global open-atmosphere pathway simultaneously closes pressure, composition, warming, power, throughput, retention, sink, and operations constraints.

DOI: 10.1103/krb8-h3v3

I. INTRODUCTION

Terraforming Mars denotes deliberate, sustained modification of Mars' surface environment—atmospheric pressure and composition, radiative balance, and water stability—toward states that support progressively less technologically mediated surface habitability. The concept is widely discussed, but feasibility is governed by a compact set of planet-scale feasibility constraints that any proposed pathway must satisfy simultaneously.

*Contact author: turyshev@jpl.nasa.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

TABLE I. Acronyms used in this paper.

Acronym	Meaning
CIA	Collision-induced absorption
GCM	General circulation model
IR	Infrared
ISRU	*In situ* resource utilization
LW	Longwave
PFC	Perfluorocarbon (supergreenhouse gas class)
PV	Photovoltaics
SW	Shortwave
TOA	Top of atmosphere
GHG	Greenhouse gas
SRP	Solar radiation pressure

We use four governing constraints as a common quantitative basis for comparing mechanisms: (1) *mass inventory constraint*, linking target pressures and partial pressures to required volatile inventories; (2) *radiative-balance constraint*, linking target surface temperatures to required top-of-atmosphere (TOA) forcing ΔF_{TOA} and/or effective/window-weighted longwave (LW) optical depth τ_{IR}; (3) *throughput and power constraint*, linking inventories and radiative agents to required sustained throughput $\dot{M}$ and power P over specified build times; and (4) *stability/retention constraint*, requiring persistence against atmospheric collapse/condensation, escape to space, and geochemical sequestration. These simplified scalings are intentionally simple: They provide optimistic order-of-magnitude bounds that expose dominant bottlenecks and feasibility floors. Table I provides a list of acronyms used here.

This paper is an engineering-level synthesis. We avoid black-box climate and mission-design optimization tools and instead use transparent scalings to map proposals onto common, easily checked metrics $\{M,\ \tau_{\mathrm{IR}}/\Delta F_{\mathrm{TOA}},\ \dot{M},\ P\}$. The results are therefore lower bounds and feasibility discriminators, not predictions of a specific Mars climate state; quantitative state prediction requires coupled three-dimensional (3D) general circulation model (GCM), aerosol microphysics, and photochemical/geochemical cycling.

Classic terraforming treatments [1–4] introduced key concepts and qualitative feasibility arguments, including early emphasis on warming Mars to enable oxygenic photosynthesis and longer-timescale biological transformation. A synthesis in Ref. [5] explicitly linked the physics of warming Mars, the expected role of biology in longer-timescale oxygenation, and the environmental-ethics questions raised by making Mars habitable. More recent work has expanded the design space to include limits on accessible CO_2 inventories [6–10], regional solid-state greenhouse paraterraforming [11], self-sustaining living-habitat concepts and biologically generated enclosure materials [12,13], engineered aerosol warming proposals [14,15], and renewed discussion of Mars-terraforming research as a coupled climate-biosphere-engineering problem [16]. Work has also begun to formulate explicit *research roadmaps* for assessing nonbiological Mars-warming pathways, including local membranes, orbiting reflectors, and engineered aerosols [17].

The present paper is complementary: Rather than prioritizing research sequencing or downselect criteria, it asks what mass, radiative, throughput, power, and retention conditions would have to be satisfied for specific atmospheric end states to be achieved and maintained. The contribution here is not an additional warming mechanism; it is a unified systems/architecture trade framework that (1) keeps the governing constraints explicit and (2) reduces disparate proposals to common quantitative requirements in mass, forcing/opacity, throughput, and power. This framing highlights which constraint dominates for a given end state and makes key uncertainties (e.g., accessible nitrogen inventory, aerosol lifetimes, and sink capacities) explicit.

Three quantitative facts motivate our constraint-based framework. First, accessible endogenous CO_2 inventories are currently best described as being of *order tens of millibars*, rather than as a secure multibar reservoir; a representative ∼20 mbar case yields $\lesssim$10 K warming under present insolation [6–10]. Second, mechanisms with high radiative leverage at low added atmospheric mass split into two distinct industrial classes: Short-lived aerosol pathways are maintenance limited, whereas CO_2-H_2 collision-induced absorption (CIA) can be fill dominated during 10^2–10^3 yr global build-out even though it still requires replenishment on long hold times. Third, breathable end points are dominated by the required masses of O_2 and buffer gas (N_2/Ar) and by $\gtrsim 10^{25}$-J-class minimum energy requirements for oxygen production even at reversible limits.

The scope of this paper is deliberately narrower than the full terraforming literature. That two-stage logic—abiotic environmental modification followed by longer-timescale biological transformation—was emphasized in Ref. [5]; the present paper addresses the first stage quantitatively and treats the second mainly as a bounding comparator. In particular, biologically mediated oxygenation and ecosystem bootstrapping are important pathways in classic and recent Mars-terraforming discussions [1,16,18], but their full quantitative analysis requires additional variables—net areal productivity, nutrient supply, burial/export efficiency, ecological stability, and sink competition—that are treated here only at bounding-scaling level. The oxygenation numbers in this paper should therefore be read as lower bounds for *abiotic* industrial oxygenation, not as a claim that biology is irrelevant.

This paper is organized as follows: Section II defines end states (E0–E4) and baseline Mars parameters. Section III develops atmospheric mass bookkeeping and buffer-gas requirements. Section IV establishes radiative requirements, including a minimal greenhouse mapping linking T_e and T_s. Section V bounds what can be achieved by mobilizing endogenous CO_2 and H_2O. Section VI evaluates engineered and exogenous warming options [perfluorocarbon (PFC)-class gases, CO_2-H_2 CIA, engineered aerosols, mirrors/albedo]. Section VII quantifies the mass and minimum energy requirements for oxygenation and breathable end points. Section VIII discusses loss, collapse, and sequestration constraints and introduces a minimal state-space control model. Section IX treats energy delivery and system architecture. Section X discusses implications, industrial scaling, timescale-power trade, and cost floors; and Sec. XI concludes.

II. WHAT "TERRAFORMING MARS" MEANS IN QUANTITATIVE TERMS

A. End states and success metrics

The term *terraforming* is used loosely in the literature, but technical feasibility depends strongly on the end point being targeted. We therefore distinguish several *end states*, labeled E0–E4, where the letter "E" denotes an *end state* (equivalently, an end point label). These labels are bookkeeping shorthand, not additional physical variables. For technical analysis, it is useful to define them with measurable criteria:

(1) *E0:"Robotic/human assisted operations" (today).* Current mean $P_s \sim 610$ Pa and $T_s \sim 210$ K, requiring pressure suits and thermal control.

(2) *E1: "Recurrent open-surface liquid water."* In the present climate, Mars can already cross the triple-point condition locally and transiently, so merely exceeding 611.657 Pa is not by itself a useful progression criterion. In this paper, E1 denotes a more useful engineering target: recurrent or seasonally accessible *surface* liquid water in some regions under open-sky conditions, which requires both pressure above the triple-point threshold and a local energy balance that can sustain liquid water for longer than brief instantaneous episodes [19–21]. A minimum necessary condition is

$$P_s > 611.657\,\text{Pa at } 273.16\,\text{K}, \tag{1}$$

following the triple-point relation for water [22], but this is far from sufficient for useful surface-water stability (Sec. IV C).

(3) *E2: "Enclosed-volume agriculture/paraterraforming."* A representative benchmark is an *internal* habitat pressure $P_{\rm in} \sim 10$–30 kPa, with the remaining thermal and compositional margins supplied by enclosure design and environmental control. For present Mars ambient pressure, such operating envelopes reduce the differential structural load by roughly a factor of $\sim$3 to $\sim$10 relative to a 101 kPa habitat, depending on the chosen internal pressure. E2 is therefore primarily an enclosure and area-deployment problem rather than a global atmospheric-pressure problem.

(4) *E2.5: Biological bootstrapping and productive biosphere support.* This end point denotes biologically productive systems that are decoupled from direct human exposure requirements. Examples include enclosed or semiprotected phototrophic ecosystems, engineered microbial or algal systems, biologically generated materials, food production, local O_2 production, soil conditioning, and ecological succession. E2.5 does not require pressure-unassisted human exposure or a human-breathable open atmosphere. Its success metrics are areal productivity, water and nutrient closure, radiation shielding, ecological stability, useful biomass/material output, and net O_2 export after respiration and sinks.

(5) *E3: "Pressure-unassisted human exposure."* The hard lower bound for avoiding ebullism is total pressure $P_s \geqslant 6.27$ kPa (the Armstrong limit) [23]. E3 is a pressure threshold, not a breathable-atmosphere criterion. Human activity at E3 still requires a supplied breathing mixture with acceptable p_{O_2} and p_{CO_2}, together with thermal and decompression management. Ambient CO_2 is not the defining E3 constraint unless the ambient atmosphere itself is being breathed; that case belongs to the E4 composition problem.

(6) *E4: "Human-breathable surface atmosphere."* A useful distinction is between (a) lower-pressure O_2-rich operating envelopes that can be adequate in controlled habitats or vehicles, and (b) buffered open-surface atmospheres with $p_{O_2} \sim 21$ kPa and tens of kilopascals of inert gas to control physiology, flammability, and broader ecological function. The representative open-surface benchmark used here is $p_{O_2} \sim 21$ kPa plus N_2/Ar buffer gas to reach $P_s \sim 50$–100 kPa. This end point is dominated by the required mass of O_2 and buffer gases (Sec. VII).

Unless explicitly qualified otherwise, "habitability" in this paper means *human-relevant surface habitability or protected biosphere utility*, not the broader astrobiological definition of any environment capable of supporting some form of life. Microbial or phototrophic habitability can occur under substantially different pressure, temperature, radiation, and compositional requirements [12].

For feasibility accounting, E4 should be decomposed rather than treated as a single colloquial end point. At minimum, it combines (1) a *pressure* requirement, (2) an *oxygen* requirement, (3) a *buffer-gas* requirement, and (4) a *thermal* requirement. This decomposition matters because the dominant bottleneck is not the same for each subtarget. Warming can, in principle, be assisted by low-mass radiative agents, whereas breathable composition remains controlled by 10^{18}-kg-scale O_2 and buffer-gas inventories. In other words, a proposal can satisfy the warm-state criterion while still failing the breathable-atmosphere criterion by many orders of magnitude in mass and energy.

Classic terraforming proposals and feasibility discussions span several decades and include early quantitative treatments of Mars warming and atmospheric engineering [2–4]. This paper emphasizes the distinction between (1) *global atmospheric modification* and (2) *regional/paraterraforming* strategies that create habitable microclimates without altering planetary-scale atmospheric mass.

We use representative mean values (Table II) for Mars' physical and radiative parameters. Solar irradiance at Mars varies substantially between perihelion and aphelion; we use the orbit-averaged TOA value $S_{\rm Mars} \approx 589\ {\rm W\,m^{-2}}$ and a mean planetary albedo $A_{\rm B} \approx 0.25$ [24]. These yield an effective radiating temperature $T_{e0} \approx 210$ K, consistent with Mars' cold climate.

B. Key results at a glance

Table III lists reference pressure thresholds and representative atmospheric inventories relevant to the end point hierarchy developed in Sec. II A. Table IV summarizes the requirement targets used throughout this paper as a compact "dashboard" linking end states (E0–E4) to (1) required inventories, (2) radiative control targets, and (3) implied industrial rates. This table makes explicit a

TABLE II. Nominal Mars parameters used in first-order calculations.

Quantity	Symbol	Value
Mean radius	R_{Mars}	3389.5 km
Surface gravity	g_{Mars}	$3.71\ \mathrm{m\,s^{-2}}$
Surface area	$4\pi R_{\mathrm{Mars}}^2$	$1.44 \times 10^{14}\ \mathrm{m^2}$
Mean TOA solar irradiance	S_{Mars}	$589\ \mathrm{W\,m^{-2}}$
Bond albedo (mean)	A_{B}	∼0.25
Effective radiating temperature[a]	T_{e0}	$[S_{\mathrm{Mars}}(1-A_{\mathrm{B}})/(4\sigma_{\mathrm{SB}})]^{1/4} \approx 210\ \mathrm{K}$

[a] σ_{SB} is the Stefan-Boltzmann constant.

central theme of the paper: As the end point advances from E2 to E4, the dominant constraint shifts from local radiative control and deployment area to 10^{18} -kg-scale inventories and multicentury industrial throughput and power.

C. End-point-normalized lower bounds and dimensionless feasibility numbers

A central difficulty in comparing terraforming pathways is that they are usually presented in incompatible variables: pressure release, trace-gas radiative forcing, aerosol optical depth, mirror area, or industrial power. A more reusable formulation is to express every pathway relative to the same end-point-normalized lower bounds (see Table V).

For any global mean partial-pressure target p_i, define

$$M_{\mathrm{req}}(p_i) = \mathcal{K}_{\mathrm{Mars}}\, p_i, \quad \mathcal{K}_{\mathrm{Mars}} \equiv \frac{4\pi R_{\mathrm{Mars}}^2}{g_{\mathrm{Mars}}} = 3.89 \times 10^{13}\ \mathrm{kg\,Pa^{-1}} = 3.89 \times 10^{16}\ \mathrm{kg\,kPa^{-1}}. \quad (2)$$

Hence, the corresponding minimum build throughput over a build time t_{build} is

$$\dot{M}_{\mathrm{req}}(p_i, t_{\mathrm{build}}) = \frac{\mathcal{K}_{\mathrm{Mars}}\, p_i}{t_{\mathrm{build}}} \approx 1.23 \times 10^6 \left(\frac{p_i}{1\ \mathrm{kPa}}\right)\left(\frac{10^3\ \mathrm{yr}}{t_{\mathrm{build}}}\right)\mathrm{kg\,s^{-1}}. \quad (3)$$

TABLE III. Reference pressure thresholds and representative atmospheric inventories on Mars. Atmospheric mass uses $M_{\mathrm{atm}} \simeq (4\pi R_{\mathrm{Mars}}^2/g_{\mathrm{Mars}})\, P_s$ from Eq. (19) and column mass uses $m_{\mathrm{col}} = P_s/g_{\mathrm{Mars}}$. [Notes: (1) 1 bar is shown as a convenient Earth sea-level reference. (2) The E1 row tabulates only the triple-point pressure threshold associated with E1; the full E1 end state additionally requires recurrent or seasonally accessible open-surface liquid water under favorable local energy balance (Sec. IV C). (3) The end state E2 is omitted here because it is typically implemented regionally inside habitats (paraterraforming); if applied globally, $P_s \sim 10$ kPa would imply $M_{\mathrm{atm}} \approx 3.9 \times 10^{17}$ kg [Eq. (19)]. (4) The CO_2 row is a representative accessible endogenous reference case consistent with current *order-of-tens-of-millibar* estimates, not a secure hard ceiling. (5) The E4 row is a pressure reference only; the full breathable end point is decomposed in Table IV into oxygen, buffer-gas, pressure, and thermal requirements.]

Label	Meaning	P_s (Pa)	P_s (mbar)	M_{atm} (kg)	m_{col} ($\mathrm{kg\,m^{-2}}$)
E0	Present mean	6.1×10^2	6.1	2.37×10^{16}	1.64×10^2
E1 (threshold)	Triple-point threshold	6.12×10^2	6.12	2.38×10^{16}	1.65×10^2
–	Representative accessible CO_2 case	2.0×10^3	20	7.78×10^{16}	5.39×10^2
E3 (threshold)	Armstrong limit (no ebullism)	6.27×10^3	62.7	2.44×10^{17}	1.69×10^3
–	0.2 bar (thick but not human breathable)	2.0×10^4	200	7.78×10^{17}	5.39×10^3
E4 (pressure ref.)	1 bar (Earth sea-level reference)	1.0×10^5	1000	3.89×10^{18}	2.70×10^4

TABLE IV. Requirement dashboard for representative end points. Inventories use $M \simeq (4\pi R_{\rm Mars}^2/g_{\rm Mars})\,P$. Radiative targets assume that $T_{e0} \approx 210$ K unless noted. Where ranges are shown, they correspond to build times $t_{\rm build} = 10^3$ yr (left) to 10^2 yr (right). All power numbers are thermodynamic minima unless stated otherwise. $\dot{M}$ assumes global deployment; E2 is typically regional. *Notes*: (1) If $P_s \sim 10$ kPa were applied *globally*, Eq. (19) implies $M_{\rm atm} \approx 3.9 \times 10^{17}$ kg; E2 is typically implemented regionally (paraterraforming/habitats), so global inventory is not required. (2) E3 is a pressure-only threshold. Human activity at E3 requires a supplied breathing mixture and thermal/decompression management; a breathable ambient atmosphere is an E4 composition requirement.

End state	Success criterion (global unless noted)	Pressure/p_i	Inventory M (kg)	Radiative/control target	Industrial implication
E0	Current Mars baseline	$P_s \approx 0.611$ kPa (6.1 mbar)	2.37×10^{16}	$T_{e0} \approx 210$ K; $T_s \approx 210$ K	None (baseline)
E1	Recurrent open-surface liquid water (regional)	$P_s \gtrsim 0.612$ kPa	$\approx 2.38 \times 10^{16}$	Local pressure + energy balance; triple-point crossing alone is insufficient	Requires accessible H_2O plus regional heat/insulation and water-vapor control; not limited by global atmospheric inventory
E2	Enclosed-volume agriculture (typically regional)	$P_{\rm in} \sim$ 10–30 kPa (inside habitats)	–	Local greenhouse/insulation/ enclosure control	Deployment area + local power; scalable in phases
E2.5	Biological bootstrapping/ productive biosphere support	Regional/enclosed or semiprotected	Area, water, nutrient, and ecology limited	Productive photosynthesis, food, biomaterials, local O_2, soil conditioning	Not bounded by abiotic electrolysis power; requires biological productivity, water access, nutrient supply, protection, and sink closure
E3	No ebullism (Armstrong limit)	$P_s = 6.27$ kPa	2.44×10^{17}	To approach $T_s \gtrsim 250$ K: $\tau_{\rm IR,eff} \sim$ 2–3 (or high-authority $\Delta F_{\rm TOA}$)	If pursued globally: $\dot{M} \sim 7.7 \times 10^6$ to 7.7×10^7 kg s^{-1} (inventory limited)
E4 (breathable): inventories + energy dominate; warming still required as in E3.					
E4a	O_2 partial-pressure target	$p_{O_2} = 21$ kPa	9.0×10^{17}	Noncollapse warm-state climate + sink management (Sec. VII D)	$\dot{M}_{O_2} \sim 2.8 \times 10^7$ to 2.8×10^8 kg s^{-1}; $\bar{P}_{\rm min} \sim 0.42$–4.2 PW
E4b	Buffer-gas example	$p_{N_2} = 50$ kPa	1.9×10^{18}	Inventory (climate independent)	$\dot{M}_{N_2} \sim 6.0 \times 10^7$ to 6.0×10^8 kg s^{-1}; if exogenous, momentum/transport dominates

TABLE V. Dimensionless feasibility numbers introduced in Sec. II C. The mechanism survey can be reexpressed in this language, making the dominant bottleneck explicit and falsifiable.

Quantity	Definition	Feasible regime	Meaning
Π_M	$M_{\rm avail}/M_{\rm req}$	$\gtrsim 1$	Inventory sufficiency
Π_F	$\Delta F_{\rm avail}/\Delta F_{\rm req}$ or $\tau_{\rm IR,eff}/\tau_{\rm IR,req}$	$\gtrsim 1$	Radiative control sufficiency
$\Pi_{\dot M}$	$\dot M_{\rm avail}/\dot M_{\rm req}$	$\gtrsim 1$	Industrial throughput sufficiency
Π_P	$P_{\rm avail}/P_{\rm req}$	$\gtrsim 1$	Power sufficiency
Π_S	$\dot M_{\rm repl}/\dot M_{\rm loss}$	$\gtrsim 1$	Ability to hold the engineered state against loss
$\Pi_{\rm min}$	$\min\{\Pi_M, \Pi_F, \Pi_{\dot M}, \Pi_P, \Pi_S\}$	$\gtrsim 1$	Conjunctive feasibility criterion
$\Lambda_{\rm maint}$	$t_{\rm build}/\tau_{\rm loss}$	$\ll 1$ fill dominated; $\gtrsim 1$ maintenance dominated	Distinguishes one shot from continuously replenished pathways

For a constituent i embedded in a well-mixed multicomponent atmosphere, the inventory associated with a target partial pressure p_i is not, in general, species independent. Instead,

$$M_i \simeq \mathcal{K}_{\rm Mars}\,\frac{\mu_i}{\bar\mu}\,p_i, \tag{4}$$

where μ_i is the molecular mass of constituent i and $\bar\mu$ is the mean molecular mass of the atmospheric mixture. Equations (2) and (3) therefore apply directly to *total pressure* or to a dominant constituent; for light minor constituents such as H_2 in a CO_2-dominated atmosphere, inventory and fill-rate requirements are reduced by the factor $\mu_i/\bar\mu$.

If the source terms $\dot M_i$ are physical mass flow rates, then conversion from mass flow to partial-pressure tendency includes the same multicomponent factor used in Eq. (4):

$$\dot p_i \simeq \frac{g_{\rm Mars}}{4\pi R_{\rm Mars}^2}\,\frac{\bar\mu}{\mu_i}\,\dot M_i, \tag{5}$$

for slowly varying $\bar\mu$. Equivalently, if $\dot M_i^{\rm phys}$ is a physical mass flow rate, the equations below may be read in terms of pressure-equivalent mass flows

$$\dot M_i^{\rm pe} \equiv \frac{\bar\mu}{\mu_i}\,\dot M_i^{\rm phys}. \tag{6}$$

For oxygen made from water electrolysis, using the reversible minimum specific work $\varepsilon_{O_2} \simeq 14.8\,\mathrm{MJ\,kg^{-1}}$ of O_2, the minimum energy associated with an O_2 partial-pressure target is

$$E_{O_2,\rm min}(p_{O_2}) = \varepsilon_{O_2}\mathcal{K}_{\rm Mars}\,\frac{\mu_{O_2}}{\bar\mu}\,p_{O_2} \approx 5.76\times 10^{23}\left(\frac{\mu_{O_2}/\bar\mu}{1}\right)\left(\frac{p_{O_2}}{1\ \mathrm{kPa}}\right)\mathrm{J}. \tag{7}$$

For a representative O_2-N_2 molecular-weight range, $\mu_{O_2}/\bar\mu \simeq 1.1$, so this correction changes the oxygenation energy by only $\sim$10%, but it keeps the bookkeeping internally consistent. The corresponding build-time-averaged reversible power floor is

$$\bar P_{O_2,\rm min}(p_{O_2}, t_{\rm build}) \approx 18.2\left(\frac{\mu_{O_2}/\bar\mu}{1}\right)\left(\frac{p_{O_2}}{1\ \mathrm{kPa}}\right)\left(\frac{10^3\ \mathrm{yr}}{t_{\rm build}}\right)\mathrm{TW}. \tag{8}$$

If a buffer gas is imported exogenously with characteristic delivery speed Δv, a useful *transport-energy benchmark* is

$$E_{\rm imp,min}(p_i, \Delta v) \simeq \frac{1}{2}\mathcal{K}_{\rm Mars}\,p_i\,\Delta v^2 \approx 4.86\times 10^{23}\left(\frac{p_i}{1\ \mathrm{kPa}}\right)\left(\frac{\Delta v}{5\ \mathrm{km\,s^{-1}}}\right)^2\mathrm{J}, \tag{9}$$

with corresponding average power

$$\bar{P}_{\rm imp,min}(p_i,\,\Delta v,\,t_{\rm build}) \approx 15.4\Big(\frac{p_i}{1\ {\rm kPa}}\Big)\Big(\frac{\Delta v}{5\ {\rm km\ s^{-1}}}\Big)^2\Big(\frac{10^3\ {\rm yr}}{t_{\rm build}}\Big)\ {\rm TW}. \tag{10}$$

The intent of Eqs. (9) and (10) is not to define a universal mission-energy floor, because gravity assists, aerocapture, and sail-assisted transfers can greatly reduce propulsive expenditure. Rather, they provide a compact normalization for the transport-scale energy associated with moving planetary-atmosphere-class mass inventories.

These end-point-normalized lower bounds motivate five dimensionless feasibility numbers:

$$\Pi_M \equiv \frac{M_{\rm avail}}{M_{\rm req}}, \quad \Pi_{\dot M} \equiv \frac{\dot M_{\rm avail}}{\dot M_{\rm req}}, \quad \Pi_P \equiv \frac{P_{\rm avail}}{P_{\rm req}}, \tag{11}$$

$$\Pi_F \equiv \begin{cases} \Delta F_{\rm avail}/\Delta F_{\rm req}, & \text{TOA forcing pathway,} \\ \tau_{\rm IR,eff}/\tau_{\rm IR,req}, & \text{greenhouse pathway,} \end{cases} \quad \Pi_S \equiv \frac{\dot M_{\rm repl}}{\dot M_{\rm loss}}. \tag{12}$$

Because planetary modification is a conjunctive problem rather than a weighted-average problem, the relevant bottleneck metric is

$$\Pi_{\rm min} \equiv \min\{\Pi_M,\,\Pi_F,\,\Pi_{\dot M},\,\Pi_P,\,\Pi_S\}. \tag{13}$$

A pathway is feasibility complete for a given end point only if $\Pi_{\rm min} \gtrsim 1$. This minimum-operator formulation is deliberate: Failure in any one constraint vetoes the end point, regardless of strength in the others.

To distinguish one-shot fill operations from truly maintenance-limited pathways, define

$$\Lambda_{\rm maint} \equiv \frac{t_{\rm build}}{\tau_{\rm loss}} = \frac{\dot M_{\rm loss}}{M^\star/t_{\rm build}}, \tag{14}$$

where $M^\star$ is the maintained inventory of the active climate agent. Pathways with $\Lambda_{\rm maint} \ll 1$ are fill dominated over the build phase; pathways with $\Lambda_{\rm maint} \gtrsim 1$ are already maintenance dominated during buildup.

D. Worked examples under the Π rubric

The utility of Eqs. (11)–(14) is that they classify which constraint fails first, or which constraint dominates, before a full climate simulation is attempted. Three examples show how the framework is intended to be used.

Example 1: Endogenous CO_2 is not a robust route to warm global E3/E4 states. Taking a representative accessible endogenous CO_2 reference case of $P_s \approx 20$ mbar (Sec. V) and comparing with the E3 threshold $P_s = 62.7$ mbar gives

$$\Pi_M^{(CO_2 \to E3)} \lesssim \frac{20}{62.7} \approx 0.32. \tag{15}$$

This should not be read as a strict upper bound: Recent carbonate and polar inventory work permits larger order-of-tens-of-millibar estimates, and the pressure-only E3 threshold could in principle be approached if accessible CO_2 inventories are near the high end of current uncertainty. The more robust conclusion is radiative and compositional. Tens of millibars of CO_2 do not provide melt-class global warming under current insolation, and a CO_2-dominated atmosphere is not a human-breathable E4 atmosphere. Endogenous CO_2 can therefore help pressure and warming, but it does not by itself close the warm, breathable, open-surface end point.

Example 2: Global CO_2-H_2 CIA is fill dominated on century-to-millennial build times. For the illustrative case $P_{\rm tot} = 0.5$ bar and $f_{H_2} = 0.05$ (Sec. VI B 3), the required H_2 inventory is $M^\star_{H_2} \approx 4.7 \times 10^{15}$ kg once the multicomponent correction $\mu_{H_2}/\bar\mu$ is included. Building that inventory over

10^3 yr requires an average fill throughput

$$\dot{M}_{\mathrm{H_2,fill}} \simeq \frac{M^{\star}_{\mathrm{H_2}}}{t_{\mathrm{build}}} \approx 1.5 \times 10^5 \ \mathrm{kg\,s^{-1}}, \tag{16}$$

whereas the diffusion-limited replenishment benchmark is only $\dot{M}_{\mathrm{H_2,loss}} \sim 6 \times 10^3 \ \mathrm{kg\,s^{-1}}$. Equivalently, using Eq. (52),

$$\Lambda_{\mathrm{maint}} = \frac{t_{\mathrm{build}}}{\tau_{\mathrm{H_2,loss}}} \approx 4 \times 10^{-2} \quad (t_{\mathrm{build}} = 10^3 \ \mathrm{yr}). \tag{17}$$

Hence, CO_2-H_2 CIA is not primarily a replenishment-limited pathway during build-out; it is primarily a global-fill pathway with a long-duration maintenance tail.

Example 3: Particle warming is maintenance dominated from the outset. If a particle pathway has atmospheric residence time $\tau_p \sim$ 30–100 days, then for any century-scale build schedule

$$\Lambda^{(p)}_{\mathrm{maint}} = \frac{t_{\mathrm{build}}}{\tau_p} \sim 10^3\text{–}10^4 \gg 1. \tag{18}$$

Thus, the dominant requirement is not cumulative injected mass but the sustained industrial rate needed to hold the atmospheric column against removal. This is why particle pathways and H_2-CIA pathways should not be grouped together merely because both are "mass-efficient warming levers": Their industrial signatures are fundamentally different.

III. ATMOSPHERIC MASS BOOKKEEPING: PRESSURE TARGETS IMPLY 10^{17}–10^{18}-kg-SCALE GASES

A. Pressure-mass relation

For a thin atmosphere in hydrostatic equilibrium, the global atmospheric mass required to produce a mean surface pressure P_s is approximately

$$M_{\mathrm{atm}} \simeq \frac{4\pi R^2_{\mathrm{Mars}}}{g_{\mathrm{Mars}}} P_s. \tag{19}$$

Using Table II, the conversion is

$$M_{\mathrm{atm}} \approx 3.89 \times 10^{15} \ \mathrm{kg\,mbar^{-1}} \ (\mathrm{Mars}). \tag{20}$$

Figure 1 plots Eq. (19) over 0.1–1000 mbar.

Table III highlights several pressure end points relevant to "success" criteria. The central message is scale: Moving from present-day Mars (∼6 mbar) to even modest human-relevant pressures is a 10^{17}–10^{18}-kg-class problem, depending on the end point.

Also, from Eq. (19), we derive atmospheric mass per unit pressure

$$\frac{dM_{\mathrm{atm}}}{dP_s} = \frac{4\pi R^2_{\mathrm{Mars}}}{g_{\mathrm{Mars}}} \approx 3.89 \times 10^{13} \ \mathrm{kg\,Pa^{-1}} \approx 3.89 \times 10^{15} \ \mathrm{kg\,mbar^{-1}}. \tag{21}$$

Equations (19) and (20) are more than bookkeeping. They show that every additional 1 mbar of global mean surface pressure on Mars costs about 3.89×10^{15} kg of atmosphere, so discussions framed only in pressure units can hide the true industrial scale of the problem. In particular, the step from present-day Mars to the Armstrong-limit threshold is not a modest atmospheric adjustment but a jump to $\sim 2.4 \times 10^{17}$ kg, and the step to an Earth sea-level atmosphere is nearly 2 orders of magnitude larger again. The hydrostatic conversion therefore already fixes the class of the problem: Once the goal is global E3 or E4, the atmosphere is a 10^{17}–10^{18}-kg-scale engineering object.

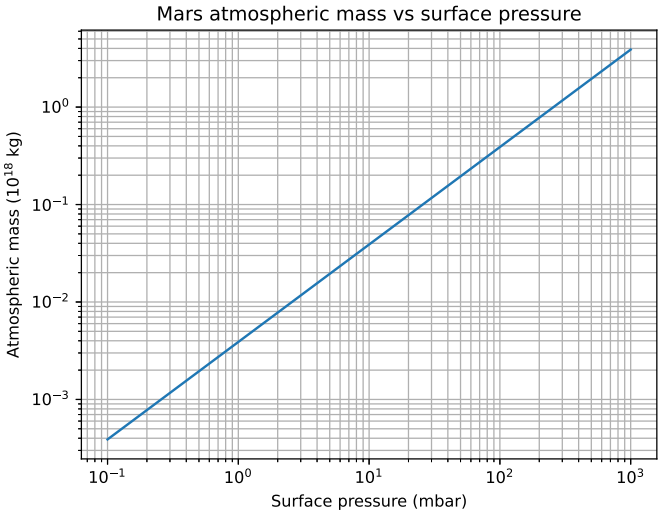


FIG. 1. Atmospheric mass required for a given mean surface pressure on Mars from Eq. (19). A 1 bar atmosphere corresponds to $M_{\rm atm} \approx 3.89 \times 10^{18}$ kg.

B. Composition matters: CO_2 is a poor buffer gas

Table VI summarizes end-point-normalized lower bounds on Mars. Pressure targets do not specify composition, but composition determines physiology and climate.

TABLE VI. End-point-normalized lower bounds on Mars, expressed per 1 kPa of target pressure. The inventory/throughput rows apply directly to *total pressure* or to a dominant constituent. For a minor constituent i in a multicomponent atmosphere, the corresponding inventory and fill rate must be multiplied by $\mu_i/\bar{\mu}$ [Eq. (4)]. Species-specific production/import energies enter through the relevant specific work or delivery architecture. Here, $\varepsilon_{O_2} \simeq 14.8\ \mathrm{MJ\,kg^{-1}}$ is the reversible minimum work to make O_2 from water, and $\varepsilon_{H_2} \simeq 1.2 \times 10^8\ \mathrm{J\,kg^{-1}}$ is the reversible minimum work to make H_2 from water. For a representative CO_2-dominated CIA atmosphere with $\bar{\mu} \approx 41.9\ \mathrm{g\,mol^{-1}}$, the H_2 factor is $\mu_{H_2}/\bar{\mu} \approx 0.0477$, so the H_2 energy/power requirements per kilopascal are about $\sim 21\times$ smaller than in pure H_2.

Quantity	Expression	Numerical value
Global inventory per 1 kPa	$\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})$	3.89×10^{16} kg
Global fill rate per 1 kPa over $t_{\rm build}$	$\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})/t_{\rm build}$	$1.23 \times 10^6\,(10^3\ \mathrm{yr}/t_{\rm build})\ \mathrm{kg\,s^{-1}}$
Reversible O_2 production energy per 1 kPa	$\varepsilon_{O_2}\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})$	5.76×10^{23} J
Reversible O_2 production power floor per 1 kPa	$\varepsilon_{O_2}\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})/t_{\rm build}$	$18.2\,(10^3\ \mathrm{yr}/t_{\rm build})$ TW
Kinetic-energy transport benchmark per 1 kPa at $\Delta v = 5\ \mathrm{km\,s^{-1}}$	$\frac{1}{2}\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})\Delta v^2$	4.86×10^{23} J
Transport-power benchmark per 1 kPa at $\Delta v = 5\ \mathrm{km\,s^{-1}}$	$\frac{1}{2}\mathcal{K}_{\rm Mars}(1\ \mathrm{kPa})\Delta v^2/t_{\rm build}$	$15.4\,(10^3\ \mathrm{yr}/t_{\rm build})$ TW
Electrolysis energy to produce H_2 for 1 kPa H_2 partial pressure in a multicomponent atmosphere	$\varepsilon_{H_2}\mathcal{K}_{\rm Mars}(\mu_{H_2}/\bar{\mu})(1\ \mathrm{kPa})$	$4.67 \times 10^{24}(\mu_{H_2}/\bar{\mu})$ J
Electrolysis power floor for 1 kPa H_2 partial pressure in a multicomponent atmosphere	$\varepsilon_{H_2}\mathcal{K}_{\rm Mars}(\mu_{H_2}/\bar{\mu})(1\ \mathrm{kPa})/t_{\rm build}$	$148(\mu_{H_2}/\bar{\mu})(10^3\ \mathrm{yr}/t_{\rm build})$ TW

1. CO_2 *toxicity constraint*

High CO_2 partial pressures are directly hazardous to humans and many plants. Therefore, even if CO_2 is the primary available endogenous greenhouse gas, a breathable end point cannot consist of "$CO_2 + O_2$" alone at tens of kilopascals. A viable human-breathable E4 atmosphere requires a buffer gas (N_2 and/or Ar) at the tens of kilopascal level.

2. *Available buffer gas on Mars*

Mars' present atmosphere is ∼95% CO_2 with minor N_2 and Ar. Scaling these minors to tens of kilopascals requires either (1) importing buffer gas from external sources or (2) mining and releasing geochemically stored nitrogen (nitrates, ammonium minerals) if sufficient inventories exist (uncertain at planetary scale). This is a first-order roadblock for E4.

Not every human-use atmosphere requires tens of kilopascals of inert diluent. Historical spacecrafts have used low-pressure O_2-rich atmospheres for crewed operations, and this design space remains relevant for controlled habitats and vehicles. However, such cases should be distinguished from an open planetary surface end point: Removing the N_2/Ar burden can reduce the *buffer-gas* mass requirement, but it does not eliminate the 10^{18}-kg-scale O_2 requirement itself, and it introduces stronger flammability, decompression, and ecosystem-composition constraints. In this paper, buffered atmospheres remain the representative E4 open-surface benchmark, whereas low-buffer-gas cases are better interpreted as habitat/vehicle operating envelopes rather than planetary end states.

A useful quantitative comparison is that eliminating the inert diluent removes the *buffer-gas* mass requirement, but it does not remove the oxygen mass requirement itself. Because $M_{O_2} \simeq \mathcal{K}_{\rm Mars} p_{O_2}$ for an open planetary atmosphere, a pure-O_2 operating envelope at $p_{O_2} \sim 34.5\,$kPa would still imply $M_{O_2} \sim 1.34 \times 10^{18}\,$kg if applied planetwide, i.e., larger than the human-breathable $p_{O_2} = 21\,$kPa benchmark. Low-buffer-gas cases therefore change the *composition split* of E4, but they do not remove the 10^{18}-kg-scale oxygen inventory if interpreted as an open-surface planetary end point [23].

C. Column mass as a proxy for radiation and aerodynamic shielding

Atmospheric column mass is $m_{\rm col} = P_s/g_{\rm Mars}$. This quantity is directly relevant to (1) attenuation of energetic particles and the mass loading relevant to entry/descent/aerodynamic performance, and only *indirectly* relevant to UV shielding. In particular, energetic-particle shielding scales approximately with column mass, whereas surface UV depends strongly on atmospheric composition (e.g., ozone) and on dust/cloud opacity rather than on column mass alone [25,26]. For representative end points,

$$m_{\rm col} \approx 2.7 \times 10^3\ {\rm kg\,m^{-2}} \left(\frac{P_s}{10\ {\rm kPa}} \right), \tag{22}$$

so $P_s = 6.27\,$kPa (Armstrong limit) corresponds to $m_{\rm col} \approx 1.7 \times 10^3\ {\rm kg\,m^{-2}}$, while $P_s = 101\,$kPa corresponds to $m_{\rm col} \approx 2.7 \times 10^4\ {\rm kg\,m^{-2}}$ on Mars (because $g_{\rm Mars} < g_\oplus$). Thus, pressure-building strategies also directly modify shielding and entry/descent regimes, even before composition is made breathable.

D. Buffer-gas mass scales and import energetics

A breathable E4 atmosphere requires a *buffer gas* (typically N_2 and/or Ar) at tens of kilopascals. The mass requirement follows directly from Eq. (19). For a well-mixed O_2-N_2 atmosphere, the buffer-gas inventory associated with a target surface partial pressure is

$$M_{N_2} \simeq \mathcal{K}_{\rm Mars} \frac{\mu_{N_2}}{\bar{\mu}}\, p_{N_2}. \tag{23}$$

For the representative buffered case $p_{O_2} = 21\,\mathrm{kPa}$ and $p_{N_2} = 50\,\mathrm{kPa}$, one has $\bar{\mu} \simeq 29.2\,\mathrm{g\,mol^{-1}}$, so

$$M_{N_2} \approx 1.87 \times 10^{18}\,\mathrm{kg}\left(\frac{p_{N_2}}{50\,\mathrm{kPa}}\right)\left(\frac{\mu_{N_2}/\bar{\mu}}{28/29.2}\right), \tag{24}$$

comparable to the mass of a 0.5 bar planetary atmosphere.

1. External sourcing

If buffer gas must be imported (as N_2, NH_3, or N-bearing ices), a useful first benchmark is the characteristic transport energy associated with delivering mass M with characteristic delivery speed Δv,

$$E_k \sim \tfrac{1}{2} M\,\Delta v^2. \tag{25}$$

For $M \sim 10^{18}\,\mathrm{kg}$ and $\Delta v \sim 5\,\mathrm{km\,s^{-1}}$, one finds $E_k \sim 10^{25}\,\mathrm{J}$, comparable to the oxygenation energy scale in Eq. (78). This should *not* be interpreted as a universal lower bound on mission-energy expenditure: Gravity assists, aerocapture, sail-assisted trajectories, and staged capture architectures can all reduce the propulsive component substantially.

The more robust conclusion is logistical rather than propulsive. Even when the transfer architecture is energetically favorable, importing $\sim 10^{18}\,\mathrm{kg}$ of volatiles remains a planetary-scale *mass-handling* problem. The hard questions are capture of volatile-bearing bodies or cargo streams, controlled breakup or processing where needed, storage/transfer into Mars-bound logistics, and finally global atmospheric delivery and distribution. For feasibility classification, the dominant issue is therefore not a single mission Δv, but whether a Mars-centered industrial system can repeatedly process and redistribute planetary-atmosphere-class mass inventories.

2. NH_3 *as an N carrier*

If nitrogen is imported primarily as ammonia, photolysis and chemistry can convert it into N_2 while hydrogen escapes, leaving N_2 as a buffer gas. Stoichiometrically, producing one mole of N_2 (28 g) requires two moles of NH_3 (34 g) to supply two N atoms, so the imported mass is at least $\sim 1.21\times$ the desired N_2 mass, before losses.

These mass and momentum scales imply that E4 atmospheres are difficult to reach without either enormous exogenous mass logistics or a yet-unquantified endogenous nitrogen inventory.

E. How much fixed nitrogen would Mars need to supply a buffer gas?

If Mars were to supply buffer gas endogenously, the limiting question is the global inventory of fixed nitrogen (e.g., nitrates) accessible to mining and release. Let an accessible regolith layer of depth d and bulk density ρ contain a mean nitrogen mass fraction w_N. The total nitrogen mass is

$$M_N \simeq 4\pi R_{\mathrm{Mars}}^2\,\rho\,d\,w_N, \tag{26}$$

and because N_2 is nitrogen by mass, the achievable buffer-gas pressure is

$$p_{N_2} \simeq \frac{g_{\mathrm{Mars}}\,M_N}{4\pi R_{\mathrm{Mars}}^2} \simeq g_{\mathrm{Mars}}\,\rho\,d\,w_N. \tag{27}$$

Solving for the required nitrogen mass fraction,

$$w_N \simeq \frac{p_{N_2}}{g_{\mathrm{Mars}}\rho d} \approx 6.7 \times 10^{-2}\left(\frac{p_{N_2}}{50\,\mathrm{kPa}}\right)\left(\frac{\rho}{2000\,\mathrm{kg\,m^{-3}}}\right)^{-1}\left(\frac{d}{100\,\mathrm{m}}\right)^{-1}. \tag{28}$$

Thus, supplying $p_{N_2} = 50\,\mathrm{kPa}$ from the top $\sim 100\,\mathrm{m}$ would require an implausibly high planetwide mean nitrogen fraction (percent level by mass). Even $d = 1\,\mathrm{km}$ requires $w_N \sim 6.7 \times 10^{-3}$, still extremely large on a global average. This does not rule out regional nitrogen extraction, but it strongly suggests that E4 buffer-gas inventories are more naturally met by exogenous sourcing

unless deep, nitrogen-rich deposits exist. *In situ* observations do indicate fixed nitrogen (nitrate) in Martian sediments, but the global inventory remains poorly constrained [27].

Equation (27) demonstrates that supplying $\mathcal{O}(10–50)$ kPa of buffer gas as a *global mean* from shallow regolith requires implausibly large planetwide w_N. This does not exclude the possibility that localized nitrogen-rich deposits could support regional E2/E2.5 covered systems, because required gas mass then scales with enclosed volume rather than planetary surface area. For a covered region of area $A_{\rm reg}$ and height H,

$$M_{\rm N_2,encl} \simeq A_{\rm reg} H \frac{\mu_{\rm N_2} p_{\rm N_2}}{R_g T} \simeq 2.0 \times 10^{12}\ {\rm kg} \left(\frac{A_{\rm reg}}{10^6\ {\rm km}^2}\right)\left(\frac{H}{3\ {\rm m}}\right)\left(\frac{p_{\rm N_2}}{50\ {\rm kPa}}\right)\left(\frac{250\ {\rm K}}{T}\right). \tag{29}$$

The corresponding open-column mass over the same area would be $A_{\rm reg} p_{\rm N_2}/g_{\rm Mars}$, larger by $H_s/H \sim 10^4$. Thus, regional covered systems are not controlled by the same buffer-gas mass scaling as global E4.

As we have seen, Sec. III showed that E3 and above are intrinsically 10^{17}–10^{18}-kg-scale mass problems: $P_s = 6.27$ kPa implies $M_{\rm atm} \approx 2.4 \times 10^{17}$ kg (Table III), while E4 O_2 plus buffer-gas inventories exceed 10^{18} kg. However, pressure alone does not imply habitability; the next question is whether Mars can be warmed to $T_s \gtrsim 250$–273 K under current insolation. Section IV therefore derives radiative targets in the two natural control variables: $\Delta F_{\rm TOA}$ for insolation changes and $\tau_{\rm IR}$ for greenhouse pathways.

F. Local log sensitivities and uncertainty propagation

Because the paper's objective is feasibility discrimination rather than high-fidelity state prediction, the most useful uncertainty metric is not a raw parameter range but the *local logarithmic sensitivity* of each derived requirement to each uncertain input. For a quantity $Y(\{x_a\})$, define

$$S_{x_a}^{(Y)} \equiv \frac{\partial \ln Y}{\partial \ln x_a}. \tag{30}$$

For independent fractional input uncertainties, first-order propagation gives

$$\left(\frac{\sigma_Y}{Y}\right)^2 \approx \sum_a \left[S_{x_a}^{(Y)}\right]^2 \left(\frac{\sigma_{x_a}}{x_a}\right)^2. \tag{31}$$

This representation is especially useful here because many requirement scalings are approximately separable power laws. For example, if engineered particles are treated through an effective forcing-per-column parameter $\gamma_p \equiv \partial \Delta F/\partial \Sigma_p$, then

$$\Sigma_p \propto \frac{\Delta F_{\rm req}}{\gamma_p}, \quad \dot{M}_p \propto \frac{\Delta F_{\rm req}}{\gamma_p\, \tau_p}, \quad A_m \propto \frac{\Delta F_{\rm TOA}}{\eta_m S_{\rm Mars}}, \quad p_{\rm N_2,achievable} \propto \rho\, d\, w_N. \tag{32}$$

Thus, the dominant uncertainties differ by mechanism: IR-active particle pathways are most sensitive to residence time τ_p and effective radiative efficacy γ_p; mirror pathways are linear in required forcing and inverse in optical/geometric efficiency; endogenous N_2 pathways are linear in accessible depth and nitrogen mass fraction; and O_2 sink estimates are linear in the effective oxidizable depth and FeO fraction. Table VII provides relevant details.

Reporting these sensitivities makes the uncertainty structure explicit and prevents false equivalence between uncertainties that merely move a coefficient by a factor of a few and uncertainties that change the feasibility class.

TABLE VII. Local log-sensitivity coefficients for key requirement quantities. Nonzero coefficients are shown explicitly. This table identifies which uncertain inputs can realistically change a pathway classification and which only perturb prefactors.

Derived quantity Y	Leading scaling	Nonzero log sensitivities $S_x^{(Y)} = \partial \ln Y / \partial \ln x$
Maintained particle column Σ_p	$\Delta F_{\rm req}/\gamma_p$	$S_{\Delta F} = +1,\ S_{\gamma_p} = -1$
Particle injection rate $\dot{M}_p$	$\Delta F_{\rm req}/(\gamma_p \tau_p)$	$S_{\Delta F} = +1,\ S_{\gamma_p} = -1,\ S_{\tau_p} = -1$
Mirror area A_m	$\Delta F_{\rm TOA}/(\eta_m C_m S_{\rm Mars})$	$S_{\Delta F} = +1,\ S_{\eta_m} = -1,\ S_{C_m} = -1,\ S_{S_{\rm Mars}} = -1$
Achievable endogenous $p_{\rm N_2}$	$\rho\, d\, w_N$	$S_\rho = +1, S_d = +1, S_{w_N} = +1$
O_2 sink capacity $M_{\rm O_2,sink}$	$w_{\rm FeO}\rho\, d_{\rm eff}(t)$	$S_{w_{\rm FeO}} = +1, S_\rho = +1, S_{d_{\rm eff}} = +1$
H_2 replenishment power $P_{\rm H_2,min}$	$\Phi_0 f_{\rm H_2}$	$S_{\Phi_0} = +1, S_{f_{\rm H_2}} = +1$
H_2 loss time $\tau_{\rm H_2,loss}$	$P_{\rm tot}\Phi_0^{-1}$	$S_{P_{\rm tot}} = +1, S_{\Phi_0} = -1$

IV. THERMAL REQUIREMENTS: RADIATIVE FORCING AND THE ∼60 K PROBLEM

A. Global mean forcing needed for large temperature changes

A minimal, transparent way to relate temperature targets to energy balance is through the effective radiating temperature T_e at TOA:

$$\sigma_{\rm SB} T_e^4 = \frac{S_{\rm Mars}}{4}(1 - A_{\rm B}) + \Delta F_{\rm TOA}, \tag{33}$$

where $\Delta F_{\rm TOA}$ is an imposed perturbation to the net absorbed flux (e.g., from orbital mirrors or albedo change). In this simple model, raising T_e from T_{e0} to T_e requires

$$\Delta F_{\rm TOA} \ = \ \sigma_{\rm SB}\left(T_e^4 - T_{e0}^4\right). \tag{34}$$

For Mars, increasing T_e by 30 K (from 210 to 240 K) requires $\Delta F_{\rm TOA} \approx 78\ {\rm W\,m^{-2}}$, while 60 K requires $\Delta F_{\rm TOA} \approx 191\ {\rm W\,m^{-2}}$. These are extremely large global perturbations.

1. Surface temperature T_s versus T_e: why the same "60 K" implies different requirements?

Equations (33) and (34) quantify the TOA forcing required to raise the effective radiating temperature T_e directly (e.g., via mirrors or albedo changes). This is the correct metric for *insolation-modifying* approaches. However, greenhouse approaches primarily seek to raise the surface temperature T_s at approximately fixed absorbed-solar power by increasing longwave optical depth, thereby increasing T_s relative to T_e. Accordingly, the $\Delta F_{\rm TOA}$ values computed above should be interpreted as bounds for *direct TOA forcing* methods, whereas greenhouse methods are more naturally expressed in terms of the required infrared optical depth $\tau_{\rm IR}$ (Sec. IV B).

2. Why "waste heat" is not a global terraforming lever?

A useful sanity check is to compare industrial waste heat to planetary radiative fluxes. A global mean forcing of just $\Delta F = 10\ {\rm W\,m^{-2}}$ corresponds to total power $\Delta P = 4\pi R_{\rm Mars}^2 \Delta F \approx 1.4 \times 10^{15}$ W, i.e., ∼1400 TW. Therefore, even a civilization dissipating ∼10 TW of waste heat would supply only ∼0.07 ${\rm W\,m^{-2}}$ globally. Direct heat injection can matter locally, but *global* warming requires radiative (albedo/insolation) or greenhouse modifications.

The surface temperature exceeds T_e when the atmosphere provides greenhouse effect. Atmospheric-modification proposals therefore focus on increasing greenhouse opacity (longwave trapping) rather than directly increasing absorbed-solar flux. Greenhouse pathways are better compared through the additional longwave opacity required to obtain a target T_s at nearly fixed absorbed-solar power, whereas mirror/albedo pathways are naturally compared through explicit $\Delta F_{\rm TOA}$.

B. Relating TOA forcing to surface temperature: a minimal greenhouse mapping

Terraforming objectives are typically stated in terms of surface temperature T_s, whereas Eqs. (33)–(37) constrain the effective radiating temperature T_e at TOA. Greenhouse interventions primarily alter the mapping $T_s(T_e)$ by increasing longwave optical depth, rather than by imposing a direct TOA forcing ΔF_{TOA}. The window/saturation language is most relevant for thickened or engineered atmospheres; it is not intended to imply that present-day ∼6 mbar Mars already behaves like a dense, band-saturated CO_2 atmosphere outside the strongest 15-μm region.

A standard first-order approximation is the Eddington gray-atmosphere relation for radiative equilibrium:

$$T_s^4 \simeq \tfrac{3}{4} T_e^4 \left(\tau_{\mathrm{IR}} + \tfrac{2}{3}\right), \tag{35}$$

where τ_{IR} is the broadband infrared optical depth from the surface to space. Throughout the remainder of this paper, τ_{IR} should be read as an effective, window-weighted longwave optical depth, $\tau_{\mathrm{IR,eff}}$, unless a spectrally resolved optical depth is explicitly specified. It is a diagnostic architecture parameter, not a claim of gray, spectrally uniform opacity.

Solving Eq. (35) for the optical depth required to reach a target T_s gives

$$\tau_{\mathrm{IR}} \simeq \frac{4}{3}\left(\frac{T_s}{T_e}\right)^4 - \frac{2}{3}. \tag{36}$$

Note that the gray-atmosphere relation in Eq. (35) is used here only to define an *architecture-level, order-of-magnitude target* for the longwave opacity required to raise T_s at fixed absorbed-solar power; it is *not* a substitute for line-by-line radiative transfer or 3D GCM predictions of a specific Mars climate state.

For Mars' present $T_{e0} \approx 210\,\mathrm{K}$ (Table II), achieving $T_s = 273\,\mathrm{K}$ implies $\tau_{\mathrm{IR}} \approx 3.1$, while $T_s = 250\,\mathrm{K}$ implies $\tau_{\mathrm{IR}} \approx 2.0$. These values provide a transparent target for greenhouse strategies: CO_2 alone at $P_s \sim \mathcal{O}(10\ \mathrm{mbar})$ does not supply sufficient τ_{IR} under current insolation, consistent with detailed inventory and climate constraints [6].

A further useful mapping relates τ_{IR} to pressure through a mass absorption coefficient κ_{IR}:

$$\tau_{\mathrm{IR}} \sim \kappa_{\mathrm{IR}} \frac{P_s}{g_{\mathrm{Mars}}}, \tag{37}$$

highlighting the trade between (1) increasing P_s (at large gas mass cost) and (2) increasing κ_{IR} via trace supergreenhouse gases or aerosol absorption in spectral window regions. Equation (37) is not intended as a substitute for line-by-line radiative transfer, but as a first-order engineering scaling relation that exposes which proposals require 10^{18}-kg-scale pressure increases versus those that require high-κ_{IR} engineered constituents.

For real gases such as CO_2, Eq. (37) should be read as an *effective bookkeeping relation*, not as a literal statement that forcing grows linearly with pressure at all P_s. In spectrally saturated bands, additional warming is controlled by line wings, window closure, and pressure broadening, so the incremental greenhouse effect can be strongly sublinear in P_s even when an effective $\kappa_{\mathrm{IR,eff}}$ is convenient for architecture-level trades.

Uncertainties in translating a target T_s into a required $\tau_{\mathrm{IR,eff}}$ are generically $\mathcal{O}(\pm 30\%\text{–}50\%)$, driven by nongray spectral structure (windows versus saturated bands), dust/cloud feedbacks, lapse-rate changes, and orbital/seasonal variability. This uncertainty does *not* affect the 10^{18}-kg-scale mass conclusions for E3/E4, which follow directly from hydrostatic pressure-mass relation (Sec. III) and composition requirements (Sec. VII).

1. Spectral interpretation of the gray τ_{IR}

In a nongray Mars atmosphere, incremental surface warming is often controlled by (1) *spectral window regions* that are weakly absorbing in the baseline state, (2) *pressure broadening* that extends absorption into line wings, and (3) cloud/dust feedbacks that redistribute LW opacity and shortwave

(SW) heating. Accordingly, the gray $\tau_{\rm IR}$ in Eq. (35) should be interpreted as an *effective, window-weighted longwave optical depth* required to obtain a target T_s/T_e, not as a literal spectrally uniform opacity.

For architecture trades, it is therefore useful to introduce a "window-filling efficacy" factor $w_{\rm win} \in (0, 1]$ such that a mechanism with mass absorption coefficient $\kappa_{\rm IR}$ produces

$$\tau_{\rm IR,eff} \sim w_{\rm win}\,\kappa_{\rm IR}\,\frac{P_s}{g_{\rm Mars}}, \tag{38}$$

where $w_{\rm win} \sim 1$ for constituents that add opacity primarily in LW window regions and $w_{\rm win} \ll 1$ for constituents that add opacity mainly in already-saturated bands. Equation (38) makes explicit that the relevant engineering question is not "how much $\tau_{\rm IR}$ is added," but "how much *window-weighted* opacity can be added per kg m^{-2} column."

2. Sensitivity to orbital, seasonal, latitudinal, and diurnal radiative state

The required optical depth depends on T_e through Eq. (36). Using $T_e = [S(1 - A_{\rm B})/(4\sigma_{\rm SB})]^{1/4}$, Mars' orbit-averaged value is $T_{e0} \approx 210$ K (Table II), but T_e varies with heliocentric distance and albedo. For illustration, if T_e spans ∼201–221 K across Mars' orbit for fixed $A_{\rm B} = 0.25$, then the required $\tau_{\rm IR}$ for $T_s = 273$ K spans

$$\tau_{\rm IR}(273\ {\rm K}) \approx 2.5\text{–}3.9, \tag{39}$$

compared to $\tau_{\rm IR} \approx 3.1$ at $T_e = 210$ K. This sensitivity underscores that the greenhouse mapping is an engineering-scale target range rather than a single predictive number; a full treatment would additionally resolve latitude, local time, and cloud/dust structure with a 3D GCM.

C. Why the water triple point is not a sufficient criterion?

Exceeding the triple point of water does not ensure stable surface liquid water. Even at $P_s \gtrsim 6.1$ mbar, liquid water can remain metastable only briefly because of the following factors:

(1) *Evaporation into an unsaturated atmosphere.* Mars' atmosphere is extremely dry; evaporation rates can be high even when boiling is suppressed.

(2) *Radiative cooling and freezing.* Without sustained warming, transient melt refreezes quickly.

(3) *Local energy balance dominates.* Insolation varies strongly with latitude, season, time of day, and dust loading; liquid water stability is therefore spatially and temporally localized.

The conclusion that crossing the triple point is insufficient for useful *surface* liquid water is consistent with a large literature on transient melting, evaporation, brines, and slope microclimates on Mars [19–21]. Even if polar CO_2 mobilization raises pressure modestly above the triple point, the temperature shortfall remains large and evaporation/sublimation cooling is severe [6]. Even if polar CO_2 is mobilized to exceed the triple point, liquid water would not be stable at the surface because the temperature shortfall remains $\mathcal{O}(60\ {\rm K})$ and evaporation is rapid.

D. Model validity and uncertainty: what the simplified scalings do and do not claim?

The constraints used here are intended as feasibility discriminators, not detailed climate predictions. The gray-atmosphere mapping from T_e to T_s [Eq. (35)] compresses spectral structure (window regions, line saturation, pressure broadening) and dynamical feedbacks (dust, clouds, lapse-rate changes) into a single effective $\tau_{\rm IR}$. Consequently, inferred $\tau_{\rm IR}$ targets should be interpreted as order-of-unity ranges (e.g., $\tau_{\rm IR} \sim 2$–4 for $T_s \sim 250$–273 K at present T_e), not point values. Mechanism feasibility is therefore assessed by whether a proposal can plausibly supply $\mathcal{O}(1)$ optical-depth changes [or $\mathcal{O}(10^2)$ W m^{-2} TOA forcing for mirror/albedo approaches] at the required timescale and maintenance burden.

Where the scaling indicates feasibility, higher-fidelity coupled 3D GCM + aerosol microphysics + photochemistry is required to determine stability, spatial structure, and control authority. Where

the scaling indicates severe shortfall (e.g., endogenous CO_2 inventories versus E3/E4 pressures), additional modeling cannot remove the mass deficit.

E. Synthesis: the feasibility rubric implied by the governing constraints

Sections III and IV reduce terraforming to a small set of quantitative targets. For any end point E*i* (Sec. II A), a proposal must simultaneously satisfy the following:

(1) *Mass inventory constraint*: Achieve the required pressure and composition inventories. For E3 ($P_s = 6.27$ kPa), Eq. (19) implies $M_{\rm atm} \approx 2.4 \times 10^{17}$ kg (Table III). For E4, $p_{O_2} = 21$ kPa implies $M_{O_2} \approx 9.0 \times 10^{17}$ kg [Eq. (73)], and a modest $p_{N_2} = 50$ kPa buffer implies $M_{N_2} \approx 1.9 \times 10^{18}$ kg [Eq. (24)].

(2) *Radiative requirement*: Provide either (a) direct TOA forcing $\Delta F_{\rm TOA}$ (mirrors/albedo) or (b) sufficient effective/window-weighted longwave optical depth $\tau_{\rm IR,eff}$ to raise T_s at fixed absorbed-solar power. At current insolation, Eq. (36) implies $\tau_{\rm IR} \sim 2$–4 for $T_s \sim 250$–273 K.

(3) *Industrial throughput/power constraint*: Supply required inventories/absorbers at rate $\dot{M} = M/t_{\rm build}$ [Eq. (98)] and power $P = E/t_{\rm build}$ [Eq. (99)]. Even the reversible minimum for oxygenation is $E_{\rm min} \approx 1.3 \times 10^{25}$ J [Eq. (78)], corresponding to ∼420 TW averaged over 10^3 yr.

(4) *Stability/maintenance constraint*: Ensure that the engineered state is stable against loss and sinks. For H_2-based warming, diffusion-limited escape implies a long-duration replenishment requirement [Eqs. (81)–(83); Table IX], even if the initial global build-out is fill dominated for $t_{\rm build} \ll \tau_{\rm H_2,loss}$. For CO_2-centric pathways, collapse/condensation introduces hysteresis risk (Sec. VIII D).

This rubric turns the mechanism survey (Secs. V and VI) into a controlled comparison: Each mechanism is evaluated by *which constraint dominates* (inventory limited, radiatively limited, throughput/power limited, or maintenance limited) and whether that dominance shifts with the target end point. Here "control authority" means the achievable change in a planetary output variable (e.g., T_s, P_s, or p_{CO_2}) per unit deployed actuation, together with the ability to maintain that output against seasonal forcing, stochastic disturbances, and internal feedbacks.

We have seen that Sec. IV provided the radiative yardstick: Global "melt-class" climates (i.e., a climate warm enough for recurrent or regionally stable open-surface liquid water) require either very large direct forcing ($\Delta F_{\rm TOA} \sim 10^2$ W m^{-2} for large T_e shifts) or a greenhouse optical depth $\tau_{\rm IR} \sim$ 2–4 at current absorbed-solar power. The remainder of the paper evaluates candidate mechanisms by whether they can plausibly deliver these targets given the mass constraints of Sec. III and the industrial constraints developed later in Sec. IX. We begin with what Mars can supply endogenously (Sec. V).

V. ENDOGENOUS VOLATILE MOBILIZATION: WHAT MARS CAN PROVIDE

A. CO_2 reservoirs and the ∼20 mbar reference case within an order-of-tens-of-millibar budget

The most cited quantitative constraint on Mars terraforming is the inventory of accessible CO_2. Jakosky and Edwards [6] argued that readily accessible reservoirs (polar deposits, adsorbed regolith CO_2, and near-surface carbonates) plausibly sum to only ∼0.02 bar under conservative assumptions. More recent work has refined the inventory picture, including buried south polar CO_2 exchange and stratigraphic constraints [8–10] and recently reported carbonate-bearing strata observed by Curiosity [7]. Taken together, the present evidence supports an *order-of-tens-of-millibar* accessible endogenous CO_2 budget rather than a secure multibar reservoir. In what follows, ∼20 mbar is used as a representative reference case for scaling.

Two consequences follow:

(1) *Pressure shortfall relative to E4 and uncertain closure for E3.* A representative ∼20 mbar accessible CO_2 case corresponds to $M_{\rm atm} \sim 7.8 \times 10^{16}$ kg (Table III), below the Armstrong threshold of 62.7 mbar. Higher order-of-tens-of-millibar estimates could reduce or even close the

TABLE VIII. Mechanism comparison matrix (order of magnitude). "One shot" indicates a fill/build operation dominated by initial deployment, "maintenance" indicates a sustained industrial rate set by loss/removal timescales, and "hybrid" denotes pathways for which build-out is fill dominated but long-duration hold requires replenishment. Values indicate which constraint dominates: inventory limited (I), radiatively limited (R), throughput/power limited (T), or maintenance limited (M).

Mechanism	Control variable/efficacy	Required scale (illustrative)	Lifetime/replenishment	Dominant bottleneck/limiting constraint
Endogenous CO_2 mobilization	P_s increase via released CO_2, modest greenhouse	Representative accessible case $P_s \lesssim \mathcal{O}(10\text{–}20\ \text{mbar})$	One shot (mostly), but collapse risk	Accessible-inventory limit; limited warming [6–10]; I + R
PFC-class supergreenhouse gas (super-GHG)	Strong window absorption; κ_{IR} in window bands	Feedstock + synthesis at extreme scale	Long-lived but chemistry-dependent; maintenance	Industrial chemistry/feedstock; lifetime [29,30]; T (+M if short-lived)
CO_2-H_2 CIA	CO_2-rich background pressure P_{bg} plus H_2 mole fraction f_{H2}; published warm cases generally require $P_{bg} \sim 0.5\text{–}1$ bar and percent-level H_2	Dense background inventory plus H_2 fill: $M_{bg} \sim K_{Mars} P_{bg} \sim 2\text{–}4 \times 10^{18}$ kg for $P_{bg} = 0.5\text{–}1$ bar, plus multimillibars to multikilopascals H_2	Hybrid: fill dominated for $t_{build} \ll \tau_{H_2,loss}$; maintenance on long hold times	Background-atmosphere inventory first; H_2 fill and escape second. Known accessible endogenous CO_2 does not supply the required P_{bg} [31]; I/T during build-out, M + T in hold phase
Engineered aerosols/ nanoparticles	IR-active particles with high thermal-IR/visible extinction ratio; maintained column Σ_p	Σ_p at mg m^{-2} to sub-g m^{-2} class depending on lifetime/design	Residence time months to years (design- and dynamics-dependent); maintenance	Lifetime, lofting, and radiative efficacy [14,15]; M + T
Orbital mirrors/reflector constellations	Direct ΔF_{TOA}; aggregate area + control	A_m continent scale for moderate global forcing; 10^{13}-m^2 -class for melt-class forcing	Long-duration control, degradation, and partial replacement over hold times	Structure mass + deployment count + pointing + orbit control + replacement (global); R + T (+M on long hold times)
Regional solid-state greenhouse/ paraterraforming	Local radiative transfer	Centimeter-scale layers over km^2-continent regions	One shot (regional) + maintenance of coverage	Manufacturing/placement at area; local operations [11]; T (area/logistics)

TABLE IX. Illustrative diffusion-limited H_2 replenishment requirements using Eq. (82) and a reversible electrolysis specific energy of 1.2×10^8 J kg^{-1} of H_2.

f_{H_2}	$\dot{M}_{H_2}$ (kg s^{-1})	$P_{H_2,\min}$ (TW)
0.01	1.2×10^3	0.14
0.05	6.0×10^3	0.72
0.10	1.2×10^4	1.44

pressure-only E3 gap, but they would still not provide a human-breathable E4 atmosphere and would not by themselves solve the global warming requirement.

(2) *Warming shortfall.* For a 20-mbar-CO_2 atmosphere, climate models predict warming *less than* 10 K at current solar output [6]. Reaching temperatures close to melting would require ∼1 bar CO_2, well beyond currently discussed accessible inventories [6].

a. Implied gray κ_{IR} scale for CO_2 under present insolation

Equation (37) provides a useful inverse inference. If detailed studies indicate that approaching melting under present solar output would require $p_{CO_2} \sim \mathcal{O}(1\ \text{bar})$ (as discussed in Ref. [6]), then achieving a representative $\tau_{IR} \sim 3$ at $P_s \sim 10^5$ Pa implies an effective gray

$$\kappa_{IR,\text{eff}} \sim \frac{\tau_{IR} g_{\text{Mars}}}{P_s} \sim 1 \times 10^{-4}\ \text{m}^2\,\text{kg}^{-1}. \tag{40}$$

With this $\kappa_{IR,\text{eff}}$, a $P_s = 20$-mbar-CO_2 atmosphere corresponds to $\tau_{IR} \sim 0.05$, consistent with the conclusion that mobilizable CO_2 inventories cannot provide near-melting global climates.

B. Energetics of CO_2 release: sublimation versus mining

Releasing polar CO_2 ice is energetically straightforward in principle (sublimation), but not necessarily easy in practice. If the buried south polar CO_2 deposit corresponds to ∼0.006 bar (6 mbar) [6], its mass is $M \approx 2.3 \times 10^{16}$ kg. Using a representative CO_2 latent heat of sublimation $L_{\text{sub}} \sim 6 \times 10^5$ J kg^{-1} [28], the ideal energy to sublimate this reservoir is $\sim 10^{22}$ J, or several terrawatts sustained over a century. Mining carbonates is harder in an *industrial throughput* sense: Although the per-kilogram heating energy for carbonate feedstock need not exceed the latent heat of CO_2 sublimation, carbonate release requires excavation, beneficiation, bulk solids handling, reactor processing, and heat delivery to geologic material rather than direct phase change of a volatile deposit. The difficulty is therefore dominated by material handling and processing logistics, not by the simple per-kilogram enthalpy comparison alone.

The representative accessible endogenous CO_2 reference case (∼20 mbar) implies both a pressure shortfall relative to E3/E4 and a warming shortfall relative to melt-class climates. Therefore, any global strategy must rely on engineered opacity/absorbers (PFCs, aerosols, CIA) and/or direct insolation modification (mirrors/albedo), which are treated in Sec. VI.

VI. EXOGENOUS AND ENGINEERED WARMING MECHANISMS

Given the representative accessible endogenous CO_2 reference case discussed in Sec. V, any pathway toward $T_s \gtrsim 250$ K or toward E3/E4 end points must rely on engineered radiative levers and/or imported volatiles. Table VIII summarizes the mechanism classes considered here and identifies whether each is (1) one shot/fill dominated, (2) maintenance dominated, or (3) hybrid, with build-out and long-duration hold controlled by different constraints.

Table VIII groups the mechanism classes that can plausibly alter Mars' radiative balance or atmospheric inventory at architecture scale. Other proposed levers are not omitted; they are screened by first-principles constraints. Waste heat is local, because 10 W m^{-2} of global forcing requires

$\sim 1.4 \times 10^{15}$ W. H_2O vapor is a feedback rather than a first-step greenhouse lever, because it depends on prior warming, ice redistribution, clouds, and cold traps. CH_4, NH_3, SO_2, HNO_3, H_2O_2, C_2H_4, and related trace gases can be radiatively active, but their terraforming relevance is set by source inventory, atmospheric lifetime, condensation/solubility, toxicity, and photochemical by-products. Exogenous volatile import is not a separate shortcut: It is the buffer-gas and mass logistics problem written in another form. Biological pathways are distinct because they can create useful regional biosphere services before open-air breathability, but they remain controlled by productive area, nutrients, water, burial/export efficiency, and sink competition.

A. Synthetic supergreenhouse gases (PFC class)

Artificial greenhouse gases (perfluorocarbons and related species) have been proposed to warm Mars efficiently because of strong absorption in infrared window regions and long photochemical lifetimes. Quantitative radiative-convective modeling has been performed for this class of gases [29,30]. Two engineering constraints dominate:

(1) *Feedstock and synthesis scale (PFC).* For a trace greenhouse gas X with volume mixing ratio q_X in a background atmosphere of pressure $P_{\rm tot}$, the required inventory is

$$M_X \simeq \mathcal{K}_{\rm Mars}\,\frac{\mu_X}{\bar{\mu}}\,q_X P_{\rm tot} \approx 3.9\times 10^{12}\ {\rm kg}\left(\frac{\mu_X/\bar{\mu}}{2}\right)\left(\frac{q_X}{1\ {\rm ppmv}}\right)\left(\frac{P_{\rm tot}}{0.5\ {\rm bar}}\right). \tag{41}$$

For CF_4 in a CO_2-dominated atmosphere, this corresponds to $\sim 3 \times 10^{12}$ kg of fluorine per ppmv at $P_{\rm tot} = 0.5$ bar. Thus, ppm-to-tens-of-ppm supergreenhouse scenarios imply that 10^{13}–10^{14}-kg-class feedstock chemistry before losses, photolysis, separations, and plant inefficiencies are included. Even optimistic assessments imply fluorine availability and industrial throughput far beyond near-term capability. For context, recent analysis notes that PFC-based warming may require volatilizing $\sim 10^5$ megatons of fluorine [14], i.e., $\sim 10^{14}$ kg of feedstock element.

(2) *Photochemistry and loss.* Some candidate gases are long-lived, but lifetimes depend on UV flux, atmospheric composition, and catalytic cycles; therefore, the required production rate is set by *replacement* in addition to initial fill.

Because PFC warming does not directly provide buffer gas, it must be coupled to a separate pressure-building pathway to achieve E3/E4.

B. CO_2-H_2 collision-induced absorption

CO_2-H_2 collision-induced absorption can provide substantial greenhouse warming only when embedded in a sufficiently dense CO_2 background atmosphere. In the published early-Mars literature, tens of kelvin of warming generally require percent-level H_2 in $\gtrsim$(0.5–1)-bar-CO_2 atmospheres rather than in present-day thin Mars conditions [31–33]. For terraforming, CIA is attractive because H_2 can in principle be produced *in situ* from water, with O_2 as a coproduct. However, hydrogen escapes efficiently from Mars, so any H_2-assisted warm state eventually requires replenishment. The key architecture question is whether the dominant industrial burden is the initial global H_2 fill or the steady replenishment needed to hold the warm state against escape. Accordingly, a complete system analysis couples

$$\text{required } p_{\rm H_2} \iff \text{required inventory } M_{\rm H_2} \iff \text{required production rate } \dot{M}_{\rm H_2} \iff \text{power } P. \tag{42}$$

As shown below, CO_2-H_2 CIA is best treated as a coupled fill-plus-maintenance pathway: fill dominated during century-to-millennial build-out, but replenishment limited on long hold times.

It is useful to separate the CIA-active background-atmosphere requirement from the H_2 requirement. Define

$$\Pi_{\rm bg}^{\rm CIA} \equiv \frac{P_{\rm bg,avail}}{P_{\rm bg,req}}, \quad P_{\rm bg,req} \sim 0.5\text{–}1\ {\rm bar}, \tag{43}$$

for the class of CO_2-H_2 warm states usually invoked in early-Mars radiative-convective studies. For the representative accessible endogenous CO_2 reference case $P_{\rm bg,avail} \sim 20$ mbar, $\Pi_{\rm bg}^{\rm CIA} \sim 0.02$–$0.04$. Thus, CO_2-H_2 CIA is not a shortcut around the background-atmosphere inventory problem; it is a high radiative-leverage mechanism only after a dense CO_2-rich background has already been supplied, imported, or manufactured.

1. Bracketed mapping: CIA warming implies a required f_{H_2} and an eventual replenishment load

CO_2-H_2 collision-induced absorption scales with the frequency of CO_2-H_2 collisions and therefore requires both (1) a sufficiently dense background atmosphere and (2) an H_2 mixing ratio f_{H_2} at the percent-to-tens-of-percent level in published early-Mars radiative-convective studies [31]. Here and throughout this subsection, f_{H_2} denotes the *mole/number mixing ratio* of H_2 in a well-mixed multicomponent atmosphere. For architecture-level accounting, we treat f_{H_2} as the primary control variable and translate it into a diffusion limited replenishment requirement:

$$\dot{M}_{H_2} \approx 1.2 \times 10^5\, f_{H_2}\ {\rm kg\,s^{-1}}, \tag{44}$$

which gives $\dot{M}_{H_2} \sim 6 \times 10^3$ kg s^{-1} for $f_{H_2} = 0.05$ and $\dot{M}_{H_2} \sim 1.2 \times 10^4$ kg s^{-1} for $f_{H_2} = 0.10$. If replenishment is provided by water electrolysis, the reversible minimum work per kg of H_2 is $\sim 1.2 \times 10^8$ J kg^{-1}, implying a minimum maintenance power

$$P_{H_2,\rm min} \sim (0.7\text{–}1.4)\,{\rm TW}, \quad \text{for} \quad f_{H_2} \sim 0.05\text{–}0.10, \tag{45}$$

before plant inefficiencies and compression/storage. Thus, any CIA pathway that requires f_{H_2} in the few-to-ten percent range couples directly to TW-class power and multi-10^3 kg s^{-1} hydrogen production capability, in addition to the background atmospheric mass inventory constraint required to reach the CIA-effective pressure regime.

2. Closing the loop: CIA warming implies a required p_{H_2} and a replenishment power

For a well-mixed multicomponent atmosphere, the mass of constituent i associated with a target partial pressure p_i is not, in general, species independent. Instead,

$$M_i \simeq \frac{4\pi R_{\rm Mars}^2}{g_{\rm Mars}} \frac{\mu_i}{\bar{\mu}}\, p_i = \mathcal{K}_{\rm Mars}\, \frac{\mu_i}{\bar{\mu}}\, p_i, \tag{46}$$

where μ_i is the molecular mass of species i and $\bar{\mu}$ is the mean molecular mass of the atmospheric mixture. Thus, for an engineered atmosphere with total pressure $P_{\rm tot}$ and hydrogen mole fraction f_{H_2}, one has $p_{H_2} = f_{H_2} P_{\rm tot}$ and

$$M_{H_2} \simeq \mathcal{K}_{\rm Mars}\, \frac{\mu_{H_2}}{\bar{\mu}}\, f_{H_2} P_{\rm tot}, \quad \bar{\mu} \simeq (1 - f_{H_2})\mu_{CO_2} + f_{H_2}\mu_{H_2} \quad \text{(for a } CO_2\text{–}H_2 \text{ atmosphere)}. \tag{47}$$

For the representative case $P_{\rm tot} = 0.5$ bar and $f_{H_2} = 0.05$, $\bar{\mu} \simeq 41.9$ g mol^{-1}, $p_{H_2} = 25$ mbar, and

$$M_{H_2} \approx 4.7 \times 10^{15}\ {\rm kg}, \tag{48}$$

about $21\times$ smaller than the single-component estimate.

The reversible minimum electrolysis work to create this H_2 inventory is then

$$E_{H_2,\rm fill,min} \simeq \varepsilon_{H_2} M_{H_2} \approx (1.2 \times 10^8\ {\rm J\,kg^{-1}})(4.7 \times 10^{15}\ {\rm kg}) \approx 5.6 \times 10^{23}\ {\rm J}. \tag{49}$$

This is still extremely large, but it is no longer comparable to the global O_2-production floor in Eq. (78). Over a 10^3 yr build time, it corresponds to an average reversible fill power of $\sim$18 TW.

Hydrogen escape then sets the ongoing replenishment requirement. Using the diffusion-limited upper bound [Eq. (82)], $f_{H_2} = 0.05$ implies $\dot{M}_{H_2} \sim 6 \times 10^3\,\mathrm{kg\,s^{-1}}$, corresponding to a reversible electrolysis power floor $P_{H_2,\min} \sim 0.7$ TW (Table IX), before real inefficiencies and compression. Thus, even if CIA provides high radiative leverage, it implies (1) a substantial H_2 inventory tied directly to the total pressure and (2) a persistent TW-class industrial metabolism to maintain f_{H_2} against loss.

3. *Fill-dominated versus maintenance-dominated climate agents*

A useful architecture-level distinction is whether the industrial burden is dominated by the initial fill of the active climate agent or by its continuous replenishment against loss. For an agent with target inventory $M^\star$ and steady loss rate $\dot{M}_{\rm loss}$ at that inventory, define

$$\tau_{\rm loss} \equiv \frac{M^\star}{\dot{M}_{\rm loss}}, \quad \Lambda_{\rm maint} \equiv \frac{t_{\rm build}}{\tau_{\rm loss}} = \frac{\dot{M}_{\rm loss}}{M^\star/t_{\rm build}}. \tag{50}$$

If $\Lambda_{\rm maint} \ll 1$, the pathway is fill dominated during buildup; if $\Lambda_{\rm maint} \gtrsim 1$, it is maintenance dominated already during build-out.

For H_2-based CIA warming, take a target atmosphere with total pressure $P_{\rm tot}$ and H_2 mole fraction f_{H_2}, so that $p_{H_2} = f_{H_2} P_{\rm tot}$ and

$$M^\star_{H_2} = \mathcal{K}_{\rm Mars}\,\frac{\mu_{H_2}}{\bar{\mu}}\,f_{H_2} P_{\rm tot}. \tag{51}$$

Using the diffusion-limited escape form $\dot{M}_{{\rm loss}H_2} \simeq 4\pi R^2_{\rm Mars} m_{H_2} \Phi_0 f_{H_2}$, the H_2 loss time is

$$\tau_{H_2,\rm loss} = \frac{M^\star_{H_2}}{\dot{M}_{{\rm loss}H_2}} = \frac{P_{\rm tot}}{g_{\rm Mars}\,\bar{m}\,\Phi_0}, \tag{52}$$

where $\bar{m} = \bar{\mu} m_u$ is the mean molecular mass in kilograms. For the representative $P_{\rm tot} = 0.5$ bar, $f_{H_2} = 0.05$ case, this gives $\tau_{H_2,\rm loss} \approx 2.5 \times 10^4$ yr, $\dot{M}_{{\rm loss}H_2} \approx 6 \times 10^3\,\mathrm{kg\,s^{-1}}$, and therefore

$$\Lambda_{\rm maint} \approx \begin{cases} 4 \times 10^{-3}, & t_{\rm build} = 10^2\ \mathrm{yr}, \\ 4 \times 10^{-2}, & t_{\rm build} = 10^3\ \mathrm{yr}. \end{cases} \tag{53}$$

The fill-versus-maintenance crossover for this representative case is shown in Fig. 2. Therefore, for century-to-millennial global build schedules, a CIA pathway is *fill dominated* rather than maintenance dominated: The dominant burden is creating the required global H_2 inventory, which can range from multimillibars to multikilopascals depending on $P_{\rm tot}$ and f_{H_2}. Maintenance remains important for long hold times, but it is not the primary scaling bottleneck during build-out.

In contrast, for an aerosol pathway with residence time $\tau_p \sim$ 30–100 days,

$$\Lambda^{(p)}_{\rm maint} = \frac{t_{\rm build}}{\tau_p} \sim 10^2\text{–}10^4 \quad (t_{\rm build} = 10^2\text{–}10^3\ \mathrm{yr}), \tag{54}$$

so aerosol warming is maintenance dominated from the outset. This distinction sharpens the mechanism taxonomy: Aerosols are truly replenishment limited, whereas global H_2 CIA is typically inventory/fill limited on civilization timescales and only secondarily maintenance limited.

4. *Electrolysis-based* $\mathbf{H_2}$ *fill is intrinsically coupled to* $\mathbf{O_2}$ *management*

If the required H_2 inventory is produced from water electrolysis, O_2 is generated stoichiometrically as a coproduct:

$$2\,H_2O \rightarrow 2\,H_2 + O_2. \tag{55}$$

By mass, $M_{O_2,\rm co} = 8 M_{H_2}$, but the corresponding partial-pressure relation is controlled by *mole number*, not by the 8:1 mass ratio. If both gases are retained in the atmosphere, one has $n_{O_2} = n_{H_2}/2$,

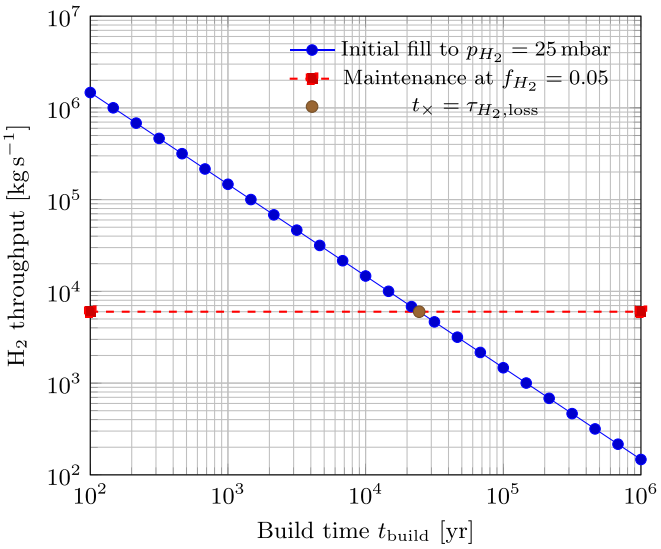


FIG. 2. Fill-vs-maintenance crossover for a representative global H_2 CIA pathway. The example assumes that $P_{\mathrm{tot}} = 0.5$ bar and $f_{H_2} = 0.05$, so that $p_{H_2} = 25$ mbar and $M^\star_{H_2} \approx 4.7 \times 10^{15}$ kg once the multicomponent correction $\mu_{H_2}/\bar{\mu}$ is included. The horizontal dashed line is the diffusion-limited replenishment rate $\dot{M}_{\mathrm{loss}H_2} \approx 6 \times 10^3$ kg s^{-1}; the solid line is the fill rate required to build the inventory over time t_{build}. The crossover occurs at $t_\times \simeq \tau_{H_2,\mathrm{loss}} \approx 2.5 \times 10^4$ yr. Century-to-millennial build schedules lie well to the left of this crossover, so global H_2 CIA is still fill dominated during build-out even though long-term replenishment remains a terawatt-class obligation.

implying to first approximation

$$p_{O_2,\mathrm{co}} \sim \tfrac{1}{2} p_{H_2}, \tag{56}$$

prior to accounting for total-pressure changes and chemical sinks. For the representative $p_{H_2} = 25$ mbar case, the retained coproduct is therefore of order 12.5 mbar of O_2, not a human-breathable O_2 partial pressure. The system's point remains important: Large-scale H_2 generation from water is inseparable from large-scale O_2 management and associated safety, sink-filling, storage, export, and compositional-control issues.

CO_2-H_2 CIA is analyzed here because its radiative physics is well studied and H_2 can in principle be produced from water *in situ*. This should not be read as prioritizing CIA over all imported-greenhouse strategies: If the dense CO_2 background itself must be supplied exogenously, then CH_4, NH_3, or other gases may offer higher radiative leverage per imported kilogram, subject to lifetime, condensation, toxicity, and photochemical constraints.

C. Engineered IR-active aerosols and nanoparticles

A recently quantified proposal is to introduce engineered *IR-active* particles that interact strongly with upwelling thermal radiation while incurring a comparatively smaller visible penalty [14]. The attraction of this pathway is high radiative leverage at low added atmospheric mass; the cost is that the particles must remain aloft and dispersed globally. Recent 3D modeling further suggests that radiative-dynamical feedbacks can assist lofting and global spread for suitably designed particles [15].

As an illustrative conversion, [14] discuss sustained particle injection of order tens of liters per second. A volumetric injection rate $\dot{V} \sim 301\,\mathrm{s}^{-1}$ corresponds to $\dot{V} \approx 0.03$ m^3 s^{-1}. For a represen-

tative particle bulk density $\rho \sim 3 \times 10^3\,\mathrm{kg\,m^{-3}}$,

$$\dot{M}_\mathrm{p} \sim \rho \dot{V} \sim 90\,\mathrm{kg\,s^{-1}} \sim 2.8 \times 10^9\,\mathrm{kg\,yr^{-1}}. \tag{57}$$

This is *small* compared to atmospheric mass requirements (Sec. III), but large as an industrial mass flow, comparable to terrestrial mining/processing streams. The power required depends on mining, comminution, transport, and lofting energy, which are architecture specific. The maintained global mean atmospheric column is then

$$\Sigma_p \simeq \frac{\dot{M}_p\,\tau_p}{4\pi R^2_\mathrm{Mars}}. \tag{58}$$

For $\tau_p \sim 30$–100 days, $\Sigma_p \sim 1.6$–$5.4\,\mathrm{mg\,m^{-2}}$; for $\tau_p \sim 10$ yr, the same injection sustains $\Sigma_p \sim 0.2\,\mathrm{g\,m^{-2}}$, comparable to the maintained columns invoked for tens of kelvin of warming [14]. This provides a direct bridge between industrial injection rate and the "$\mathrm{mg\,m^{-2}}$ class" entries in Table VIII.

Particle-induced warming is a *maintenance* mechanism. If particles have an atmospheric residence time τ_p set by sedimentation and scavenging, maintaining a global column mass Σ_p requires a source flux $\dot{\Sigma}_p \approx \Sigma_p/\tau_p$. At Mars, dust residence times can range from days (local storms) to months (global events), and engineered particles must be designed to balance (1) optical efficiency, (2) coagulation resistance, and (3) acceptable health and environmental impacts. Thus, the required injection rates scale inversely with achievable residence time.

Detailed climate outcomes depend on particle optical properties and atmospheric dynamics, but the appropriate figure of merit is not shortwave absorption. Instead, the key design objective is a particle population with strong thermal-IR interaction, limited solar penalty, and a long effective atmospheric residence time. For architecture-level bookkeeping, the practical conclusion is simple: Particle pathways are maintenance limited, with industrial requirements set primarily by the pair (τ_p, Σ_p), and therefore by lifetime extension mechanisms such as slow sedimentation, photophoretic support, and radiative self-lofting [15]. The literature values should therefore be read as highly lifetime sensitive rather than as one-size-fits-all injection prescriptions.

For carbon-based particle concepts, the feedstock itself need not come from mineral mining: In principle, it can be derived from atmospheric CO_2 processing, which shifts the industrial burden from excavation toward gas processing, particle fabrication, and lofting. This does not remove the maintenance character of the pathway, but it can materially change the supply-chain architecture [15].

D. Orbital mirrors and albedo modification

Directly increasing absorbed-solar flux is conceptually simple but scale limited. Solar-reflector and sail-based climate-engineering concepts for Mars have a substantial prior literature, including displaced non-Keplerian reflector orbits and mass-produced sail architectures launched from Earth's orbit [34,35]. The scaling developed here is therefore not intended to claim novelty for the mirror concept itself, but to place such approaches on the same forcing/area architecture axis as the other mechanisms.

Suppose one seeks an average TOA forcing ΔF_TOA by adding reflected sunlight. The additional absorbed power is

$$\Delta P = 4\pi R^2_\mathrm{Mars}\,\Delta F_\mathrm{TOA}. \tag{59}$$

A mirror of area A_m at Mars orbit intercepts solar power $S_\mathrm{Mars}A_m$. If an overall efficiency η_m accounts for reflectivity, pointing, and geometric losses, then

$$A_m \simeq \frac{4\pi R^2_\mathrm{Mars}\Delta F_\mathrm{TOA}}{\eta_m C_m S_\mathrm{Mars}}, \tag{60}$$

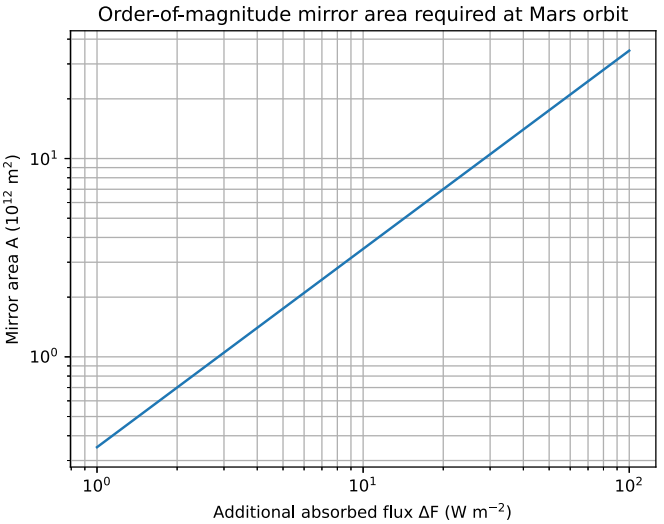


FIG. 3. Order-of-magnitude mirror area required to supply a global mean additional absorbed flux ΔF_{TOA} at Mars orbit, using Eq. (60) with $\eta_m = 0.7$ and $S_{\mathrm{Mars}} = 589\,\mathrm{W\,m^{-2}}$ [24].

where $C_m \leqslant 1$ is a geometry/duty-cycle factor collecting orbital geometry, phase angle, illumination pattern, specular versus diffuse behavior, target selection, eclipses, pointing losses, and time-dependent degradation. The numerical values below use $C_m = 1$ and are therefore aperture lower bounds. Combining the definition of global mean forcing with intercepted solar flux yields Eq. (60). For a quick estimate, the mirror area scaling is given by

$$A_m \approx 7.0 \times 10^{12} \left(\frac{\Delta F_{\mathrm{TOA}}}{20\,\mathrm{W\,m^{-2}}}\right)\left(\frac{1}{C_m}\right)\left(\frac{0.7}{\eta_m}\right)\left(\frac{589\,\mathrm{W\,m^{-2}}}{S_{\mathrm{Mars}}}\right)\mathrm{m}^2. \tag{61}$$

Figure 3 shows this scaling for $\eta_m = 0.7$. Even a modest global forcing $\Delta F_{\mathrm{TOA}} = 20\,\mathrm{W\,m^{-2}}$ implies $A_m \sim 7 \times 10^{12}\,\mathrm{m^2}$ ($\sim 7 \times 10^6\,\mathrm{km^2}$), a continent-scale structure.

Albedo reduction is a related lever, with historical precedent in Sagan's polar-darkening proposal for Mars climate modification [36]. Because absorbed solar is $S_{\mathrm{Mars}}(1 - A_{\mathrm{B}})/4$, a change ΔA_{B} yields

$$\Delta F_{\mathrm{TOA}} \approx -\frac{S_{\mathrm{Mars}}}{4}\,\Delta A_{\mathrm{B}}. \tag{62}$$

Reducing A_{B} by 0.05 provides only $\sim 7\,\mathrm{W\,m^{-2}}$, illustrating that large temperature changes require either very large albedo changes, very large mirrors, or strong greenhouse feedbacks.

1. Mirror mass scaling

Mirror area is not the only constraint: Total mirror mass scales as

$$M_m \sim \sigma_m A_m, \tag{63}$$

where σ_m is areal density. For $A_m \sim 7 \times 10^{12}\,\mathrm{m^2}$ (Fig. 3) and $\sigma_m \sim 1$–$10\,\mathrm{g\,m^{-2}}$, the mirror mass is $M_m \sim 10^{10}$–$10^{11}\,\mathrm{kg}$, comparable to a large terrestrial megaproject and requiring either very large launch cadence, modular in-space assembly, or in-space manufacturing for credibility.

2. Mirror scale for melt-class forcing

The $\sim$60 K "melt-class" deficit corresponds to $\Delta F_{\mathrm{TOA}} \approx 191\,\mathrm{W\,m^{-2}}$ for a direct T_e increase from 210 to 270 K [Eq. (34)]. Using Eq. (60), this implies a mirror area larger by a factor of $\sim 191/20 \approx 9.6$ than the $\Delta F_{\mathrm{TOA}} = 20\,\mathrm{W\,m^{-2}}$ example, i.e., $A_m \sim 7 \times 10^{13}\,\mathrm{m^2} \sim 7 \times 10^7\,\mathrm{km^2}$,

before additional geometric and control losses. This highlights why mirrors are best viewed as a high-authority but extreme-scale reflector architecture. The burden need not be a single monolithic structure—distributed reflector or sail constellations are possible—but the aggregate area, mass, deployment, and control requirements remain extreme at global-forcing levels.

As we have seen, Sec. VI has shown that substantial warming can, in principle, be achieved with relatively small added atmospheric mass, but typically at the cost of either large initial fills plus long-term replenishment (CO_2-H_2 CIA), sustained maintenance (aerosol injection), or extreme aggregate area and operations (mirrors). Importantly, warming alone does not deliver a breathable atmosphere: E4 is dominated by O_2 and buffer-gas inventories and their associated energy and throughput requirements. Section VII quantifies these composition-driven requirements.

3. Operational realism: reflector control, stationkeeping, degradation, and replacement

Area scaling alone is not a sufficient feasibility test for Mars reflectors. A viable reflector architecture must also be delivered to Mars space, deployed at large area, pointed continuously, kept in the required orbit family, monitored, and partially replaced as membranes and interfaces degrade over hold times of 10^2–10^3 yr. Solar sails are relevant to this discussion because they can provide propellant-free transfer and non-Keplerian control authority for *small* sailcraft [34,37,38]; they do *not* by themselves close the aggregate area and operations problem for climate-scale Mars reflectors.

At Mars, the photon pressure on a perfectly reflecting sheet normal to the Sun is

$$p_{\mathrm{rad,M}} \simeq \frac{2S_{\mathrm{Mars}}}{c} \approx 3.93 \times 10^{-6}\,\mathrm{N\,m^{-2}}, \tag{64}$$

using $S_{\mathrm{Mars}} \approx 589\,\mathrm{W\,m^{-2}}$ (Table II). The corresponding characteristic acceleration for a sail or reflector of areal mass density σ is

$$a_{\mathrm{SRP,M}} \simeq \frac{p_{\mathrm{rad,M}}}{\sigma} \approx 3.93 \times 10^{-4}\left(\frac{10\,\mathrm{g\,m^{-2}}}{\sigma}\right)\mathrm{m\,s^{-2}}. \tag{65}$$

For comparison, NASA's ACS3 technology demonstration has sail area $\sim$80 m^2, total spacecraft mass $\sim$16 kg, and effective characteristic acceleration $a_c \simeq 0.045\,\mathrm{mm\,s^{-2}}$ at 1 AU [39,40]; this corresponds to a total system areal density $\sigma \sim 0.20\,\mathrm{kg\,m^{-2}}$, for which Eq. (65) gives $a_{\mathrm{SRP,M}} \sim 2 \times 10^{-5}\,\mathrm{m\,s^{-2}}$ at Mars. NASA's current solar-sail roadmap discusses controlled sails with areas $>2.5 \times 10^3$ m^2 (Generation 2) and $>10^4$ m^2 (Generation 3a), still many orders of magnitude below the 10^{12}–10^{13} m^2 areas implied by Eqs. (34)–(60) [41].

It is important to separate aggregate-area feasibility from local force balance. For a reflecting sail with areal density σ, the solar-radiation-pressure (SRP) acceleration at Mars is $a_{\mathrm{SRP,M}} = p_{\mathrm{rad,M}}/\sigma$, whereas Mars gravity decreases as μ_{Mars}/r^2. A simple radial force-balance distance is therefore

$$r_{\mathrm{bal}} \simeq \left(\frac{\mu_{\mathrm{Mars}}\sigma}{p_{\mathrm{rad,M}}}\right)^{1/2} \simeq 3.3 \times 10^8\,\mathrm{m}\left(\frac{\sigma}{10\,\mathrm{g\,m^{-2}}}\right)^{1/2}, \tag{66}$$

where r_{bal} is measured from Mars' center. Thus, force balance is not the fundamental objection to reflector concepts; SRP-assisted and displaced non-Keplerian reflector geometries are physically plausible for sufficiently low areal density [34]. The hard constraint is the simultaneous closure of aggregate area, module count, pointing, orbit control, lifetime, and replacement over 10^2–10^3 yr.

Near-surface body-fixed hovering is a different and more restrictive problem than displaced Mars-centered reflector orbits; it is not the architecture assumed in the area and operations estimates below.

The aggregate operations burden is also extreme. For the representative forcing levels already used in this paper, $A_m \sim 7 \times 10^{12}$ m^2 for $\Delta F_{\mathrm{TOA}} = 20\,\mathrm{W\,m^{-2}}$ and $A_m \sim 7 \times 10^{13}$ m^2 for melt-class

forcing. If one imagines 80 m^2 ACS3-class units, the required module count is

$$N_{\mathrm{ACS3}} \sim \frac{A_m}{80\ \mathrm{m}^2} \sim 9 \times 10^{10} \text{ to } 9 \times 10^{11}, \tag{67}$$

while even 1×10^4-m^2-class future sail/reflector units imply

$$N_{10^4} \sim \frac{A_m}{10^4\ \mathrm{m}^2} \sim 7 \times 10^8 \text{ to } 7 \times 10^9. \tag{68}$$

At $\sigma = 10\,\mathrm{g\,m^{-2}}$, the associated reflector mass is $M_m \sim 7 \times 10^{10}$ to 7×10^{11} kg, consistent with Eq. (63). If an illustrative module lifetime of only $\tau_{\mathrm{mod}} = 30$ yr is assumed, the steady replacement flow is

$$\dot{M}_{m,\mathrm{repl}} \sim \frac{M_m}{\tau_{\mathrm{mod}}} \sim 7 \times 10^1 \text{ to } 7 \times 10^2\ \mathrm{kg\,s^{-1}} \quad (\sigma = 10\,\mathrm{g\,m^{-2}}), \tag{69}$$

before accounting for failed deployment, attitude-control hardware, metrology, or disposal/reconfiguration.

A convenient operations metric is the required *area replacement rate*:

$$\dot{A}_{m,\mathrm{repl}} \sim \frac{A_m}{\tau_{\mathrm{mod}}} \sim 7 \times 10^3 \text{ to } 7 \times 10^4\ \mathrm{m^2\,s^{-1}} \quad (\tau_{\mathrm{mod}} = 30\,\mathrm{yr}), \tag{70}$$

for the representative forcing range already considered in this paper. In module units, the same replacement burden corresponds to

$$\dot{N}_{10^4} \sim \frac{N_{10^4}}{\tau_{\mathrm{mod}}} \sim 0.7\text{–}7\ \mathrm{s^{-1}}, \quad \dot{N}_{\mathrm{ACS3}} \sim \frac{N_{\mathrm{ACS3}}}{\tau_{\mathrm{mod}}} \sim 10^2 \text{ to } 10^3\ \mathrm{s^{-1}}, \tag{71}$$

for 1×10^4-m^2-class and ACS3-class units, respectively. This makes explicit that the reflector problem is not merely a launch or transfer problem: It is a long-duration manufacturing, deployment, metrology, control, and retirement/replacement problem at rates that remain extreme even if the individual modules are assumed to be lightweight.

This local-versus-global distinction is also reflected in recent roadmap analyses: Orbiting reflectors may be relevant for warming specific sites or volatile reservoirs, but their local-base examples already require 100s of km^2 of reflector area, far below yet qualitatively consistent with the extreme aggregate-area burden implied here for global forcing [17].

Finally, long-duration reflector operations are limited by environmental durability as well as by orbital control. NASA solar-sail materials work emphasizes that mission success depends on membrane/interface durability over packaging, deployment, and operations; radiation-driven adhesive degradation and associated reflectivity loss are active engineering concerns rather than solved issues [42]. The correct interpretation is therefore not that Mars reflectors are ruled out by physics, but that they belong to a *far-horizon space-manufacturing and operations* regime: Solar sails may help with transfer and trim of modular reflectors, yet the dominant bottlenecks remain aggregate area, mass fabrication, deployment count, pointing/control, durability, and replacement over centuries to millennia.

VII. OXYGENATION AND BREATHABLE END POINTS

A. O_2 mass required for breathable partial pressures

A breathable atmosphere requires substantial O_2 partial pressure. For a well-mixed multicomponent atmosphere, the mass of O_2 corresponding to a target surface partial pressure p_{O_2} is

$$M_{O_2} \simeq \mathcal{K}_{\mathrm{Mars}} \frac{\mu_{O_2}}{\bar{\mu}}\, p_{O_2}. \tag{72}$$

For the representative buffered case $p_{O_2} = 21$ kPa, $p_{N_2} = 50$ kPa, and $\bar{\mu} \simeq 29.2\ \mathrm{g\,mol^{-1}}$,

$$M_{O_2} \approx 9.0 \times 10^{17}\ \mathrm{kg}\left(\frac{p_{O_2}}{21\ \mathrm{kPa}}\right)\left(\frac{\mu_{O_2}/\bar{\mu}}{32/29.2}\right). \tag{73}$$

If instead O_2 is treated as a pure single-component atmosphere, the corresponding pressure-equivalent inventory is 8.2×10^{17} kg, comparable to the total mass of a 0.2 bar atmosphere. The distinction is only ∼10%, but Eq. (72) is the internally consistent well-mixed-atmosphere form.

1. Covered-area ("worldhouse") crossover

A sealed enclosure trades the open-atmosphere hydrostatic mass requirement into a finite-volume gas requirement. For a covered area A_{cov} with mean interior height H, interior pressure P_{in}, mean gas molecular mass $\bar{\mu}$, and temperature T, the enclosed-gas mass is

$$M_{encl} \simeq A_{cov} H\,\frac{\bar{\mu} P_{in}}{RT} = A_{cov}\,\frac{H}{H_s}\,\frac{P_{in}}{g_{Mars}}, \quad H_s \equiv \frac{RT}{\bar{\mu} g_{Mars}}. \tag{74}$$

Thus, the ratio of enclosed-gas mass to open-atmosphere hydrostatic mass over the same area is simply H/H_s. For a representative breathable enclosure with $H = 3$ m, $P_{in} = 50$ kPa, $T = 250$ K, and $\bar{\mu} \approx 29\ \mathrm{g\,mol^{-1}}$, Mars gives $H_s \approx 19$ km and hence $H/H_s \approx 1.6 \times 10^{-4}$. If applied over the full planetary surface area, this corresponds to

$$M_{encl,global} \approx 3.0 \times 10^{14}\ \mathrm{kg}, \tag{75}$$

about 6.4×10^3 smaller than an open global 50 kPa atmosphere. This does not make a global worldhouse easy—structural mass, support spacing, puncture tolerance, maintenance, and thermal control become the dominant burdens—but it reinforces the paper's broader conclusion that covered-area strategies can beat open-atmosphere strategies by trading atmospheric mass into engineered area and support infrastructure.

The atmospheric-mass saving is purchased by structural and maintenance load. A covered region must carry or anchor a pressure load

$$\frac{F_{anchor}}{A} \simeq P_{in} - P_{amb} \approx 50\ \mathrm{kN\,m^{-2}}\left(\frac{P_{in} - P_{amb}}{50\ \mathrm{kPa}}\right), \tag{76}$$

before dust loading, creep, puncture tolerance, radiation aging, thermal cycling, and repair logistics are included. Worldhouse concepts therefore trade gas inventory into area, supports, anchoring, and maintenance rather than eliminating the engineering burden.

B. Thermodynamic minimum energy to produce O_2 from water

Producing O_2 by electrolyzing water,

$$2\,H_2O \rightarrow 2\,H_2 + O_2, \tag{77}$$

requires a minimum Gibbs free energy input ΔG° per mole of O_2. Using standard thermochemical data [43], the reversible minimum at 298 K is $\Delta G^\circ \approx 474\ \mathrm{kJ\,mol^{-1}}$ per mole of O_2 (≈14.8 MJ kg^{-1} of O_2). Therefore, the idealized minimum energy to produce Eq. (73) is

$$E_{min} \approx \left(\frac{M_{O_2}}{0.032\ \mathrm{kg\,mol^{-1}}}\right)\Delta G^\circ \approx 1.3 \times 10^{25}\ \mathrm{J}. \tag{78}$$

Real systems require larger energy due to overpotentials, compression/liquefaction, plant losses, and ancillary mining/processing. Even the theoretical minimum corresponds to an average power of ∼420 TW sustained for 1000 yr, illustrating that E4 end points are *energy-dominated*.

Stoichiometrically, producing 1 kg of O_2 from water requires $36/32 = 1.125$ kg of H_2O processed. Thus, producing $M_{O_2} \approx 9.0 \times 10^{17}$ kg implies processing $M_{H_2O} \approx 1.0 \times 10^{18}$ kg of water, equivalent to a global water layer of thickness $\sim M_{H_2O}/(\rho_w 4\pi R_{Mars}^2) \approx 7$ m for $\rho_w \approx 1000\ \mathrm{kg\,m^{-3}}$.

Biologically mediated O_2 accumulation is treated separately in Sec. X using areal productivity and burial/export efficiency. The present subsection defines the abiotic electrochemical lower bound, not a limit on biological bootstrapping.

C. MOXIE as a scale anchor

MOXIE demonstrated oxygen production from atmospheric CO_2 on Mars [44,45]. Such demonstrations are crucial technology anchors, but MOXIE is not invoked here as a direct terraforming route; rather, it serves as a scale benchmark for the gulf between a successful Mars oxygen-production demonstration and the 10^{18}-kg-scale oxygen inventories relevant to atmospheric transformation. A breathable-atmosphere inventory $M_{O_2} \sim 10^{18}\,\mathrm{kg}$ is $\sim 10^{19}$ times larger than an $\mathcal{O}(10^2)$-g-class technology demonstration output, even before accounting for compression, storage, distribution, and O_2 sinks into the regolith and crust. Therefore, oxygenation is not a "scale-up by a factor of ten" problem; it is an orders-of-magnitude industrial-civilization problem.

D. O_2 sinks: oxidation is a kinetic and operational constraint

A breathable end point requires not only producing M_{O_2} (Sec. VII), but also overcoming sinks as O_2 oxidizes reduced minerals. An upper-envelope *capacity estimate* is obtained by assuming an accessible oxidizable layer of thickness d and density ρ with representative FeO mass fraction w_{FeO} (e.g., FeO $\sim$18 wt % in Mars-relevant silicates [46,47]).

If FeO is oxidized to Fe_2O_3 via $4\mathrm{FeO} + O_2 \rightarrow 2\mathrm{Fe_2O_3}$, the required moles of O_2 per mole of FeO is $1/4$. The corresponding O_2 mass sink capacity of a global layer is approximately

$$M_{O_2,\mathrm{sink}}(t) \sim \frac{1}{4}\left(\frac{w_{\mathrm{FeO}}\, 4\pi R_{\mathrm{Mars}}^2 \rho\, d_{\mathrm{eff}}(t)}{\mu_{\mathrm{FeO}}}\right)\mu_{O_2}, \tag{79}$$

where μ_{FeO} and μ_{O_2} are molar masses. Using shallow regolith rather than crustal-depth values,

$$M_{O_2,\mathrm{sink}} \lesssim 1.4 \times 10^{16}\,\mathrm{kg}\left(\frac{w_{\mathrm{FeO}}}{0.18}\right)\left(\frac{\rho}{1000\,\mathrm{kg\,m^{-3}}}\right)\left(\frac{d_{\mathrm{eff}}}{5\,\mathrm{m}}\right). \tag{80}$$

This is well below the $M_{O_2} \simeq 9 \times 10^{17}\,\mathrm{kg}$ inventory associated with the representative E4 O_2 target. The appropriate conclusion is therefore not that a breathable atmosphere must disappear stoichiometrically into the soil. Rather, oxidation is an uncertain kinetic and operational sink: Uptake depends on permeability, reaction rates, already oxidized phases, surface gardening, and the amount of fresh reduced material exposed by excavation and construction. E4 oxygenation must therefore include sink management—passivation, sealing/vitrification where appropriate, and controlled excavation.

VIII. ATMOSPHERIC LOSS, RETENTION, AND CONTROL

A terraformed Mars must retain its engineered atmosphere against escape to space and against surface sequestration (carbonates, adsorption, polar condensation).

A. Escape to space: why it matters more for H_2 than for CO_2?

MAVEN observations quantify present-day atmospheric loss and its dependence on solar conditions [48]. Mars Express ASPERA-3 measurements provide a complementary constraint on heavy-ion escape and its extrapolation through time [49]. For heavy species (CO_2, N_2), bar-level atmospheres have enormous mass and are likely to be lost on geological timescales even without a global magnetic field, whereas for hydrogen the relevant benchmark escape timescale can be much shorter. In this paper, we use the diffusion-limited form as a conservative upper-bound maintenance benchmark for H_2-rich warm states; actual escape may be smaller if other limits intervene. Thus,

H_2-based warming (Sec. VI B) implies a persistent replenishment requirement on long hold times, even if the initial global build-out is fill dominated for $t_{\rm build} \ll \tau_{\rm H_2,loss}$.

B. Artificial magnetic shielding

Artificial magnetic shielding may reduce ion escape and alter the radiation environment [50], but it does not by itself remove the dominant neutral-H diffusion-limited escape benchmark used here for CO_2-H_2 maintenance. We therefore treat magnetic shielding as a secondary retention/radiation-environment concept rather than as a primary solution to the H_2 maintenance burden.

C. Hydrogen escape sets a long-duration replenishment benchmark for CO_2-H_2 warming

A key distinction between CO_2-mass-based warming and CO_2-H_2 CIA warming is that hydrogen can be lost rapidly, implying a persistent replenishment requirement. A standard reference ceiling for hydrogen escape through a heavier background atmosphere is the diffusion-limited flux, often expressed in terms of the hydrogen atom flux ("H equivalents") (e.g., [51,52]):

$$\Phi_{\rm DL,H} \approx \Phi_0\, f_{\rm H}, \tag{81}$$

where $f_{\rm H}$ is the total hydrogen mixing ratio expressed as H atoms (summing over all H-bearing species). If hydrogen is supplied primarily as H_2, then $f_{\rm H} \approx 2f_{\rm H_2}$ and the corresponding H_2 molecular flux is $\Phi_{\rm DL,H_2} \approx \frac{1}{2}\Phi_{\rm DL,H} \approx \Phi_0 f_{\rm H_2}$. The corresponding global H_2 mass loss rate is then

$$\dot{M}_{\rm H_2} \approx 4\pi R_{\rm Mars}^2\, m_{\rm H_2}\, \Phi_{\rm DL,H_2} \approx 1.2\times 10^5 \left(\frac{\Phi_0}{2.5\times 10^{13}\ {\rm cm^{-2}\,s^{-1}}}\right) f_{\rm H_2}\ {\rm kg\,s^{-1}}, \tag{82}$$

where $m_{\rm H_2}$ is the molecular mass of H_2 and we used Mars' surface area for scaling.

This expression should be interpreted as a diffusion-limited upper-bound benchmark, not as a guaranteed loss rate. At percent-level H_2, the actual escape rate depends on photochemical production of H-bearing species, homopause structure, extreme-ultraviolet (EUV) energy input, exobase geometry, and whether the escape regime remains diffusion limited or becomes energy limited or partly hydrodynamic. We therefore use Eq. (82) to size a conservative maintenance-capability envelope; detailed escape modeling is required for a specific atmosphere.

Even modest $f_{\rm H_2}$ can therefore imply $\dot{M}_{\rm H_2} \sim 10^3$ to $10^4\,{\rm kg\,s^{-1}}$ class replenishment. If H_2 is produced by water electrolysis, the reversible specific energy is $\sim 1.2\times 10^8\,{\rm J\,kg^{-1}}$ of H_2, implying a minimum maintenance power

$$P_{\rm H_2,min} \sim 0.1\text{–}1\ {\rm TW} \quad \text{for} \quad \dot{M}_{\rm H_2} \sim 10^3 \text{ to } 10^4\ {\rm kg\,s^{-1}}, \tag{83}$$

before compression, storage, and plant inefficiencies. Therefore CO_2-H_2 warming should be analyzed as a coupled fill-plus-maintenance system: Section VI B 3 shows that the initial global H_2 fill dominates for $t_{\rm build} \ll \tau_{\rm H_2,loss}$, while the equations in this subsection quantify the diffusion-limited replenishment benchmark relevant to long-duration hold of the warm state.

D. CO_2 condensation and atmospheric collapse: a stability constraint on CO_2-centric pathways

CO_2-centric terraforming pathways must contend with the possibility of atmospheric collapse, i.e., condensation of CO_2 into permanent polar deposits that reduce P_s and weaken greenhouse warming. Three-dimensional climate simulations of early Mars with 0.1–7 bar CO_2 atmospheres show that CO_2 ice clouds, obliquity, dust loading, and surface properties can lead to climates where CO_2 condensation and cold trapping limit warming and can induce collapselike behavior under some conditions [53].

In a first-order control sense, CO_2 stability requires maintaining polar temperatures above the CO_2 condensation threshold for the prevailing $p_{\rm CO_2}$. Thus, even if CO_2 is available in principle, the engineered climate may require active control authority (e.g., directed insolation, albedo man-

agement, or supplemental absorbers) to avoid pressure-loss hysteresis and to keep the system on a warm branch.

E. Regional enclosed-volume scaling

A laterally bounded region cannot maintain overpressure without a roof, cover, membrane, or transition to a globally pressurized atmosphere; otherwise, gas escapes over the boundary. The relevant regional paraterraforming inventory is therefore the enclosed-volume mass, not a full hydrostatic atmospheric column:

$$M_{\rm encl} \simeq A_{\rm reg} H \frac{\bar{\mu} P_{\rm in}}{R_g T} = A_{\rm reg} \frac{H}{H_s} \frac{P_{\rm in}}{g_{\rm Mars}}, \quad H_s \equiv \frac{R_g T}{\bar{\mu} g_{\rm Mars}}. \tag{84}$$

Relative to an open hydrostatic column over the same area, the gas inventory is reduced by H/H_s. For $H \simeq 3\,\mathrm{m}$, $T \simeq 250\,\mathrm{K}$, and $\bar{\mu} \simeq 29\,\mathrm{g\,mol^{-1}}$, $H_s \simeq 19\,\mathrm{km}$ and $H/H_s \simeq 1.6 \times 10^{-4}$. Thus, regional habitability trades global gas inventory into roof/membrane area, pressure load, anchoring, puncture tolerance, dust management, and repair logistics. This is the quantitative reason covered-area strategies scale far more favorably than open-atmosphere terraforming: Gas fill is no longer the leading mass term once the enclosed height is small compared with the atmospheric scale height.

F. A minimal state-space model for Mars as a controlled climate-industrial system

The preceding sections motivate treating terraforming as a coupled dynamical system rather than as a sequence of isolated static budgets. To avoid redundancy, we treat the partial pressures as states and the total surface pressure as an algebraic output. A minimal slow state vector may therefore be written as

$$\mathbf{x} \equiv \begin{bmatrix} p_{\rm CO_2} & p_{\rm H_2} & p_{\rm O_2} & p_{\rm N_2} & M_p & \alpha_{\rm pol} & T_s \end{bmatrix}^{\mathsf{T}}, \tag{85}$$

where M_p is the atmospheric mass of engineered particles and $\alpha_{\rm pol}$ is the fraction of the accessible CO_2 reservoir that is cold trapped in permanent polar deposits. A corresponding control vector is

$$\mathbf{u} \equiv \begin{bmatrix} \dot{M}_{\rm CO_2} & \dot{M}_{\rm H_2} & \dot{M}_{\rm O_2} & \dot{M}_{\rm N_2} & \dot{M}_{p,\rm inj} & A_m \end{bmatrix}^{\mathsf{T}}. \tag{86}$$

The total surface pressure is then treated diagnostically as

$$P_s = p_{\rm CO_2} + p_{\rm H_2} + p_{\rm O_2} + p_{\rm N_2} + p_{\rm other}, \tag{87}$$

where $p_{\rm other}$ denotes any neglected minor constituents and may be set to zero in the minimal closure.

A first-order architecture model is then

$$\dot{p}_{\rm CO_2} = \frac{g_{\rm Mars}}{4\pi R_{\rm Mars}^2}\left(\dot{M}^{\rm pe}_{\rm CO_2} - \dot{M}^{\rm pe}_{\rm CO_2,loss} - \dot{M}^{\rm pe}_{\rm CO_2,sink,chem} - M_{\rm CO_2,acc}\,\dot{\alpha}_{\rm pol}\right), \tag{88}$$

$$\dot{p}_{\rm H_2} = \frac{g_{\rm Mars}}{4\pi R_{\rm Mars}^2}\dot{M}^{\rm pe}_{\rm H_2} - \frac{p_{\rm H_2}}{\tau_{\rm H_2,loss}}. \tag{89}$$

$$\dot{p}_{\rm O_2} = \frac{g_{\rm Mars}}{4\pi R_{\rm Mars}^2}\dot{M}^{\rm pe}_{\rm O_2} - \frac{p_{\rm O_2}}{\tau_{\rm O_2,sink}}, \tag{90}$$

$$\dot{p}_{\rm N_2} = \frac{g_{\rm Mars}}{4\pi R_{\rm Mars}^2}\left(\dot{M}^{\rm pe}_{\rm N_2} - \dot{M}^{\rm pe}_{\rm N_2,loss} - \dot{M}^{\rm pe}_{\rm N_2,sink}\right), \tag{91}$$

$$\dot{M}_p = \dot{M}_{p,\rm inj} - \frac{M_p}{\tau_p}, \tag{92}$$

$$\dot{\alpha}_{\rm pol} = \mathcal{C}\left(T_{\rm pol}, p_{\rm CO_2}\right) - \mathcal{S}\left(A_m, \dot{M}_{p,\rm inj}, M_{\rm H_2}, \dots\right), \tag{93}$$

$$C_{\rm eff}\dot{T}_s = F_{\rm in}(A_B, A_m, t) + F_{\rm GH}\left(p_{\rm CO_2}, p_{\rm H_2}, M_p, \dots\right) - \mathrm{OLR}(T_s, \mathbf{x}), \tag{94}$$

TABLE X. Qualitative sign structure of the steady-state gain matrix $G_{kj} = \partial y_k/\partial u_j$ around a warm operating point, with P_s treated as an algebraic output via Eq. (87) rather than as an independent state. Parentheses indicate a weaker or indirect coupling.

Control input u_j	T_s	P_s	p_{CO_2}	p_{O_2}	p_{H_2}	α_{pol}
$\dot{M}_{CO_2}$	+	+	+	0	0	±
$\dot{M}_{H_2}$	+	(+)	0	0	+	−
$\dot{M}_{O_2}$	(+)	+	0	+	0	0
$\dot{M}_{N_2}$	(+)	+	0	0	0	(−)
$\dot{M}_{p,inj}$	+	0	0	0	0	−
A_m	+	0	0	0	0	−

where $M_{CO_2,acc}$ is the accessible CO_2 reservoir against which α_{pol} is defined and C_{eff} is an effective surface-atmosphere heat capacity. If $\dot{M}_i^{phys}$ is a physical mass flow rate, define the pressure-equivalent mass flow as in Eq. (6).

Linearizing about an operating point $(\mathbf{x}_0, \mathbf{u}_0)$ gives

$$\delta\dot{\mathbf{x}} = \mathbf{A}\,\delta\mathbf{x} + \mathbf{B}\,\delta\mathbf{u}, \quad \delta\mathbf{y} = \mathbf{C}\,\delta\mathbf{x}, \tag{95}$$

where $\mathbf{y}$ collects the controlled outputs of interest, e.g.,

$$\mathbf{y} = \left[T_s P_s p_{O_2} p_{CO_2} \alpha_{pol}\right]^{\mathsf{T}}, \quad P_s \text{ from Eq. (87)}. \tag{96}$$

The steady-state control-authority matrix is then

$$\mathbf{G}(0) = -\mathbf{C}\mathbf{A}^{-1}\mathbf{B}, \tag{97}$$

which quantifies how strongly each industrial control lever changes each planetary output near the operating point. The qualitative sign structure of this steady-state gain matrix is summarized in Table X.

The value of Eqs. (88)–(97) is not that they replace high-fidelity climate modeling, but that they define the minimum dynamical structure that any credible terraforming architecture must close: source terms, loss terms, sink terms, reservoir bookkeeping, and control authority. This makes the phrase "planetary-scale control system" mathematically precise.

For E4, this state-space model is an inventory-and-control skeleton, not a closed atmospheric-chemistry model. A predictive breathable-atmosphere model must add photochemical states such as O_3, CO, HO_x, NO_x, H_2O, oxidants, and long-lived synthetic gases, together with wavelength-dependent photolysis, heterogeneous dust/aerosol chemistry, and surface sink terms. The inventories derived here are therefore necessary conditions; photochemical closure is an additional feasibility requirement.

IX. ENERGY DELIVERY AND SYSTEM ARCHITECTURE

A. Energetic "success" thresholds

Table XI summarizes representative energy scales for several canonical operations. These numbers should be interpreted as *lower bounds*; real architectures add large multiplicative factors for inefficiency, infrastructure, and operations.

TABLE XI. Representative energetic scales (order of magnitude).

Operation	Scale assumption	Energy/size
Sublimate ∼6 mbar CO_2 polar deposit	$M \sim 2.3 \times 10^{16}$ kg, $L_{sub} \sim 6 \times 10^5$ J/kg	$\sim 1 \times 10^{22}$ J
Produce human-breathable benchmark $p_{O_2} = 21$ kPa from water	$M_{O_2} \sim 9 \times 10^{17}$ kg, reversible ΔG°	$\sim 1.3 \times 10^{25}$ J
Global forcing by orbital mirrors	$\Delta F_{TOA} = 20\,\mathrm{W\,m^{-2}}$, $\eta_m = 0.7$	$A_m \sim 7 \times 10^{12}$ m^2
Decadal engineered nanoparticle injection	$\dot{V} = 30\,\mathrm{l\,s^{-1}}$, $\rho = 3 \times 10^3\,\mathrm{kg\,m^{-3}}$	$\sim 3 \times 10^{10}$ kg/decade

B. Throughput requirement: required mass flow and average power versus build time

For any proposed end point with target inventory M_{target} achieved over a build time t_{build}, the implied mean industrial throughput is

$$\dot{M}_{req} \equiv \frac{M_{target}}{t_{build}}. \tag{98}$$

Likewise, for an intervention requiring integrated energy E, the implied mean power is

$$P_{avg} \equiv \frac{E}{t_{build}}. \tag{99}$$

These relations are decisive because they convert "large but abstract" quantities into concrete kg s^{-1} and TW requirements that can be compared with industrial analogs and with plausible Mars power architectures.

Table XII emphasizes that breathable end points (E4) are not merely "large": They imply sustained, multidecade to multimillennial industrial mass flow rates orders of magnitude above the maintenance rates implied by aerosol-based warming proposals [14]. This does not prove impossibility, but it places E4 in the category of long-duration planetary industry rather than "scale-up of demonstrated *in situ* resource utilization (ISRU)."

Equations of the form $\dot{M} \sim M/t_{build}$ convert a thermodynamic/atmospheric end-state requirement into a required *continuous industrial mass flow*. This is a useful discriminator for feasibility because it folds together mining, chemical processing, plant uptime, logistics, maintenance, and scaling laws for an industrial base capable of sustaining the required rates. To make these rates interpretable, Sec. X D benchmarks them against present-day terrestrial energy use and bulk material production.

C. Power generation scaling: solar versus nuclear

For photovoltaic power, the relevant quantity is incident solar irradiance on the collecting aperture, not the planetary Bond-albedo-weighted absorbed flux. A conservative global-horizontal reference uses $S_{Mars}/4$; additional losses from latitude, dust, atmospheric opacity, night-side storage, curtailment, and downtime are collected into an availability factor $C_{PV} \leqslant 1$. The average electrical

TABLE XII. Illustrative throughput and power requirements for representative end points, showing the scaling with assumed build time. Values use $M_{O_2} \approx 9.0 \times 10^{17}$ kg for $p_{O_2} = 21$ kPa [Eq. (73)], $M_{N_2} \approx 1.9 \times 10^{18}$ kg for $p_{N_2} = 50$ kPa [Eq. (24)], and $E_{min} \approx 1.3 \times 10^{25}$ J for reversible electrolysis [Eq. (78)].

Target	M_{target} (kg)	t_{build}	$\dot{M}_{req}$ (kg s^{-1})	P_{avg} (TW)
O_2 for $p_{O_2} = 21$ kPa	9.0×10^{17}	10^3 yr	2.8×10^7	420
O_2 for $p_{O_2} = 21$ kPa	9.0×10^{17}	10^2 yr	2.8×10^8	4200
N_2 buffer for $p_{N_2} = 50$ kPa	1.9×10^{18}	10^3 yr	6.0×10^7	–
N_2 buffer for $p_{N_2} = 50$ kPa	1.9×10^{18}	10^2 yr	6.0×10^8	–

TABLE XIII. Photovoltaic (PV) area required to supply average electrical power on Mars using Eq. (101), with $S_{\rm Mars} = 589\ {\rm W\,m^{-2}}$. Values are $C_{\rm PV} = 1$ aperture lower bounds; realistic installations require additional area for latitude, dust, night-side storage, atmospheric opacity, and downtime.

Average power	$A_{\rm PV}$ at $\eta_{\rm PV} = 0.2$	$A_{\rm PV}$ at $\eta_{\rm PV} = 0.3$
1 TW	$3.4 \times 10^4\ {\rm km^2}$	$2.3 \times 10^4\ {\rm km^2}$
10 TW	$3.4 \times 10^5\ {\rm km^2}$	$2.3 \times 10^5\ {\rm km^2}$
100 TW	$3.4 \times 10^6\ {\rm km^2}$	$2.3 \times 10^6\ {\rm km^2}$
1000 TW	$3.4 \times 10^7\ {\rm km^2}$	$2.3 \times 10^7\ {\rm km^2}$

power delivered by an area $A_{\rm PV}$ is therefore

$$P \simeq A_{\rm PV}\,\eta_{\rm PV}\,C_{\rm PV}\,\frac{S_{\rm Mars}}{4}. \tag{100}$$

Thus, the panel aperture area required for average electrical power P is

$$A_{\rm PV} \approx \frac{4P}{\eta_{\rm PV} C_{\rm PV} S_{\rm Mars}} \approx 3.4 \times 10^4\ {\rm km^2} \left(\frac{P}{1\ {\rm TW}}\right)\left(\frac{0.2}{\eta_{\rm PV}}\right)\left(\frac{1}{C_{\rm PV}}\right)\left(\frac{589\ {\rm W\,m^{-2}}}{S_{\rm Mars}}\right). \tag{101}$$

The values in Table XIII should thus be interpreted as $C_{\rm PV} = 1$ lower bounds; realistic systems require larger areas.

All mirror and PV areas reported here are optimistic lower bounds; realistic designs increase required area by factors of a few to account for pointing losses, dust, latitude, diurnal cycling, seasonal insolation, and storage. These areas are not physically impossible, but they imply continent-scale deployment and maintenance, plus storage for night and multiweek dust storms. High-capacity fission or fusion provides an alternative with much higher power density but large development and safety overhead. Most plausible long-term architectures are therefore *hybrid*: nuclear baseload plus solar where practical, with large-scale energy storage and distribution.

D. Architecture as a coupled control problem

Terraforming is not a single technology but a coupled planetary climate-industrial system whose source terms, loss terms, and feedbacks must be monitored and actively managed over long timescales. A credible architecture must include the following:

(1) *Energy generation and transmission.* Required power is mechanism- and end point-dependent: gigawatt-class source powers may be relevant for some warming-only particle pathways, whereas open-atmosphere E4 oxygenation requires multi-10^2 TW- to PW-class average power. Energy delivery is limited by transmission, storage, and reliability (surface grids, beamed power, or distributed reactors).

TABLE XIV. Order-of-magnitude comparison between Mars E4 build requirements (for $t_{\rm build} = 10^3$ yr) and representative present-day terrestrial industry magnitudes.

Quantity	Mars E4 (1000 yr build)	Earth today
$\dot{M}_{\rm O_2}$	$9.0 \times 10^{14}\ {\rm kg\,yr^{-1}} \approx 900\ {\rm Gt\,yr^{-1}}$	–
$\dot{M}_{\rm N_2}$	$1.9 \times 10^{15}\ {\rm kg\,yr^{-1}} \approx 1900\ {\rm Gt\,yr^{-1}}$	–
Global material extraction	–	$\sim 1.06 \times 10^2\ {\rm Gt\,yr^{-1}}$ [55]
World crude steel production	–	$\sim 1.89\ {\rm Gt\,yr^{-1}}$ [56]
Global primary energy (avg)	–	$620\ {\rm EJ\,yr^{-1}} \approx 20\ {\rm TW}$ [57]

(2) *Mass production and logistics.* Industrial throughput must be stated by pathway rather than globally. Warming-only particle pathways can lie in the megaton-per-year class, whereas O_2/buffer-gas atmospheric build-out for E4 requires gigaton-to-teraton-per-year class mass flow sustained for centuries.

(3) *Climate monitoring and feedback control.* The required degree of active control is mechanism-dependent. Some interventions may benefit from self-lofting or radiative-dynamical feedbacks, but any serious planetary-scale program still requires continuous sensing and at least some actuation authority (factory throttling, mirror pointing, gas-production control, or particle release management) to maintain targets and avoid collapse/hysteresis.

(4) *Planetary protection, biosafety, and biogeochemical coupling.* Biological components can grow rapidly, but ecological turnover, nutrient cycling, burial/export efficiency, and sink competition introduce additional system variables and long biogeochemical response times that lie outside the present minimal engineering closure.

E. Staged roadmap: from regional habitability to global modification

Given the scale barriers for E4 end points, a technically plausible progression is as follows:

(1) *Phase I (decades): regional paraterraforming.* Deploy solid-state greenhouse layers (silica aerogel) over ice-rich terrain to enable persistent melt and photosynthesis locally without global atmospheric change [11].

(2) *Phase II (century): climate nudging and resource mobilization.* Use localized warming to trigger limited volatile release, deploy pilot aerosol systems, and build power and manufacturing base.

(3) *Phase III (multicentury): sustained atmospheric engineering.* If pursued, combine an operationally sustained warming mechanism (aerosols/PFC/CIA) with buffer-gas import or extraction and large-scale oxygen production, while simultaneously investing in retention and geochemical management.

This staged view recognizes that the fastest path to meaningful surface habitability is likely *regional* rather than *planetary*.

X. DISCUSSION AND IMPLICATIONS

The constraint-based framework in Table IV shows that the dominant feasibility constraint shifts sharply with the target end state. For regional end points (E1–E2), deployment area and local power dominate; for global pressure (E3), atmospheric inventory dominates; and for breathable end points (E4), composition inventories, minimum work, and sink management dominate. For near-term habitability return per unit industrial scale, the results favor regional approaches: paraterraforming and contained biospheres, local thermal/insulation control, and build-out of power and autonomous manufacturing. If global modification is pursued, credible architectures must be evaluated as coupled control systems that explicitly close (1) inventory sourcing, (2) radiative authority, (3) sustained throughput and power, and (4) retention/sink management over multicentury timescales.

A. Which constraint dominates depends on the end point

A central implication of the constraint-based framework is that the limiting factor shifts with the target end state:

(1) For E1–E2 (regional water stability, enclosed-volume agriculture), the dominant constraints are local radiative control and deployment area; global 10^{18}-kg-scale gas inventories are not required.

(2) For E3 (no ebullism), the dominant constraint becomes atmospheric mass: $M_{\mathrm{atm}} \approx 2.4 \times 10^{17}$ kg (Table III), plus enough warming to avoid CO_2 collapse.

(3) For E4 (breathable), the dominant constraints are composition inventories and energy: $M_{O_2} \approx 9.0 \times 10^{17}$ kg and $M_{N_2} \approx 1.9 \times 10^{18}$ kg, with reversible oxygenation energy $E_{\min} \approx 1.3 \times 10^{25}$ J (Sec. VII) and additional sink filling (Sec. VII D).

Where the scaling suggests feasibility (aerosols/CIA), quantitative climate outcomes still require coupled 3D GCM + aerosol microphysics for better than an order-of-magnitude accuracy.

B. Design rules implied by the numbers

The results above imply several robust design rules:

(1) *Warming-only proposals cannot shortcut buffer-gas and oxygen inventories.* Even if T_s targets are met via aerosols or CIA, E4 remains dominated by 10^{18} -kg-class O_2 and N_2/Ar.

(2) *Mass-efficient radiative levers split into maintenance-dominated and fill-dominated classes.* For IR-active particles, the representative source rate is $\dot{M}_p \sim 90\ \mathrm{kg\,s^{-1}}$ (i.e., $\sim 3 \times 10^9\ \mathrm{kg\,yr^{-1}}$ for the illustrative $30\ \mathrm{L\,s^{-1}}$ case), and the maintained column mass scales as $\Sigma_p \sim \dot{M}_p \tau_p / (4\pi R_{\mathrm{Mars}}^2)$ (Sec. VI C); these pathways are genuinely maintenance limited. In contrast, Fig. 2 and Sec. VI B 3 show that global CO_2-H_2 CIA is typically fill dominated during century-to-millennial build-out, with TW-class replenishment becoming the relevant benchmark for long hold times rather than for initial construction.

(3) *Initial deployment is structure-mass-and-operations limited.* Equation (60) implies continent-scale mirrors for ΔF_{TOA} of order 10–20 $\mathrm{W\,m^{-2}}$, with total mass $M_m \sim \sigma_m A_m$ in the 10^{10}–10^{11} kg range for $\sigma_m \sim 1$–$10\ \mathrm{g\,m^{-2}}$. However, the harder problem is not launch energy alone: It is aggregate deployment count, pointing/control, non-Keplerian orbit maintenance, reflector degradation, and replacement over long hold times. Solar sails can assist transfer and trim of modular reflectors, but current and roadmap sail systems remain many orders of magnitude below the aggregate area implied by global Mars forcing targets.

(4) *E4 is a planetary industry problem: power + throughput + control.* Even at thermodynamic minima, $E_{\min}$ implies ~420 TW averaged over 10^3 yr, and Table XII implies $\dot{M} \sim 10^7$–$10^8\ \mathrm{kg\,s^{-1}}$ for century-to-millennial build times—far beyond "scaled ISRU."

C. Technology implications and a credible maturation path

If the objective is to maximize near-term habitability impact per unit industrial scale, the results favor (1) regional paraterraforming (solid-state greenhouse) and enclosed agriculture; (2) MW-GW-class power and localized thermal control; and (3) atmospheric monitoring/control infrastructure. If global end points are pursued, the pacing technologies become continent-scale energy generation, multigigaton/year materials handling, and long-duration closed-loop climate control (sensing + actuation), not a single "magic gas."

A complementary class of concepts is quasiglobal paraterraforming or "worldhouse" enclosure of the surface [54]. Such approaches largely avoid the 10^{18}-kg-scale open-atmosphere mass requirement by trading it into membrane/covering area, fabrication, puncture tolerance, and long-term maintenance. Regional/paraterraformed Mars is a difficult but technically plausible long-horizon development path; a global open-air breathable Mars is not currently a closed engineering program under known volatile inventories and present-to-forecast manufacturing capability.

A practical way to summarize the abiotic architecture result is therefore as follows: Use space industry to build habitable area before trying to build habitable planet. The staging logic is area first, atmosphere later: local coverings, local water access, local biosystems, local pressure control, and only then any attempt at larger-scale atmospheric or orbital climate intervention. This ordering is not merely conservative; it is what the quantitative mass and power bookkeeping favors.

With current inventories, current solar-sail/reflector technology, and plausible 21st-century industrial assumptions, there is no currently credible path to a *global open-atmosphere* E3/E4 Mars. The simultaneous requirements—10^{18}-kg-scale volatile supply, sustained $\mathcal{O}(10^2)$ TW- to

PW-class power, century-to-millennial climate control, and sink/retention management—are not jointly closed by any proposal surveyed in this paper.

A plausible path does exist, but it is narrower: staged *regional* habitability via covered or partly covered systems. This includes local aerogel-style solid-state greenhouse deployment over ice-rich terrains [11], enclosed agriculture at $P_{\rm in} \sim$ 10–30 kPa, and, at larger scales, worldhouse-type coverings that trade atmospheric mass into membrane area and repair logistics [12,13,54]. In this regime, microbes and plants are useful because they operate inside controlled environments as oxygen recyclers, food producers, biomaterial sources, and bioprocessing agents, rather than as exposed planetary-scale terraforming agents.

The gas-inventory scale of this plausible path is quantified by the enclosed-volume scaling in Eq. (84). For $H = 3\,\mathrm{m}$, $T = 250\,\mathrm{K}$, and $\bar{\mu} \simeq 29\,\mathrm{g\,mol^{-1}}$, the required gas mass per unit covered area is

$$\frac{M_{\rm encl}}{A_{\rm reg}} \simeq 0.42\ \mathrm{kg\,m^{-2}} \left(\frac{P_{\rm in}}{10\ \mathrm{kPa}}\right)\left(\frac{H}{3\ \mathrm{m}}\right)\left(\frac{250\ \mathrm{K}}{T}\right).$$

Thus, a 10^6 km^2 covered region requires only $\sim 4 \times 10^{11}$ kg of gas at 10 kPa, or $\sim 2 \times 10^{12}$ kg at 50 kPa, before leakage and reserve margins. These are large engineering inventories, but they are not planetary-atmosphere inventories. For covered regional systems, the dominant burdens become membrane/roof area, pressure load, anchoring, repair, water access, power, thermal control, and ecological management rather than global atmospheric mass.

The practical implication is that "Mars terraforming" should be disaggregated into two very different classes. Regional/paraterraformed Mars is a difficult but technically plausible long-horizon development. A global open-air breathable Mars is not currently a plausible engineering program on known inventories and present-to-forecast sail/manufacturing technology.

This end-point-based feasibility split is broadly consistent with recent roadmap-style assessments of Mars warming [17], which emphasize local membranes, reflectors for key sites, and engineered aerosols as early research priorities. Our contribution is to show that even if such warming pathways prove workable, they do not by themselves close the 10^{18}-kg-scale inventory and oxygen/buffer-gas requirements associated with open-atmosphere global E3/E4 end points.

D. Industrial scaling, timescale-power trade, and energy and installed-power lower bounds

A central result of this paper is that once the end point moves beyond E2–E3, feasibility is governed less by the existence of a mechanism and more by whether an industrial base can sustain the required *mass flow* and *power* for centuries to millennia. For the nominal E4 targets summarized above, the required atmospheric inventories are $M_{\rm O_2} \simeq 9.0 \times 10^{17}$ kg and $M_{\rm N_2} \simeq 1.9 \times 10^{18}$ kg, with a minimum reversible electrochemical work $E_{\rm min} \simeq 1.3 \times 10^{25}$ J for oxygenation alone (exclusive of compression, losses, and coproducts).

1. Biological oxygenation and E2.5 bootstrapping

The abiotic oxygenation estimates above should not be interpreted as limits on biological bootstrapping. Photosynthetic pathways are better described by areal productivity and the fraction of fixed carbon that is exported or buried rather than respired. If $\phi_{\rm C}$ is the net fixed-carbon productivity available for long-term O_2 accumulation, $A_{\rm prod}$ is the productive area, and β is the fraction of fixed carbon that remains reduced on the relevant timescale, then

$$\dot{M}_{\rm O_2,bio} \simeq \frac{32}{12}\,\beta\,\phi_{\rm C}\,A_{\rm prod}. \tag{102}$$

Equivalently, for productive fractional area $f_{\rm prod}$ on Mars,

$$t_{\rm O_2,bio} \simeq \frac{M_{\rm O_2}}{(32/12)\,\beta\,\phi_{\rm C}\,f_{\rm prod}\,4\pi R_{\rm Mars}^2}. \tag{103}$$

For a column corresponding to 100 mbar O_2 on Mars,

$$m_{O_2}(100\ \mathrm{mbar}) = \frac{10^4\ \mathrm{Pa}}{g_{\mathrm{Mars}}} \simeq 2.7 \times 10^3\ \mathrm{kg\,m^{-2}}, \tag{104}$$

requiring burial or export of

$$m_C \simeq \frac{12}{32} m_{O_2} \simeq 1.0 \times 10^3\ \mathrm{kg\,C\,m^{-2}}. \tag{105}$$

Using a landlike productivity normalization $\phi_C \simeq 0.43\ \mathrm{kg\,C\,m^{-2}\,yr^{-1}}$, the corresponding idealized timescale is

$$t_{100\mathrm{mbar,bio}} \simeq 2.3 \times 10^3\ \beta^{-1} \left(\frac{0.43\ \mathrm{kg\,C\,m^{-2}\,yr^{-1}}}{\phi_C} \right) f_{\mathrm{prod}}^{-1}\ \mathrm{yr}. \tag{106}$$

This estimate is not a prediction of terraforming time. It is a productivity benchmark showing that biological oxygenation is controlled by productive area, water, nutrients, radiation protection, ecology, carbon burial/export, and sink competition. Selective breeding and synthetic biology could change ϕ_C, stress tolerance, and material production. Thus, biological bootstrapping is a distinct pathway class (i.e., especially for E2.5) for producing useful regional habitability, biomass, materials, local O_2, and soil conditioning; it can change the route to habitability, but it does not remove the terminal inventory requirement for an open human-breathable E4 atmosphere.

2. Dominant constraint as a function of end state

For E1–E2 regional habitability, the dominant constraints are deployment area and local power; global volatile inventories are not required. For E3 global pressure, the mass inventory constraint becomes dominant ($M_{\mathrm{atm}} \sim 2.4 \times 10^{17}$ kg at the Armstrong limit). For E4 breathable end points, composition inventories and energy dominate: $M_{O_2} \sim 9 \times 10^{17}$ kg and buffer gases at the 10^{18} kg scale, with oxygenation work $E_{\mathrm{min}} \sim 10^{25}$ J even at reversible limits.

3. Timescale-power-throughput trade

For any target inventory M achieved over build time t_{build},

$$\dot{M} \equiv \frac{M}{t_{\mathrm{build}}}, \quad \bar{P} \equiv \frac{E}{t_{\mathrm{build}}}. \tag{107}$$

Here and below, an overbar denotes a build-time average (e.g., $\bar{P} \equiv E/t_{\mathrm{build}}$). These relations translate end states into continuous industrial requirements. For example, $p_{O_2} = 21$ kPa implies $M_{O_2} \approx 9.0 \times 10^{17}$ kg, so $\dot{M}_{O_2} \approx 2.8 \times 10^7$ to 2.8×10^8 kg s^{-1} for $t_{\mathrm{build}} = 10^3$–$10^2$ yr. The reversible oxygenation floor implies $\bar{P}_{\mathrm{min}} \approx 0.42$–$4.2$ PW over the same range. Real systems require $\bar{P} \gtrsim \bar{P}_{\mathrm{min}}/\eta_{\mathrm{sys}}$, with $\eta_{\mathrm{sys}} \ll 1$ once compression, separation, thermal losses, and downtime are included.

4. Benchmarking E4 mass flows against terrestrial industry

For a build time t_{build}, the implied continuous production/import rates are

$$\dot{M}_{O_2} \simeq \frac{M_{O_2}}{t_{\mathrm{build}}}, \quad \dot{M}_{N_2} \simeq \frac{M_{N_2}}{t_{\mathrm{build}}}. \tag{108}$$

Table XIV compares these Mars E4 build requirements with representative present-day terrestrial industry magnitudes. For $t_{\mathrm{build}} = 10^3$ yr, this corresponds to $\dot{M}_{O_2} \approx 9.0 \times 10^{14}$ kg yr^{-1} ($\approx$900 Gt yr^{-1}) and $\dot{M}_{N_2} \approx 1.9 \times 10^{15}$ kg yr^{-1} ($\approx$1900 Gt yr^{-1}), i.e., $\mathcal{O}(10^4)$–$\mathcal{O}(10^5)$ metric tons per second.

As a reality check, global material extraction on Earth has grown from ~30 to ~106 billion metric tons per year since 1970 [55], and world crude steel production is ~1.9 billion metric tons

per year [56]. Thus, even under a 10^3-yr schedule, the *oxygen alone* production rate is of order $\sim 8\times$ present-day total terrestrial material extraction, and $\sim 400\times$ present-day global steel output. The nitrogen buffer requirement is larger still.

Importantly, $\dot{M}_{O_2}$ and $\dot{M}_{N_2}$ are not directly mined "as is"; they must be extracted from feedstocks. If oxygen is produced from mined water ice, the required water throughput is

$$\dot{M}_{H_2O} \simeq \tfrac{36}{32}\,\dot{M}_{O_2} \approx 1.125\,\dot{M}_{O_2}, \tag{109}$$

while if oxygen (or nitrogen-bearing species) is extracted from regolith with effective yield f by mass, the required regolith handling is $\dot{M}_{reg} \sim \dot{M}/f$. For plausible $f \sim 0.01$–0.1, regolith throughput rises by 1–2 orders of magnitude, pushing the problem decisively into "planetary industry" rather than "chemical plant" scaling.

5. Timescale-power trade: what does "accelerating" E4 actually require?

The minimum average power required to supply the reversible oxygenation work is

$$\bar{P}_{min} \simeq \frac{E_{min}}{t_{build}} \simeq 4.2 \times 10^{14}\,\mathrm{W}\left(\frac{10^3\ \mathrm{yr}}{t_{build}}\right)\left(\frac{E_{min}}{1.3 \times 10^{25}\ \mathrm{J}}\right). \tag{110}$$

For $t_{build} = 10^3$ yr, $\bar{P}_{min} \approx 420$ TW, which is $\sim 20\times$ today's global primary energy consumption rate ($\approx 620\,\mathrm{EJ\,yr^{-1}}$) [57]. Conversely, if one demands $t_{build} \sim 100$ yr, Eq. (110) implies $\bar{P}_{min} \sim 4.2$ PW *before* accounting for real inefficiencies. Because real systems incur conversion, compression, separation, thermal losses, and downtime, a more realistic requirement is $\bar{P} \sim \bar{P}_{min}/\eta_{sys}$ with $\eta_{sys} \ll 1$, pushing required power further upward.

Using Eq. (101), the PV area required to supply the reversible oxygenation power floor is

$$A_{PV} \sim (3.4 \times 10^4\,\mathrm{km^2/TW})\left(\frac{\bar{P}_{min}}{1\ \mathrm{TW}}\right)\left(\frac{0.2}{\eta_{PV}}\right)\frac{1}{C_{PV}}, \tag{111}$$

where $C_{PV} \leqslant 1$ is an effective capacity factor capturing diurnal cycling, latitude, dust storms, and storage limits. For $\bar{P}_{min} \approx 420$ TW (1000-yr build), even the optimistic $C_{PV} = 1$ bound implies $A_{PV} \sim 1.4 \times 10^7\,\mathrm{km^2}$ at $\eta_{PV} = 0.2$. For $t_{build} = 100$ yr, $\bar{P}_{min} \sim 4.2$ PW implies A_{PV} comparable to or exceeding Mars' total surface area unless power is provided by higher-density sources (nuclear/fusion) or by space-based collection and beaming. This illustrates that accelerating E4 is fundamentally a power-density problem, not only a chemistry problem.

6. Energy-purchase floor and installed-power lower bound

The following estimates are not a terraforming cost model. They are irreducible lower bounds: an energy-purchase floor associated with the reversible oxygenation work, and an installed-power scaling associated with sustaining the required average power. They exclude capital deployment, maintenance, mining, separations, compression, storage, distribution, construction, financing, labor, reliability margins, and replacement. The reversible oxygenation energy corresponds to

$$E_{min} \approx 3.7 \times 10^{18}\ \mathrm{kWh}, \tag{112}$$

so even at an optimistic electricity price c_e one has an energy-purchase floor

$$C_{E,min} \gtrsim (3.7 \times 10^{18}\ \mathrm{kWh})\,c_e \approx 1.9 \times 10^{17}\ \$\left(\frac{c_e}{0.05\ \$/\mathrm{kWh}}\right), \tag{113}$$

exclusive of capital expenditures, maintenance, and the additional energy required for nonreversible steps (compression, gas separation, transport, and thermal management).

Similarly, if the generation system scales with cost c_W per installed watt, then the capex scaling for sustaining $\bar{P}$ is

$$C_{\rm capex} \sim c_W \bar{P} \approx 4.2 \times 10^{14}\ \$ \left(\frac{c_W}{1\ \$/\mathrm{W}}\right)\left(\frac{\bar{P}}{420\ \mathrm{TW}}\right), \tag{114}$$

showing that even wildly optimistic c_W values imply civilization-scale investment when $\bar{P}$ is in the multi-10^{14}-W class.

7. What industrial activities are actually implied?

The constraints above map directly onto different industrial primitives depending on the end point. For warming-only aerosol pathways, published studies imply source powers and material flows that may lie in the GW class and $\sim 10^2$-kg-s^{-1} class [14,15]. In contrast, for century-to-millennial E4 oxygenation and buffer-gas build-out, the relevant primitives escalate to multi-10^2 TW- to PW-class power, very large excavation/beneficiation throughput if feedstock yields are low, and long-duration climate/composition control. The article's largest throughput and power numbers therefore apply to *E4-scale atmospheric build-out*, not to every warming-only pathway.

E. What would change the feasibility classification?

The abiotic open-atmosphere classification shifts only if one or more of the following becomes true: (1) discovery or delivery of much larger accessible CO_2 and/or fixed-nitrogen inventories; (2) a supergreenhouse constituent with strong window absorption, long lifetime, and abundant *in situ* feedstock; (3) megascale reflector production and deployment at low areal mass with manageable orbit control and replacement; or (4) an industrial ecology capable of sustaining 10^2 TW-PW power and 10^3-Gt yr^{-1}-class throughput for centuries. Separately, biological bootstrapping can shift the *useful-habitability* classification at E2.5 without closing the open-atmosphere E4 inventory problem. Absent such changes, E1–E2.5 are the credible staged-development targets, whereas abiotic global E3–E4 remains a planetary-industry, long-horizon program.

XI. CONCLUSIONS

Terraforming Mars is constrained by coupled planet-scale budgets in atmospheric inventory, radiative forcing/opacity, industrial throughput and power, and long-term stability. Using transparent order-of-magnitude scalings, we find the following:

(1) *Open-surface pressure targets are* 10^{17}–10^{18}*-kg scale.* These targets translate directly into atmospheric mass [Eq. (19)]. E3/E4 end points require 10^{17}–10^{18} kg of gas, far above endogenous CO_2 inventories. Hydrostatic balance implies $M_{\rm atm} \simeq 3.89 \times 10^{15}$ kg per mbar [Eq. (20)]. Thus, even the Armstrong-limit pressure ($P_s = 6.27$ kPa) corresponds to $M_{\rm atm} \approx 2.4 \times 10^{17}$ kg (Table III).

(2) *Accessible endogenous* CO_2 *is inventory limited and cannot deliver melt-class global climates.* Current evidence supports an *order-of-tens-of-millibar* accessible endogenous CO_2 budget rather than a secure multibar reservoir (Sec. V). A representative $\sim$20 mbar case yields $\lesssim$10 K warming under present insolation, leaving a large melt-class thermal shortfall to globally stable surface liquid water.

(3) *Melt-class* T_s *requires either* $\tau_{\rm IR} \sim 2$–4 *at current insolation or extreme direct forcing.* Thermal targets require very large effective forcing. Accessible CO_2 ($\sim$20 mbar) yields $\lesssim 10$ K warming, whereas $\sim$60 K is needed for globally stable surface liquid water [6]. In a minimal gray-atmosphere approximation, $T_s \sim 250$–273 K at $T_e \approx 210$ K requires $\tau_{\rm IR} \sim 2$–4 [Table IV; Eq. (36), Sec. IV B]. Direct insolation modification requires $\Delta F_{\rm TOA}$ at the $\mathcal{O}(10^2)$ W m^{-2} level for large T_e shifts, implying continent-scale mirrors for $\Delta F_{\rm TOA} \sim \mathcal{O}(10^2)$ W m^{-2} [Eqs. (34) and (60)].

(4) *Breathable end points are bottlenecked primarily by composition inventories and oxygenation energy.* A human-breathable benchmark $p_{O_2} = 21$ kPa implies $M_{O_2} \approx 9.0 \times 10^{17}$ kg [Eq. (73)], and modest buffer-gas targets $p_{N_2} = 50$ kPa require comparable or larger masses $M_{N_2} \approx 1.9 \times 10^{18}$ kg

[Eq. (24)]. The reversible minimum oxygenation work is $E_{\min} \approx 1.3 \times 10^{25}$ J [Eq. (78)], implying $\bar{P}_{\min} \sim 0.42$–4.2 PW for $t_{\mathrm{build}} = 10^3$–$10^2$ yr before inefficiencies and before sink-management costs (Sec. VII D).

(5) *Mass-efficient warming levers split into maintenance-dominated and fill-dominated classes, while initial reflector deployment trades into extreme aggregate area and operations.* Aerosol pathways couple to sustained injection set by residence time (Sec. VI C) and are genuinely maintenance limited. In contrast, global CO_2-H_2 CIA requires long-term replenishment against escape but is typically fill dominated during century-to-millennial build-out because a *large, design-dependent global* H_2 *inventory* must be created before the warm state can be held (Secs. VI B 3 and VIII C). For representative warm-state cases, this inventory can correspond to H_2 partial pressures ranging from multimillibars to multikilopascals, depending on the chosen total pressure P_{tot} and mixing ratio $f_{\mathrm{H_2}}$. Reflector-based TOA forcing scales to continent-class areas for global ΔF_{TOA}; reflector-based forcing pathways are therefore structure and operations intensive, while breathable global end points still require a civilization-scale industrial base operating for centuries to millennia (Sec. X D).

Taken together, these scalings imply a sharp feasibility split. Regional and covered-area pathways (E1–E2.5) are the technically favored route under known inventories. They include local solid-state greenhouse deployment, enclosed agriculture, covered-area or worldhouse concepts, and biologically productive systems used where biology is strongest: controlled environments, food production, biomaterials, O_2 recycling, soil conditioning, and bioprocessing. Their central advantage is that they scale primarily with covered area, local power, water access, and maintenance, not with planetwide atmospheric mass.

In contrast, an abiotic global open-atmosphere transformation to E3–E4 is not currently a closed engineering program. With known accessible volatiles, no surveyed mechanism simultaneously closes the pressure, composition, warming, power, sink, retention, and operations constraints. A human-breathable open Mars becomes credible only under explicit conditions: 10^{18}-kg-scale volatile supply, especially N-bearing buffer gas; sustained 10^2 TW-PW power and 10^3-Gt yr^{-1}-class material processing over centuries to millennia; and long-duration climate/composition control with sink and replacement management.

The principal contribution of this paper is a reusable architecture-level framework: end-point-normalized lower bounds, feasibility numbers ($\Pi_M, \Pi_F, \Pi_{\dot{M}}, \Pi_P, \Pi_S$), and a fill-versus-maintenance criterion $\Lambda_{\mathrm{maint}} = t_{\mathrm{build}}/\tau_{\mathrm{loss}}$. This framework separates inventory limited, radiatively limited, throughput limited, power limited, and maintenance limited concepts, and cleanly distinguishes true maintenance pathways such as short-lived aerosols from fill-dominated pathways such as large H_2 inventories for CO_2-H_2 CIA. These results should not be read as a policy recommendation for or against terraforming Mars; they define necessary physical and industrial conditions for candidate end states, while questions of desirability, planetary protection, governance, scientific preservation, and ethics remain separate constraints.

Ultimately, the burden of proof for any proposed global open-atmosphere terraforming pathway is therefore quantitative: It must close the end point inventory, radiative, throughput, power, loss, sink, and operations inequalities in Table IV, not merely identify a warming mechanism. The most technically defensible roadmap is therefore staged: Pursue E1–E2.5 deployments first, while using them to develop the power, excavation, processing, logistics, biosphere-control, and monitoring/actuation capabilities required for any credible later E3–E4 attempt.

ACKNOWLEDGMENT

The work described here was carried out at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, under Contract No. 80NM0018D0004 with the National Aeronautics and Space Administration.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

[1] C. P. McKay, Terraforming Mars, J. Br. Interplanet. Soc. **35**, 427 (1982).

[2] C. P. McKay, O. B. Toon, and J. F. Kasting, Making Mars habitable, Nature (London) **352**, 489 (1991).

[3] R. Zubrin and C. P. McKay, Technological requirements for terraforming Mars, in *Proceedings of the 29th Joint Propulsion Conference and Exhibit*, AIAA Paper 93-2005 (American Institute of Aeronautics and Astronautics, Reston, VA, Monterey, CA, 1993), pp. 1–13.

[4] M. J. Fogg, *Terraforming: Engineering Planetary Environments* (SAE International, Warrendale, PA, 1995).

[5] C. P. McKay and M. M. Marinova, The physics, biology, and environmental ethics of making Mars habitable, Astrobiology **1**, 89 (2001).

[6] B. M. Jakosky and C. S. Edwards, Inventory of CO_2 available for terraforming Mars, Nat. Astron. **2**, 634 (2018).

[7] B. M. Tutolo *et al.*, Carbonates identified by the Curiosity rover indicate a carbon cycle operated on ancient Mars, Science **388**, 292 (2025).

[8] P. B. Buhler, A. P. Ingersoll, B. L. Ehlmann, P. O. Hayne, and S. Piqueux, Coevolution of Mars's atmosphere and massive south polar CO_2 ice deposit, Nat. Astron. **4**, 364 (2020).

[9] P. B. Buhler and S. Piqueux, Obliquity-driven CO_2 exchange between Mars' atmosphere, regolith, and polar cap, J. Geophys. Res. Planets **126**, e2020JE006759 (2021).

[10] A. Broquet, M. A. Wieczorek, and W. Fa, The composition of the South Polar Cap of Mars derived from orbital data, J. Geophys. Res. Planets **126**, e2020JE006730 (2021).

[11] R. Wordsworth, L. Kerber, and C. S. Cockell, Enabling Martian habitability with silica aerogel via the solid-state greenhouse effect, Nat. Astron. **3**, 898 (2019).

[12] R. Wordsworth and C. Cockell, Self-sustaining living habitats in extraterrestrial environments, Astrobiology **24**, 1187 (2024).

[13] R. Wordsworth, R. Quayum, E. Kocharian, A. Pearson, X. Portillo, M. Yang, C. S. Cockell, S. Nangle, and G. Church, Biomaterials for organically generated habitats beyond Earth, Sci. Adv. **11**, eadp4985 (2025).

[14] S. Ansari *et al.*, Feasibility of keeping Mars warm with nanoparticles, Sci. Adv. **10**, eadn4650 (2024).

[15] M. I. Richardson, S. Ansari, B. Fan, R. Ramirez, H. Mohseni, M. A. Mischna, M. H. Hecht, L. J. Steele, F. Sharipov, and E. S. Kite, Atmospheric dynamics of IR-active particles released from Mars' surface, Geophys. Res. Lett. **53**, e2025GL121051 (2026).

[16] E. A. DeBenedictis, E. S. Kite, R. D. Wordsworth, N. L. Lanza, C. S. Cockell, P. A. Silver, R. M. Ramirez, J. Cumbers, H. Mohseni, C. E. Mason, W. W. Fischer, and C. P. McKay, The case for Mars terraforming research, Nat. Astron. **9**, 634 (2025).

[17] E. S. Kite *et al.*, A research roadmap for assessing the feasibility of warming Mars, arXiv:2604.02242.

[18] J. M. Graham, The biological terraforming of Mars: Planetary ecosynthesis as ecological succession on a global scale, Astrobiology **4**, 168 (2004).

[19] M. H. Hecht, Metastability of liquid water on Mars, Icarus **156**, 373 (2002).

[20] N. Schorghofer, Mars: Quantitative evaluation of crocus melting behind boulders, Astrophys. J. **890**, 49 (2020).

[21] L. Lange and F. Forget, On the possibility of melting water ice during the recent past of Mars: Implications for the formation of gullies, J. Geophys. Res. Planets **131**, e2025JE009512 (2026).

[22] International Association for the Properties of Water and Steam (IAPWS), Revised release on the pressure along the melting and sublimation curves of ordinary water substance, IAPWS Release (R14-08(2011)), https://iapws.org/documents/release/MeltSub.

[23] NASA Office of the Chief Health and Medical Officer, Habitable atmosphere (OCHMO-TB-003 Rev A), NASA-STD-3001, Technical Brief, 2023, https://www.nasa.gov/wp-content/uploads/2023/12/ochmo-tb-003-habitable-atmosphere.pdf.

[24] M. Alexander, Mars transportation environment definition document, Tech. Rep. NASA/TM-2001-210935, NASA Marshall Space Flight Center, NTRS Document ID: 20010046858, 2001, https://ntrs.nasa.gov/citations/20010046858.

[25] L. R. Dartnell, L. Desorgher, J. M. Ward, and A. J. Coates, Modelling the surface and subsurface Martian radiation environment: Implications for astrobiology, Geophys. Res. Lett. **34**, 2006GL027494 (2007).

[26] J. E. Moores, P. H. Smith, R. Tanner, A. C. Schuerger, and K. J. Venkateswaran, The shielding effect of small-scale martian surface geometry on ultraviolet flux, Icarus **192**, 417 (2007).

[27] J. C. Stern *et al.*, Evidence for indigenous nitrogen in sedimentary and aeolian deposits from the Curiosity rover investigations at Gale crater, Mars, Proc. Natl. Acad. Sci. USA **112**, 4245 (2015).

[28] Edited by P. J. Linstrom and W. G. Mallard, NIST Chemistry WebBook, NIST Standard Reference Database Number 69, National Institute of Standards and Technology, Gaithersburg, MD, Carbon dioxide, https://webbook.nist.gov/cgi/cbook.cgi?ID=124-38-9 (accessed 22 May 2026).

[29] M. F. Gerstell, J. S. Francisco, Y. L. Yung, C. Boxe, and E. T. Aaltonee, Keeping Mars warm with new super greenhouse gases, Proc. Natl. Acad. Sci. USA **98**, 2154 (2001).

[30] M. M. Marinova, C. P. McKay, and H. Hashimoto, Radiative-convective model of warming Mars with artificial greenhouse gases, J. Geophys. Res. **110**, 2004JE002306 (2005).

[31] R. M. Ramirez, R. Kopparapu, M. E. Zugger, T. D. Robinson, R. Freedman, and J. F. Kasting, Warming early Mars with CO_2 and H_2, Nat. Geosci. **7**, 59 (2014).

[32] M. Turbet, H. Tran, O. Pirali, F. Forget, C. Boulet, and J.-M. Hartmann, Far infrared measurements of absorptions by $CH_4 + CO_2$ and $H_2 + CO_2$ mixtures and implications for greenhouse warming on early Mars, Icarus **321**, 189 (2019).

[33] P. J. Godin, R. M. Ramirez, C. L. Campbell, T. Wizenberg, T. G. Nguyen, K. Strong, and J. E. Moores, Collision-induced absorption of CH_4–CO_2 and H_2–CO_2 complexes and their effect on the ancient Martian atmosphere, J. Geophys. Res. Planets **125**, e2019JE006357 (2020).

[34] C. R. McInnes, Mars climate engineering using orbiting solar reflectors, in *Mars: Prospective Energy and Material Resources*, edited by V. Badescu (Springer, Berlin, Heidelberg, 2010), pp. 645–659.

[35] C. J. Handmer, How to terraform Mars for $50B with solar sails, in *Tenth International Conference on Mars*, LPI Contribution No. 3007 (Lunar and Planetary Institute, Houston, TX, 2024), Abstract No. 3025, p. 3025.

[36] C. Sagan, Planetary engineering on Mars, Icarus **20**, 513 (1973).

[37] S. G. Turyshev *et al.*, Science opportunities with solar sailing smallsats, Planet. Space Sci. **235**, 105744 (2023).

[38] S. G. Turyshev, Propulsion trades for a 2035–2040 solar gravitational lens mission, arXiv:2602.04198.

[39] W. K. Wilkie, The NASA advanced composite solar sail system (ACS3) flight demonstration a technology pathfinder for practical smallsat solar sailing, in *Proceedings of the 35th Annual AIAA/USU Conference on Small Satellites, Logan, UT, 2021* (Utah State University, Logan, UT, 2021), Paper SSC21-II-10, pp. 1–15.

[40] K. Wilkie and J. Fernandez, Advanced composite solar sail system (ACS3) mission update, *Presented at the 6th International Symposium on Space Sailing* (New York City College of Technology, City University of New York, New York, NY, 2023), presentation slides, pp. 1–26.

[41] L. Johnson *et al.*, Solar photon sails roadmap, Tech. Rep., NASA; roadmap targets include deploy-and-control areas $>2.5 \times 10^3$ m^2 (Generation 2) and $>10^4$ m^2 (Generation 3a), 2025.

[42] J. H. Kang, K. L. Gordon, R. G. Bryant, O. R. Stohlman, W. K. Wilkie, A. Stark, R. Barfield, B. Sindle, M. Finckenor, and P. Craven, Durability characterization of mechanical interfaces in solar sail membrane structures, Adv. Space Res. **67**, 2643 (2021).

[43] M. W. Chase, Jr., *NIST-JANAF Thermochemical Tables*, 4th ed. (American Chemical Society and American Institute of Physics for the National Institute of Standards and Technology, Washington, DC, 1998), Journal of Physical and Chemical Reference Data, Monograph No. 9.

[44] J. A. Hoffman *et al.*, Mars oxygen ISRU experiment (MOXIE)—Preparing for human Mars exploration, Sci. Adv. **8**, eabp8636 (2022).
[45] NASA, NASA's Oxygen-Generating Experiment MOXIE Completes Mars Mission, NASA News (2023), https://www.nasa.gov/missions/mars-2020-perseverance/perseverance-rover/nasas-oxygen-generating-experiment-moxie-completes-mars-mission/.
[46] G. J. Taylor, The bulk composition of Mars, Chem. Erde Geochem. **73**, 401 (2013).
[47] J. McSween, Y. Harry, G. J. Taylor, and M. B. Wyatt, Elemental composition of the Martian crust, Science **324**, 736 (2009).
[48] B. M. Jakosky *et al.*, Loss of the Martian atmosphere to space: Present-day loss rates determined from MAVEN observations and integrated loss through time, Icarus **315**, 146 (2018).
[49] R. Ramstad, S. Barabash, Y. Futaana, H. Nilsson, and M. Holmström, Ion escape from Mars through time: An extrapolation of atmospheric loss based on 10 years of Mars Express measurements, J. Geophys. Res. Planets **123**, 3051 (2018).
[50] J. L. Green, J. L. Hollingsworth, D. Brain, V. Airapetian, A. Glocer, A. Pulkkinen, C. Dong, and R. Bamford, A future Mars environment for science and exploration, in *Planetary Science Vision 2050 Workshop*, LPI Contribution No. 1989 (Lunar and Planetary Institute, Houston, TX, 2017), Abstract No. 8250, pp. 1–2.
[51] D. M. Hunten, The escape of light gases from planetary atmospheres, J. Atmos. Sci. **30**, 1481 (1973).
[52] D. C. Catling and J. F. Kasting, *Atmospheric Evolution on Inhabited and Lifeless Worlds* (Cambridge University Press, Cambridge, 2017).
[53] F. Forget, R. Wordsworth, E. Millour, J.-B. Madeleine, L. Kerber, J. Leconte, E. Marcq, and R. M. Haberle, 3D modelling of the early Martian climate under a denser CO_2 atmosphere: Temperatures and CO_2 ice clouds, Icarus **222**, 81 (2013).
[54] R. L. S. Taylor, The Mars atmosphere problem: Paraterraforming—The Worldhouse solution, J. Br. Interplanet. Soc. **54**, 236 (2001).
[55] United Nations Environment Programme and International Resource Panel, Global resources outlook 2024, Tech. Rep. (UNEP/International Resource Panel, Nairobi, Kenya, 2024).
[56] World Steel Association, World steel in figures 2024, Tech. Rep. (World Steel Association, Brussels, Belgium, 2024).
[57] Energy Institute, Statistical review of world energy 2024, Tech. Rep. (Energy Institute, 2024).