The Fleet That Wakes on Starlight

Interstellar Panspermia and the First Cell at Another Star

Douglas C. Youvan

About this edition

Douglas C. Youvan

Revised and expanded paperback edition, 2026

The reader’s guide follows the opening flight assumptions through the tests they require. Appendix E gives worked decisions on lineage provenance and habitat renewal. Appendix F contains the complete source for the computational companion and instructions for reproducing the book’s sequence and hydropathy analyses.

Contents

Preface 5
From the First Watt to the First Report 6

I The Fleet and Its Destinations

1 The Cargo That Must Stay Alive 9
2 Choosing the Stars 13
3 What Another Ship Buys 17
4 The Ships That Came Before 22
5 Ten Thousand Years in the Dark 26
6 Dose Along the Route 29
7 Catching the Destination Star 33
8 A Message and a Receiver 36

II How to Make the First Cell

9 The Genome and the Working Cell 40
10 The First Cell After Handoff 44
11 What an Unknown Receiver Can Infer 49
12 The Address and the Search 52
13 The Human Bath Begins Before Birth 56

III A Habitat for One Person

14 The Person in the Balance Sheet 60
15 A Floor Area for Food 64
16 When Air and Water Loops Fail 68
17 A Year Inside the Boundary 72
18 The Average Year and the Failed Harvest 76

Contents continued

IV A Biosphere That Renews Itself
19 The Slow Sinks....80
20 The Second Harvest....84
21 A Biosphere for a Community....88
22 Deciding When to Hand Off....92
V The Far End of the Link
23 A Workshop That Can Arrive Alone....96
24 Booting the Receiver....99
25 Archive of a Living World....103
26 Where a Seed Is Welcome....106
27 The First Delivery....109
VI Proving That Living Systems Renew
28 What Counts as the First Cell....113
29 Building the Next Living Layer....117
VII Qualifying the Fleet
30 One Ship Through One System....121
31 Three Cargoes on the Same Route....127
32 Choosing What the Fleet Must Prove....131
33 How a Mission Earns Its Next Attempt....136
34 Keeping the Record Alive....140
Appendix A Reproducing the Crossing Calculation....143
Appendix B Reproducing the Habitat Calculation....146
Appendix C Tests on a Received Sequence....150
Appendix D Recording an Attempt and Its Outcome....155
Appendix E Comparing Lineages and Habitats....157
Appendix F Running the Computational Companion....160
References....171

Preface

A ship leaves Earth without a living thing aboard. Ten thousand years later, if its receiver works, the light of another star is enough to wake it. The ship must slow down, stay, and bring a workshop into service. Then it can attempt to make the first living thing it was sent to establish. Now imagine that this ship never wakes. A fault in one connection, an error in the archive, or the lack of a workable orbit could end the whole attempt.

I would send a fleet. Several ships would go to each of a few nearby systems. They would carry self-contained equipment for their assigned tests and could proceed on their own. Some would try different ways of waking, braking, or building. A few would go farther, toward stars worth examining even at the price of a much longer wait. The fleet spends extra effort on chances to fail differently. Another ship can survive a fault that stops its neighbor. Another star can offer a resource none of the ships can find at their first destination. An error they all inherit is a different matter.

I arrived at this mission by imagining a single bacterium appearing all at once. Its DNA can be written in bits. Its membrane, enzymes, ribosomes, gradients, and other working parts also have to come from somewhere. I call this material and its organization the *bath*. Give two builders the same genome and different starting inventories. What can each make, and how long can the result keep living after the builder steps back? The same question applies to a fleet. Send two machines with the same instructions and compare what they can establish under different conditions.

Other authors saw much of the mission. Claudius Gros proposed a probe that makes microbes at its destination. Vernor Vinge gave an AI a hundred centuries to reach Alpha Centauri. James P. Hogan and Arthur C. Clarke imagined ships that make and raise the people they were sent to establish. Their work leaves a set of questions I want to calculate. Where does the first usable watt come from? What has to be present when the first genome is read? How many independent attempts does the fleet really buy?

At the scale of a person, a crop can feed someone for a season by drawing down an initial stock. A second harvest grown from measured recovered inputs, with its reserves restored, is a harder result. The difference between those harvests is as important to the mission as the difference between one division and a lineage. We can set up both tests on Earth, in separate laboratories that start with declared materials and record what they use. The first fleet could consist of those laboratories.

From the First Watt to the First Report

What each ship can attempt

Imagine the eleven craft lined up for departure. All eleven can, in principle, find their way into an operating orbit, survey a proposed worksite, and send a record of what they find. Seven also carry equipment to establish and test a contained microbial lineage. Three of those seven carry what they need to attempt the larger renewal trial. At each of the three nearby systems there are two chances to demonstrate a lineage and one chance to demonstrate habitat renewal. The 61 Virginis craft adds a distant lineage attempt. The TRAPPIST-1 craft adds a distant survey. Each ship has the equipment required for its own assignment and carries the responsibility for returning its own evidence.

Those assignments make failure easier to read. A ship that measures a worksite from a sustained orbit can return a successful survey even if it finds the site unsuitable. A ship that misses capture may still collect useful measurements on its way through. Its flyby record belongs in the account as partial evidence. A chamber that contains a growing lineage can report where the founder came from and what the descendants kept doing once the machine stepped back. A habitat can earn a further result by replenishing the usable stocks needed to do the work again. Chapters 30 through 32 follow one craft through arrival and then ask how the different assignments add up across the fleet.

Waking in time to brake

In the worked Alpha Centauri A case, a ten-tonne craft coasts for about ten thousand years. The first useful power comes from a fixed collector. At ten astronomical units, the illustrative two square meters supply about seven watts. A small controller must use that power to establish the conditions for the next step. The approach plan puts the sail to work at about that distance. A restart that takes too long can spend the useful part of the approach even if the workshop eventually comes to life.

We can test the order of those operations on Earth. Seal the cold assembly with only the power and connections it would have in flight. Bring up the light along a modeled approach and measure when the collector can sustain the first controller, when the controller can read a usable archive, and when it can command and verify a deployment. Repeat with a missing archive copy, a noisy attitude estimate, a

reluctant actuator, or a collector delivering less than its assigned power. Record the latest useful time for each recovery. A rescue that depends on a technician touching the box has located a job the flight design still has to assign.

Separate specimens can probe material aging, radiation damage, seal integrity, and wear during intermittent operation. The combined trial can reveal which of those failures propagate across an actual restart. A more severe exposure can change the failure mechanism, so the extrapolation must follow the mechanism that the test actually exercises. The useful result is a record of which links in the chain recovered under specified conditions, along with the evidence needed to extend that result to a longer dormant interval.

The part of the sail that buys an orbit

The capture sketch asks a reflector covering 16 square kilometers to deploy, steer, and cut off the braking force. The mass allocation for the whole sail system is 1,500 kilograms. Spread evenly, that is about 0.094 gram per square meter for the reflector, supports, deployment, control, and reefing together. Near the end of the inward leg, the craft must reduce the force enough to preserve the orbit it has just earned. A test of the film establishes one property. A test of a deployed system must also account for its shape, how the load reaches the spacecraft, and how it behaves when commanded to change attitude.

Consider a shortfall in the braking work. In the simplified energy account, the incoming craft has about 8.12×10^{13} joules of positive orbital energy. The target ellipse requires removal of about 8.90×10^{13} joules. If the deployed system removes only 90 percent of that target, the craft still has about 1.05×10^{12} joules of positive orbital energy. For ten tonnes that corresponds to an excess speed of about 14.5 kilometers per second at great distance in the same point-mass model. That small change in the assigned work changes the outcome from capture to passage through the system. A real trajectory must include the timing and direction of the force, the binary companion, and the separate powered transfer into the working orbit. A near-space demonstration can test deployment, pointing, force measurement, and cutoff in sequence. The full route still needs a model that carries the demonstrated behavior into the target system.

The part of the record that gets home

From the working orbit in the example, a transmitter one astronomical unit from Alpha Centauri A appears at most about three-quarters of an arcsecond from the star

as seen from Earth. The estimated signal is about 0.8 detected photons per second under the stated link allowances. Recovering that signal requires a receiver that can point, keep time, and distinguish events associated with the transmission from events caused by the star, the sky, or the detector. The collected signal alone leaves the background and the usable observing time to be measured.

An early optical trial can give a second team a bright source near the apparent position of a weak, deliberately coded beacon. Vary the separation, inject interruptions, and ask the team to return a packet and its confidence in the identification. The onboard test should do its part as well. It should build a short report that a reader can interpret by itself, then preserve the larger evidence package for later passes. At the worked average of 0.01 bit per second, even a million-bit record takes about 3.2 years to send before its travel time is added. Every field assigned to the short report has to justify the observation time it costs.

Start with a habitat we can measure

The one-person chapters use the flows needed to support one adult as a scale for Earth experiments. A caloric load gives the crop system a demand that can be applied day after day. A reserve floor gives it a failure condition that a photograph cannot show. An unoccupied trial can reproduce the material demands with measured inputs and withdrawals, then ask whether a second harvest can be grown on recovered inputs while reserves are restored. Run that work in parallel with the cold-start and optical tests. Each can produce evidence before anyone has to commit to a crossing. The results converge when an independent team can trace a complete sequence from a declared inventory, through a working system, to a record it can judge.

1 The Cargo That Must Stay Alive

Two ships aim for the same star by different approaches. A third approaches another member of the same system. One leaves behind a working radio but cannot hold its orbit. Another stays and wakes, then discovers that its main archive is unreadable. The third has enough of its own power, records, and machinery to test a sample of local material. None can call an operator to fix the other two. Their histories have separated, and the mission can continue through the ship that remains.

Elsewhere the same departure sends workshops toward other stars. A second system gives the mission a chance to find different materials and operating conditions. A still longer voyage may have no bearing on the result at the first star, but it can carry the experiment to a kind of system the nearby missions do not sample. This fleet will have to be designed for more than one result. A single success can establish possibility. A comparable result from another workshop can help establish how much of that success depended on the first site.

At any successful destination the workshop finds an archive made by people who knew where it was going. Somewhere in the archive is a description of an organism, including its sequence, the structures that hold it, and a program for preparing an environment in which it can begin to live. The receiver can ask whether the description arrived intact. The deeper question starts when the answer is yes.

Now compare a ship crossing the same space with a different kind of cargo. Its cargo includes living organisms. Some continue to respire and grow. Others have been put into a dormant state and are expected to wake. In the hold there is an inventory of water, air, nutrients, and structures that keep the cargo in a usable condition. The ship carries a history as well as a payload. Every failure of cooling, every excursion in chemistry, and every energetic particle that leaves lasting damage can change what arrives. A complete record of the original cargo can be consulted, but a record alone will have little use if there is no way to rebuild what the ship lost.

Both ships need propulsion, navigation, and a way to remain functional. Their obligations to the cargo differ. The living ship must preserve a viable biological route through the crossing. The receiver ship can postpone the beginning of that biological route until the destination. In return it takes on a harder construction problem. It must receive enough information, possess enough capability, and have access to enough matter and energy to make something that can live.

This is the central experiment in interstellar panspermia. Can the fragile, working state of a living system be made at the far end of a signal? If so, which parts of that state must travel as instructions, which must be carried in the receiver, and which can form from local conditions? Radiation makes the choice urgent. It also makes it testable. We can ask what survives a specified radiation history, and we can compare that result with the reliability of the receiver that would build the same living system after arrival.

In 1973, Francis Crick and Leslie Orgel called the deliberate dispatch of organisms to another world *directed panspermia*. They considered whether someone elsewhere could have sent life to Earth and found the evidence inadequate to judge the likelihood (Crick and Orgel 1973). Their question was about what might already have happened. Ours is about what a sender could attempt and what an observer could verify. Their dispatch carried living organisms. The fleet in this book carries instructions and depends on machines that can put them to work at their destinations.

That change brings the place of origin into focus. Under the receiver plan, every atom of a newly built organism could come from local feedstock. Some of the arrangement would depend on a message from Earth, some on knowledge carried by the receiver, and some on the forces at work where it builds. We could trace those contributions even if no molecule from the organism's ancestor ever crossed the intervening space.

The target and the meaning of survival

Proxima Centauri is a useful distance marker, about 4.25 light-years away. The discovery of a planet candidate in its temperate orbital zone made it a concrete destination for discussion (Anglada-Escudé et al. 2016). An orbit is one piece of information about a planet. In the crossing calculations the star fixes a distance. The fleet has several other destinations, and even the Alpha Centauri group can consider different stars for power and capture. A receiver could operate in a manufactured enclosure without relying on the surface conditions of that planet. Even its choice of where to operate belongs in the mission design.

At one tenth the speed of light, an ideal coasting crossing of 4.25 light-years takes 42.5 years in the frame of the endpoints. At one hundredth, it takes 425 years. Those numbers exclude acceleration, braking, commissioning, and the wait for a suitable place to begin work. In the slower reference case, the same distance takes ten thousand years at a mean coasting speed of about 127 kilometers per second. Much of the question of what to send concerns what the payload has to do during that time.

For a microbial spore, survival might be defined by its ability to germinate and start a line of descendants. For a community, survival could mean recovery of all the functions that allow its members to reproduce together. For a human, it could mean the safe functioning of an entire body with its long-lived tissues and associated microorganisms. Those are three different endpoints. When a person says that life cannot survive a crossing, the relevant next question concerns the type of life, the radiation spectrum, the material surrounding it, the crossing time, and the test of recovery.

A single viable spore would settle the narrow claim that a spore could survive those conditions. If a mission also needed producers, decomposers, and partners that exchange essential compounds, the spore would settle very little of the wider claim. A community can lose a function when a lineage disappears or when its members recover at different rates. Arrival at one instant and the subsequent capacity to renew the system are different observations. A biosphere intended to keep a human alive must pass both tests.

What has already crossed a little space

One experiment gives a useful check on the premise. The Long Duration Exposure Facility carried *Bacillus subtilis* spores in low Earth orbit for nearly six years. Some multilayer samples protected against solar ultraviolet light showed survival as high as about eighty percent. The study described cosmic radiation exposure of about 4.8 gray among the tested conditions. The exposures, protections, and sample arrangements belonged to a specific experiment in low Earth orbit (Horneck, Bücker, and Reitz 1994).

A result like that is important precisely because it limits the claim. A dry spore carries its own specialized protection and has very different demands from a continuously functioning tissue. Its repair machinery becomes available again when it germinates. An experiment on spores aboard an Earth-orbiting facility gives a measured example of survival. A millennial flight beyond the solar environment has its own particle spectrum, exposure geometry, and storage risks. The mission has to earn an answer for each kind of cargo and each expected endpoint.

There is a similar distinction between a viable spore and a useful spore. An organism that starts to divide after retrieval can still carry changes acquired during the crossing. A recovery test can compare the descendants with the intended lineage and ask which functions remain. This matters if the original reason for carrying the organism was a particular ability, such as returning a nutrient to a form another member of a community can use. A colony counted after exposure answers one question. Its contribution to an engineered habitat answers another.

The biological case for a message is therefore strongest when the required working state is intricate, poorly tolerant of damage, or especially expensive to maintain during transit. The case becomes weaker for a compact, rugged seed that can reliably recover using material already at hand. A serious mission would compare both, including mixed plans. A preserved biological sample can be a reference for a receiver. A receiver can carry data to check which features of the sample still work. That comparison gives radiation its proper role in the argument. It can impose a cost large enough to change the choice of architecture.

A second way to define the cargo

One way to understand the distinction is to ask what happens when a single item is lost. If a packet of digital information fails a check, a second copy may be able to restore it. If an organism loses the very structures needed to read its genome or repair its damage, its surviving DNA is a much less direct path back to a growing organism. There might be another viable organism in the package. There might also be a builder capable of reconstructing the missing structures. Each possibility supplies a different kind of redundancy.

The builder can hold a small inventory of carefully selected physical machinery and use a large inventory of raw material at the destination. In that case the mass launched from Earth need not include every kilogram that will eventually become a habitat. It still has to include the equipment needed to acquire, purify, organize, and test the local material. If the builder turns out to require a substantial living factory as part of its initial equipment, the original problem has moved inside the receiver. Counting that contribution is part of the experiment.

There is a practical first question that requires no interstellar flight. Give a specified receiver a description of one ordinary bacterium and ask whether the receiver can make a lineage from the starting materials it claims to require. Give a second copy of the description to an independent team. Compare where it succeeds, where it fails, and what both teams had to supply in advance. The rest of this book extends that question outward. The chromosome is surrounded by a working cell. The body is surrounded by a working environment. Each layer adds things that are hard to maintain aboard a ship and hard to make at a destination. The crossing motivates the work because it puts a price on choosing where and when to build each layer.

2 Choosing the Stars

On Earth, an observatory can tell the fleet where a star is and what kinds of objects orbit it. The receiving workshop will need to know where it can park, how it will get power there, and what it can carry back from a mining trip. A planet in the habitable zone answers none of those questions by itself. A disk of dust answers a few of them indirectly. We choose stars for a fleet by asking what each workshop could learn on arrival and what it could do if the first site it surveys proves unusable.

Three nearby systems and two long shots

The reference fleet has eleven independently equipped craft. Three go to the Alpha Centauri system, three to Barnard’s Star, and three to Epsilon Eridani. Two others go farther, one to 61 Virginis and one to TRAPPIST-1. Each is assigned a common survey capability and must qualify its own way to wake, stay, and work without help from the other ships. Three launches give a system several chances at those steps. The two singletons buy a chance to investigate very different systems, at the cost of losing local redundancy. Eleven is a design choice for this book. A different launch budget or new observations could change the allocation.

Destination	Craft	Present distance, light-years	Coast at 127.4 km/s, years	What earns a place
Alpha Centauri system	3	4.35 to A/B, 4.25 to Proxima	~10,200 to A/B, ~10,000 to Proxima	Two relatively bright nearby stars, with a separate red-dwarf option
Barnard’s Star	3	5.96	~14,000	Nearby single star and a test of operation around a faint host
Epsilon Eridani	3	10.5	~24,700	A resolved debris belt to guide a survey for feedstock
61 Virginis	1	27.7	~65,000	A mature Sun-like star with a resolved belt and known planets
TRAPPIST-1	1	40.5	~95,000	Seven terrestrial-size planets around an ultracool star

All eleven craft are assigned to survey. Seven can also attempt to establish a verified renewing lineage in containment, two at each nearby system plus 61 Virginis. Three can attempt the renewal trial, one per nearby system. The TRAPPIST-1 craft is assigned to survey only. A human settlement would require a separate program for development, care, and continuity. Chapter 32 carries the three fleet assignments through separate probability calculations.

The source of a founder is a separate item on the record. A craft can qualify a lineage by recovering a dormant founder or by constructing one under a declared inventory. Only a documented build that meets the nonliving start requirement answers the first-cell handoff question. Chapter 31 compares those routes and Appendix E shows how to report them.

The distances are rounded from present observations (NASA Goddard Space Flight Center n.d., NASA Exoplanet Archive n.d., Booth et al. 2023). At the reference speed, each light-year takes about 2,353 years. The coast column leaves out launch, braking, and stellar motion. The flight plan needs a date, a prediction of where the star will be on arrival, and an approach with enough room and time to brake.

Barnard’s Star shows how much the correction matters. Today it approaches the Sun at about 110 kilometers per second along our line of sight and moves across it at about 90 kilometers per second. Imagine a ship starting at the Sun and coasting in a straight line at 127.4 kilometers per second, with the star also moving at constant velocity. The ship must match the star’s sideways motion. That leaves about 90 kilometers per second of the ship’s velocity along the original line of sight. The two close at roughly 200 kilometers per second. The first intercept in this sketch takes about 8,900 years, and the ship arrives at about 200 kilometers per second relative to the star. The shorter crossing comes with a harder brake. Each mission needs its own trajectory and arrival budget (NASA Exoplanet Archive n.d., SIMBAD n.d.).

Where the nearby workshops might go

Alpha Centauri. The preferred operations are around A and B. They give collectors two relatively luminous stars to work with, while the binary also restricts the stable orbits around each star. A favorable working orbit and a reachable supply of solids have to be demonstrated. Webb observations have yielded a candidate giant planet around A. The initial image and the follow-up nondetections constrain possible orbits. The fleet can survey those orbits to see what is there. Proxima is a conditional alternative. It lies thousands of astronomical units from the binary and has produced intense observed flares. A craft assigned to Proxima needs its own

arrival geometry and power and radiation budgets. The fleet counts A, B, and Proxima as one allocation. Each craft plans its work around the star it can reach directly (Kervella, Thévenin, and Lovis 2017, MacGregor et al. 2021, Beichman et al. 2025).

Barnard's Star. A nearby single star gives the fleet another approach and a different power environment. Four small planets detected by radial velocities complete their orbits in roughly two to seven days. They show there is material organized into planets. Their tight orbits give little reason to plan around temperate surfaces, and the detections say little about smaller, accessible bodies. Barnard's faintness means a collector of a given size must work much closer to the star for a given incident flux. The workshop needs to check whether that orbit, the brake, and the routes to candidate materials fit the same mass and energy budget (Basant et al. 2025).

Epsilon Eridani. ALMA has mapped its outer debris belt, with much of the emission concentrated around 70 astronomical units from the star. The map helps the workshop plan where to send a surveyor and where to avoid dust during a fast approach. At 70 astronomical units even a Sun-bright star supplies only about one five-thousandth of the sunlight at Earth. The part of the system that offers the best place to gather material may be a poor place to run a solar workshop. A receiver could establish a power base closer in and send a surveyor outward. Whether that pays depends on the size and composition of the bodies it finds and the cost of bringing useful material back (Booth et al. 2023).

What the two long shots buy

61 Virginis offers an older, Sun-like comparison with Epsilon Eridani. Observations with Herschel and at millimeter wavelengths resolve a broad outer debris disk, and radial velocities identify planets closer in. A workshop would investigate whether that disk has reachable bodies with the elements it needs. The known planets help map the system's architecture and give the survey landmarks. The receiver can do its work without landing on one (NASA Exoplanet Archive n.d., Marino et al. 2017).

TRAPPIST-1 offers the opposite experiment. Seven terrestrial-size planets orbit close to a very faint star. The system is an unusually rich place to compare how rocky worlds develop under related conditions. Atmospheric measurements are still an active investigation. The first Webb spectra of planet e left multiple possibilities open, and a newer observing program uses planet b to help separate the star's

influence from e’s spectrum. If a destination might host native biology, the scientific reason to observe it also gives the workshop a reason to keep Earth life inside its own enclosure. The craft can return a useful survey even if it never starts a culture (Gillon et al. 2017, Glidden et al. 2025, Allen et al. 2026).

What earns the first build

At every star the receiver must demonstrate a stable place to work, reliable power through foreseeable interruptions, and access to feedstocks it can turn into tested materials. The useful question about water is how much the machine can deliver to its tanks for a given expenditure of energy, equipment life, and time. The same accounting applies to carbon, nitrogen, phosphorus, metals, and shielding. A prospect can fail because its materials are too diffuse or its transport route costs too much. The machine can report that result and preserve its archive while it looks elsewhere.

A ship that succeeds may eventually send its maps and procedures to the other receivers. Every receiver can act on arrival whether or not it has heard from another ship. At the reference speed, the two outer flights require archives and hardware qualified for much longer dormant periods than the nearby flights. A faster version would need a new launch and braking budget. The first local successes could happen tens of thousands of years before the last craft begins its survey. The fleet can count those successes when they happen. Its nearest workshops have work to do on their own arrival, and its farthest ones carry questions that remain worth asking if the nearby workshops thrive.

3 What Another Ship Buys

Eleven uncrewed craft leave with a common survey obligation. Each is assigned a cold restart, a safe working location, and an independently interpretable survey report. Selected craft are assigned the harder work of establishing a renewing microbial lineage or qualifying a contained habitat. The example in this chapter compares allocations under a single assumed endpoint. Chapter 32 assigns different milestones to different craft and carries their probabilities separately.

For a reference allocation, send three craft each to the Alpha Centauri cluster, Barnard's Star, and Epsilon Eridani. Send one each to TRAPPIST-1 and 61 Virginis as longer shots. These are five destination branches for the arithmetic. Grouping the Alpha Centauri craft in one branch assumes a shared site constraint in the model. The actual mission would need to identify which constraints two craft share even when their proposed worksites differ. A group of three has the job of surviving individual losses. A distinct destination gives the experiment another set of local conditions to work with.

Independent chances on one route

Suppose a destination can support the planned enclosure and a common flaw has already been ruled out. Give each of the craft on that route the same chance p of delivering the accepted result. In this example their remaining failure events are independent. One craft fails with probability 1 − p. For n craft all to fail, that event has to occur n times, which has probability (1 − p)^n. The chance that at least one succeeds is the complement.

Chance of at least one success = 1 − (1 − p)^n.

Set p to one-half for a sensitivity case. One craft gives a 50 percent chance. Two give 75 percent, three give 87.5 percent, and four give 93.75 percent. The next copy makes a difference only when all the others fail and the new one succeeds. In this example that margin falls from 25 percentage points for the second craft to 12.5 for the third and 6.25 for the fourth. An identical copy can therefore have a lower marginal value even when the probability assigned to each craft stays the same.

The probability p includes the route, restart, commissioning, construction, and validation conditional on the stated shared requirements. An 87.5 percent chance of receiving a radio signal would tell us little about a second harvest if most functioning radios were attached to machines that could never construct a

workshop. The assigned p of one-half makes the arithmetic visible. It comes from no measurement of an interstellar ship, a target, or a biological build.

Three ways a branch can close

A particle strike can ruin one machine while another on a separate trajectory continues. A resource assumption can fail for every machine using the same target. A fabrication step can be wrong in every machine using the same design. These losses act at different scales. A second machine on the same route can cover the first kind. A second destination can cover some of the second kind. Correcting or independently implementing the fabrication step addresses the third.

Let q be the assigned chance that one destination meets the shared requirements for the experiment. Those requirements might involve access to a workable energy supply, usable materials, and an acceptable site for containment. A branch with n craft has a success chance s given by

s = q × [1 − (1 − p)^n].

For q = 0.5 and p = 0.5, three craft give s = 0.4375. Sending enough copies can push the conditional craft result toward one. The whole branch remains bounded by q because every copy faces the same target requirement. If a site survey discovers a different enclosure design that can work at that site, the receiver may have changed the relevant requirement. That is a change in the plan rather than another draw from the old model.

For distinct branches, calculate each branch's failure chance as 1 − s. If the branches are independent once the common architecture works, multiply their failure chances and subtract that product from one. Finally let g be the probability that the architecture passes a mission-wide requirement. A fatal defect in that shared architecture would stop every branch in this model. Multiply the conditional fleet result by g to include it. With g = 0.8, even unlimited copies of the same architecture cannot raise success beyond 80 percent. Here g is a stipulated sensitivity parameter. In practice a common hazard can affect some branches or operating states rather than all of them. The applicable calculation then needs those separate cases and their conditional chances.

NASA reliability work describes how dependencies can defeat redundancy when several units share a design, manufacturing, control, or environmental vulnerability (Jones 2025). A fleet can spend its extra capacity on removing some of those

dependencies. The separate curves show what replication alone buys if the shared assumptions remain fixed.

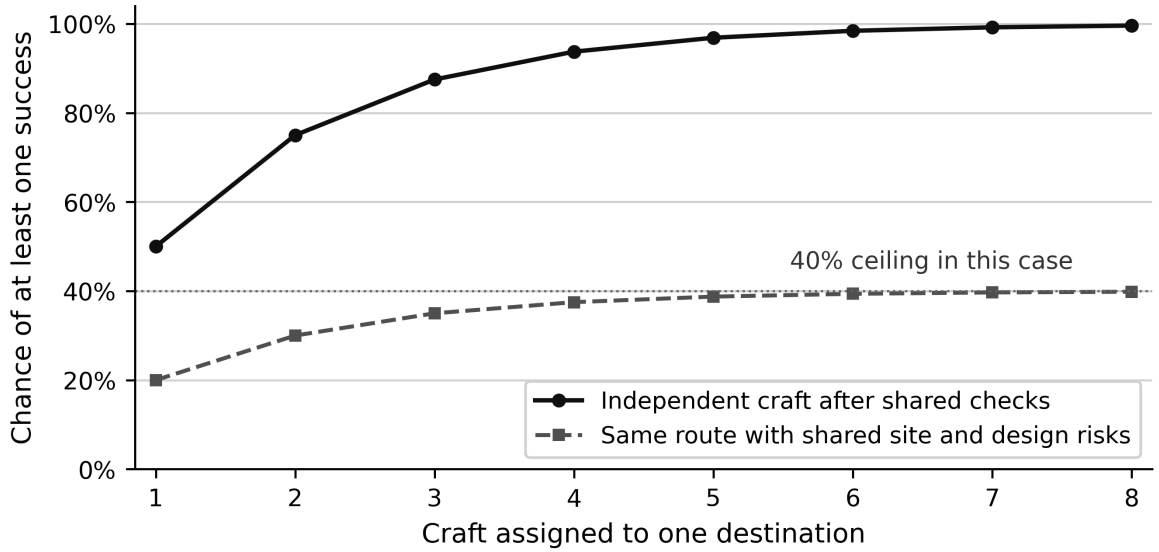


Figure 1. Illustrative odds for one destination with p = 0.5 per craft. The lower curve also applies a site factor of 0.5 and a shared design factor of 0.8.

Suppose all the receivers carry the same omission in their instructions for making a cell. They may each assemble the chromosome and fail to establish the working chemical state that allows it to function. A different star provides no independent attempt at the missing step. The same problem arises if every habitat controller measures total phosphorus and misses the transfer into a pool the next crop cannot use. Each machine might record the same reassuring number while the useful stock declines. An independent validation path checks the cell's continued operation and measures the pools that matter at the following harvest. That path can identify an error shared by otherwise separate ships.

Two allocations of eleven ships

Give each of the three nearby branches q = 0.5 and each of the two longer branches q = 0.2. Use p = 0.5 per craft after the relevant site and design requirements have been met. For this example each craft consumes the same unit of the launch budget, and the branches are conditionally independent. A comparison between three craft at each nearby destination plus the two single longer shots, and four, four, and three craft at the nearby destinations, then holds the fleet size fixed. The model assigns the longer shots a smaller chance of meeting their shared site requirement, although both allocations have the same stipulated craft probability.

Allocation of eleven craft	Conditional chance from the nearby branches	Conditional chance after all branches can qualify
3, 3, 3 nearby plus 1, 1 farther out	82.20 percent	85.58 percent
4, 4, 3 nearby	84.12 percent	84.12 percent

For the first allocation the three nearby branches each have a failure chance of 0.5625. Each long shot has success chance 0.2 × 0.5 = 0.1 and failure chance 0.9. The combined failure chance is 0.5625 cubed times 0.9 squared, about 0.1442. For the second allocation each nearby branch with four craft has failure chance 0.53125. The failure product is 0.53125 squared times 0.5625, about 0.1588. Multiplying the resulting eventual success chances by the illustrative g of 0.8 gives about 68.47 and 67.30 percent. The value of the two additional destinations is small under these assignments, and a change in their assumed chance can reverse the ranking.

When the two longer branches have the same success chance, each would have to offer greater than about 5.56 percent for the first allocation to exceed the second in this model. This follows by comparing the 0.9 factors in the first failure product with the reduction in nearby failure from changing 0.5625 to 0.53125. The ratio is 17/18, so the break-even success chance is 1/18. If both long shots had p = 0.25 with q still 0.2, their branch chances would each be 5 percent. The same eleven ships would then favor the nearby allocation under the stated objective and assumptions. An actual comparison would also assign costs to separate routes, communication support, shielding, development of independent designs, and the duration of the commitment. Eleven launches describe one budget constraint.

The deadline belongs inside the calculation

Imagine a reporting deadline by which the nearby branches could have qualified, while the two longer craft are still in transit. At that cutoff the longer ships contribute zero to the chance of an accepted result. In the simplifying case where the nearby craft have already had their full assigned opportunity, the first allocation gives 82.20 percent conditional on the architecture working and the second gives 84.12 percent. The order reverses compared with the eventual comparison in the table. A realistic calculation would attach a distribution of arrival and test completion times to each branch. Braking, commissioning, cell tests, growth, replenishment, and any required return of evidence all take time. If acceptance is evaluated on Earth, the travel time of the report belongs on the clock too.

At each date we can repeat the complement calculation with the chance that each eligible craft has reached the specified endpoint by that date. Before a craft can physically reach that endpoint its contribution is zero. The resulting cumulative schedule describes the chance of receiving at least one qualifying result by successive dates. It gives the project a basis for distinguishing an early demonstration from a later expansion of options. The two objectives can favor different uses of the next ship.

Put each redundancy where it can act

Separate launches and hardware batches can reduce exposure to a failure shared by one launch or production lot. Different approaches and worksites can reduce dependence on one local hazard if the alternative routes and sites have been measured. Independent implementations of critical fabrication, measurement, and control steps can expose errors shared by a copied design. Each measure needs a test that identifies the failure it addresses. Adding copies of an untested claim about the bath adds little evidence about whether the first cell will keep working. Adding copies of a crop room that uses one failed nutrient return adds little capacity to reach the second harvest.

The flight units can carry a common reporting contract even when their designs or routes differ. That contract names what the receivers measure, how they record reserve loss and recovery, and which observations authorize a contained trial. A mission gains evidence from failures it can place in the causal chain. A craft that reports a successful restart and a failed propagation test has told us which half of its promise held. A second craft with a different propagation method can test a specific alternative. The fleet's value comes from the attempts it keeps independent and the questions it makes possible to answer.

4 The Ships That Came Before

Vernor Vinge gave an onboard intelligence named Ilse ten thousand years to reach Alpha Centauri. In his 1972 story “Long Shot,” her ship carries embryos. She manages the crossing and chooses a place to land. As components fail, she loses memories. Near the end, a backup has to tell her what the voyage was for. The engineering problem in the story begins with an ordinary question. When the ship arrives, which of the things it knew when it left can it still put to use? (Vinge 1972, Hein and Baxter 2019)

The idea that a machine could reach another star and establish life there has a long history. Some proposals put cells in the hold. Others put descriptions of cells in the archive and a laboratory beside it. The choice determines where the work of keeping life possible takes place. It also determines what the machine must be able to do on its first day of work.

From carrying life to making it on arrival

Francis Crick and Leslie Orgel proposed in 1973 that a technological civilization could deliberately send organisms to a suitable world. They asked whether life on Earth might have begun that way and judged the available evidence insufficient to tell. Michael Mautner later asked what an actual seeding program might send. In a 1997 study he proposed small packets of microorganisms carried by solar sails, with material in dense star-forming clouds slowing and capturing the packets. He considered speeds around five ten-thousandths of the speed of light. At that speed a straight coast over 4.25 light-years would take about 8,500 years. That comparison puts the millennia of this book in perspective. Mautner’s target and method of capture were different, and his microbes had to arrive viable. He developed the broader case for deliberate panspermia in his book *Seeding the Universe with Life* (Crick and Orgel 1973, Mautner 1997, 2000).

Claudius Gros made the decisive change in the Genesis Project. In 2016 he proposed a robotic spacecraft that would carry a genome archive and a laboratory to synthesize microbes at the destination. It would investigate a planet, assess its suitability, and have the means to seed it with a mixture of microorganisms suited to the local conditions. Gros explicitly anticipated a future ability to synthesize life from scratch. In a later paper he considered travel times of a few millennia for such missions, with a magnetic sail using the sparse interstellar plasma to brake. A long

voyage, a machine in charge, and the manufacture of life after arrival already belong together in that proposal (Gros 2016, 2019).

Genesis asks what happens if a temporarily habitable planet receives a head start on biological evolution. A successful release would put the next chapter in the hands of natural selection and the planet's chemistry. The receiver in this book starts in a chamber it can sample and isolate. Its first success is a lineage it can reproduce in a second build. Later it has to sustain and repair an ecology in a bounded habitat. Both projects depend on making life at the destination. They put the boundary of the experiment in different places.

The factory and its first generation

Andreas Hein and Stephen Baxter described a wider class of missions in their 2019 paper "Artificial Intelligence for Interstellar Travel." They called one type the Founder. It surveys resources, builds facilities, and checks whether its interventions actually work. A Founder could manufacture habitats in space or on a planetary surface. Their paper also considers a converter that makes DNA from digital information at the destination and advanced bioprinting that might eventually make whole organisms there. One proposed route would carry stem cells to reprogram. Another would use a factory already at the destination to build a converter whose design could be sent from Earth. This is a precedent for putting the capacity to learn and build at the far end of the journey (Hein and Baxter 2019).

Fiction got to the first generation early. In James P. Hogan's *Voyage from Yesteryear*, an automated ship carries human genetic information electronically and uses onboard materials to produce embryos. Its machines build farms and factories and raise the children. Arthur C. Clarke's *The Songs of Distant Earth* puts the machine in charge of making and raising the first generation on Thalassa. Clarke's ship also has to provide the living community that sustains its settlers. The children need tools, education, and years of care. In these stories the first moment of life opens a far longer obligation. A habitat that can support an adult for a year would tell us little about whether machines could safely raise a child or whether a community could persist across generations (Hogan 1982, Clarke 1986, Hein and Baxter 2019).

The laboratory analogues help separate the obligations. Kent Boles and colleagues built an automated converter that used digitally transmitted sequences to produce DNA, RNA, proteins, and viral particles. Daniel Gibson and colleagues had already shown that a chemically synthesized bacterial genome could take control after

transplantation into an existing cell. Roy Sleator and Niall Smith drew on minimal bacterial cells and synthetic biology in proposing how a seeding mission might tailor organisms to another world. These are concrete steps from a file to biological function. The recipient in a genome transplant also makes plain how much organized machinery a functioning cell contributes at the start. The receiver imagined here has to supply that organization and show that the resulting cell can produce descendants (Boles et al. 2017, Gibson et al. 2010, Sleator and Smith 2017).

The star has two jobs

The destination star may supply both the first electrical power and a way to brake. These uses of starlight ask different things of the ship. A fixed collector on a favorable approach could deliver a small current to a protected controller. The controller has to wake early enough to check power, recover the archive, establish attitude, and prepare the maneuver that keeps the ship in the system. Those checks must fit the approach trajectory. An electrical system that wakes after closest passage might work perfectly and still miss the mission’s only chance at capture.

René Heller and Michael Hippke investigated how photon pressure and gravity could slow a very light interstellar sail and leave it in a bound orbit. Their calculations show how strongly the result depends on the mass carried per unit of sail area. A shielded workshop and its equipment change that ratio. The prior work gives us a physical mechanism to evaluate against the mass and approach speed of the ship. In the architecture developed here, capture gets the ship into a bound orbit. A subsequent spiral is a way to move between bound orbits if the ship has a means of exchanging energy and angular momentum (Heller and Hippke 2017).

An experiment before a voyage

The first construction site can be on Earth. Give a receiver a defined archive, a stated inventory of machinery and feedstock, and a power source it must bring into service from a cold start. The archive specifies a genome and the conditions and operations needed to assemble its first cellular bath. List every dependency of the receiver, including any biologically prepared materials, so that a reagent made elsewhere cannot quietly stand in for a step assigned to the machine. Ask for a line of descendants and repeat the build with independently prepared inputs. The first result worth sending to a star is a result another receiver can obtain.

The next test moves the boundary out from the cell. The receiver builds a small ecology, measures what enters and leaves its air, water, food, and waste streams, and demonstrates recovery after a specified failure. A habitat that works while every culture is healthy and every pump is available gives a measurement of normal operation. Removing a culture or interrupting a loop tests whether the receiver can diagnose and restore the system from the stocks and instructions it actually has. Start with one person’s resource demand as an accounting unit. The balance sheets and recovery trials have to close before a human settlement becomes a credible design problem.

The earlier ships clarify what that experiment must cover. Genesis gives us the laboratory that makes organisms after the crossing. The Founder gives us a machine that can test its own work. Ilse gives us ten thousand years in which a perfectly sensible design can lose the knowledge it needs at arrival. The question for the ground demonstration is whether a cold receiver can recover that knowledge and turn it, step by measurable step, into a living system that can recover from losses of its own.

5 Ten Thousand Years in the Dark

At the departure star the collectors can unfold and the ship can draw power while it is being checked. On the other end of the crossing it can do the same. Put the two stars 4.25 light-years apart and most of the route lies outside either useful power neighborhood. The ship I want to examine spends that interval cold, with no viable cells, spores, seeds, embryos, or crew aboard. It carries an archive, a receiver and workshop, and a declared inventory of nonliving materials. A small system has to know when to wake. The equipment that makes decisions and does biological work can wait until there is power to run it.

Take ten thousand years as a reference coasting time. The mean speed is about 127 kilometers per second, or 0.000425 of light speed. It is slow by the standards of many interstellar proposals and fast by the standards of an orbital arrival. The launch and braking phases add their own times. Even this slower ship carries roughly eight billion joules of kinetic energy per kilogram at cruise. A ten-tonne vehicle carries about 8×10^{13} joules, or 22 gigawatt-hours, in the departure frame. The comparison fixes a scale for the momentum and energy problem. It supplies neither a launch system nor a brake.

Now put eleven ships of the same illustrative ten-tonne mass on that energy account. Their combined mass at cruise is 110 tonnes, with approximately 8.9×10^{14} joules of kinetic energy, or 248 gigawatt-hours, in the departure frame. This is energy of the moving vehicles alone. It excludes propulsion inefficiency, propellant or beaming losses, and the demands of braking each vehicle. The fleet multiplies a real freight and launch burden. Using several departure windows and designs may reduce some shared hazards while adding commissioning work. Those costs belong beside the improvement in odds.

The long interval is the reason to discuss an AI-first mission. A viable organism can carry mechanisms for its own repair and reproduction. A self-supporting biological colony could eventually make machines. To recover a technological civilization after arrival, that colony would have to recreate a chain of knowledge, labor, materials, and institutions. An AI and a workshop can be tasked directly with restarting that chain. A workshop that can build living systems still has to be demonstrated, and a workshop that survives ten thousand years has to be demonstrated separately. The order of arrival is a claim about the work we want done and the dependencies each payload brings.

Biological dormancy and machine dormancy

Dormancy gives a spore a different problem from the one a crew faces. It can wait with little active metabolism, then use the chemistry available during germination to repair some damage. A crew or a managed ecosystem spends the crossing repairing bodies, handling wastes, and maintaining its physical plant. The loss of a food system or a tissue over that period can be as decisive as the failure of a pressure vessel. There is also a middle class of cargo, such as stored reproductive material, whose successful recovery depends on many operations taking place after arrival. Tests of radiation tolerance have to name both the cargo and the recovery endpoint.

Cold machinery can postpone many of its failures by avoiding moving parts and switching off components that would accumulate operating wear. The dormant machine is still exposed to energetic particles, impacts, aging of materials, and the thermal stresses of departure and revival. A redundant archive can correct some errors when checked. It cannot read itself if all copies depend on the same failed reader. A collection of spare actuators has value only if something that can still move is able to fit one. These are physical dependencies, and they can be traced as explicitly as a cell's dependence on a translation system.

Compare dormant biology with cold hardware at the same shielding mass, elapsed time, and mission endpoint. A living seed is the more economical cargo if it can be recovered with simpler equipment and supplies the required functions. Cold hardware has the stronger case if its protected records and machinery can be repaired and the necessary living state is costly to preserve. The comparison should include a mixed case. That case can tell us what we give up by requiring a lifeless reference mission. The requirement forces us to expose every imported contribution to the first cell.

The electric bill in the dark

Even a small continuous load becomes substantial when it runs for a hundred centuries. Ten thousand years contain about 87.7 million hours. One watt over that interval uses 87.7 megawatt-hours. A milliwatt uses 87.7 kilowatt-hours. A ten-watt computer uses 877 megawatt-hours. These are ideal energy deliveries with storage losses, degraded power conversion, and reserve excluded. The question for a very long cruise is first how often a function has to operate, then what means of powering it can remain usable for as long as required.

Call a ship cold if nearly all of its computational and fabrication capability remains dormant throughout normal cruise. That definition leaves choices about the minimal

functions. Navigation might require several short corrections, a protected timer or occasional attitude measurement. The ship might reserve a physical sensor that begins drawing useful energy when it sees a bright enough star in an expected direction. In each design the trigger, the pointing system, and the first switch must be listed as hardware whose survival is needed. A low average power requirement and a viable power supply are two distinct entries in the account.

A laboratory can exercise an analogue of this arrangement. It can separate the protected restart chain from the main computer, impose storage and component faults, then increase available light gradually and record which functions recover. It can run redundant readers after withholding one copy of the archive. Such exercises establish performance over the stress histories actually tested. Any extrapolation across ten thousand years still needs material evidence and an explicit account of possible shared failures.

6 Dose Along the Route

The first useful radiation number is a rate measured inside an actual spacecraft. During the cruise of the Mars Science Laboratory, the Radiation Assessment Detector recorded the energetic particle environment behind the spacecraft's shielding. The comparison published by Hassler and colleagues gives a tissue-equivalent absorbed dose from galactic cosmic rays of roughly 0.48 milligray per day and 1.84 millisievert per day of dose equivalent for the cruise. The instrument also registered solar particle events, which belong in a separate part of the accounting. The rate reflects the radiation field, shielding, and solar conditions of that particular flight (Zeitlin et al. 2013, Hassler et al. 2014).

A gray is a joule absorbed per kilogram. Dose equivalent applies a radiation-quality weighting to the absorbed dose to reflect how different radiation fields can produce different biological effects. A detector's dose equivalent has a different meaning from an individual's organ doses, lifetime health outcome, or effective dose under a specified protection standard. The absorbed dose is the more direct starting point for comparing the energy deposited in a sample. Even in gray, two exposures can differ greatly in their effects if one field leaves a very different pattern of damage.

Use the cruise rate as a yardstick and imagine that it holds throughout the 42.5 years to Proxima. At 365.25 days per year, the crossing lasts approximately 15,523 days. Multiplying gives about 7.45 gray of tissue-equivalent absorbed dose and 28.6 sievert of instrument dose equivalent. For 425 years at the same rates the arithmetic gives 74.5 gray and 286 sievert. Over ten thousand years that same multiplication gives about 1,750 gray and 6,720 sievert. This last number is a deliberately long extrapolation of the instrument rate. An actual survival prediction would require a separate exposure and recovery model. These products illustrate the effect of elapsed time. The actual interstellar field and a proposed shield could yield different rates. There are no measurements in this calculation from the route of the imagined ship.

The order of magnitude matters even with that uncertainty. A number that looks small when quoted per day can become consequential when multiplied by fifteen thousand days. At the same time, an annualized rate for one instrument cannot tell us when a person would become ill, how an embryo would develop, or how many bacteria would grow after storage. Applying a dose number to an organism requires a model or measurement for that organism, under the relevant exposure and

recovery conditions. Prolonged exposure and a brief irradiation can differ even when their totals agree.

What the energy does

A charged particle can deposit energy along a track, leaving a dense line of local damage in some cases. A different radiation field can distribute the same energy much more sparsely. One tissue may repair a fraction of the injury while the crossing continues. A dormant spore must wait to recover its active chemistry. The spore can carry DNA protection that helps it through the delay, and its capacity for repair during germination contributes to measured survival. Experiments with *Bacillus subtilis* spores exposed to X rays, protons, and energetic iron ions showed that removing some spore-core protections and DNA repair functions reduced survival. The changes depended on the radiation used (Moeller et al. 2014).

A survival curve for one sample offers a better comparison than a vague label such as radiation resistant. One experimental measure is the dose associated with a tenfold reduction in colony-forming survivors, often written D_{10}. It describes a specified strain and state, an exposure, and a recovery assay. Imagine fitting such a curve in a laboratory and using it to estimate the number of spores left after a crossing. The estimate would rely on whether the curve still applies at the flight's dose rates and spectrum. A second experiment could expose the whole stored package and follow its recovery, including the loss of useful functions that a simple survival count would overlook.

The cumulative injury problem becomes more complicated with multicellular organisms. The state that has to survive is distributed across many tissues and across processes that maintain those tissues. Even if a sample of cells could grow after exposure, that observation would address only one requirement for recovering a whole organism. A useful report would specify what is being preserved, how long it has to remain recoverable, and what a successful return to function looks like.

How much a shield has to weigh

Consider a rough geometry. Give a spherical compartment a radius of two meters and assign a uniform shielding mass of ten grams per square centimeter along its surface. That areal density is one hundred kilograms per square meter. The sphere has about fifty square meters of surface area. In a thin-shell estimate, the shielding weighs roughly five metric tons. At twenty grams per square centimeter it weighs roughly ten tons. With the same ten grams per square centimeter, increasing the

radius to five meters raises the shield estimate to about thirty-one tons. The square of the radius sets the first-order mass cost.

These numbers say how much material surrounds the compartment. To predict how much protection it offers, we would need the composition, shield geometry, and incoming spectrum. Particles can scatter, lose energy, or create secondary particles when they strike a shield. Measurements and particle transport models help estimate the dose deeper in a material. One calculation for the Martian subsurface illustrates the reason to ask for that estimate. After scaling a transport model to the measured surface field, Hassler and colleagues reported a higher modeled dose at a shallow depth in their chosen materials, followed by falling dose at greater depths. That particular Martian geometry supplies an example of competing effects in radiation transport (Hassler et al. 2014).

The geometry can change the trade between biological and digital payloads. A small vault that protects an archive or a handful of dormant samples takes far less shielding material at the same areal density than a room for an active community. A room that stays inhabited also has to remain thermally and mechanically stable, carry necessary equipment, and protect its working surfaces and passageways. Water and other supplies might serve more than one role in a design. In that case the same kilogram can contribute to storage and shielding, and the system has to function as the supplies move or are drawn down.

A fixed shield also affects propulsion. A kilogram moving at one tenth the speed of light has about 4.5×10^{14} joules of kinetic energy in the departure frame. This is the relativistic kinetic energy of the mass itself, without the energy required by an actual propulsion system. A five-ton shield would account for about 2.3×10^{18} joules at that speed. A braking system has to manage the kinetic energy at the other end as well. The numbers make a high-speed shield a major item in the mass and energy account, even when the biological design calls for only a modest-looking depth of material.

The gas in front of the ship

A fast ship faces another source of exposure. Ordinary interstellar gas approaches its forward surfaces at nearly the ship's speed in the ship's frame. At one tenth the speed of light, a hydrogen nucleus initially at rest in the surrounding gas carries about 4.7 megaelectronvolts in that frame. At one hundredth the speed of light it carries about 47 kiloelectronvolts. These values follow from the relativistic kinetic

energy of a proton. They are energies per incident particle, distinct from the absorbed dose in a particular part of the cargo.

That forward exposure depends on the gas column encountered, the energy loss in the shield, and the secondary radiation that follows. Dust adds impacts that can damage surfaces and sensors. Hoang and colleagues modeled gas and dust interactions for fast spacecraft and found strong dependence on the material and accumulated column (Hoang et al. 2017). Their model supplies a reason to include both kinds of impact in design, even for a probe that carries only machines. Its receiver can lose a lens, a radiator, or a power supply long before the storage device has reached its expected lifetime.

Speed therefore has two competing consequences in this example. Going faster shortens the time available for a background field to deposit energy and increases the impact energies of the ordinary material encountered along the path. To compare two speeds, put both histories into the same mission design and track the vulnerable parts of each payload. Biological survival, electronics reliability, and receiver readiness can then be measured under consistent assumptions. The claim that motivates this book becomes concrete when a specified living cargo fails a credible crossing test and the alternative receiver still delivers a verified build.

7 Catching the Destination Star

A spiral is a good picture of what a low-thrust ship can do once it has an orbit. In the simplest case a vessel already bound to the star points a small thrust component against its motion and gradually changes the size of its orbit. An incoming vessel has a different starting condition. It approaches on an unbound trajectory with positive energy relative to the star. If only the star's gravity acts, the energy gained on approach is returned on departure. The track bends, and the ship leaves. A capture design first has to remove energy, exchange it with another body, or otherwise arrange a bound trajectory.

In a two-body approximation the ship's speed at closest approach is the square root of the sum of the squares of its excess speed far away and the local escape speed. An instantaneous retrograde impulse would have to reduce that local speed below escape speed. This is an optimistic boundary for a particular geometry. A finite burn has its own trajectory, and a useful bound orbit needs margin beneath the boundary. A binary system or planet can exchange orbital energy with the spacecraft. The magnitude and sign of the exchange depend on the geometry and the velocity with which the spacecraft meets the body.

Approach to a Sun-like star	Local escape speed	Ideal capture impulse
1 astronomical unit at 127 km/s excess	42 km/s	About 92 km/s
0.1 astronomical unit at 127 km/s excess	133 km/s	About 51 km/s
0.1 astronomical unit after prior slowing to 10 km/s excess	133 km/s	About 0.38 km/s

The last row illustrates why the system might need separate braking stages. If the incoming ship is slowed well before its closest approach, the remaining arrival maneuver can be much smaller. At a tenth of an astronomical unit the higher escape speed is an advantage for a well-timed impulse. Intense light, heat, restricted maneuver time, and pointing constraints enter the same choice. These distances and masses are illustrative. The reference distance to Proxima does not commit the design to a Sun-like capture star.

Light for the first switch

Use 1,361 watts per square meter at one astronomical unit from a Sun-like star as a reference. A hundred square meters of collector facing that star from ten astronomical units receives about 1.36 kilowatts of light. At 25 percent electrical conversion and a 0.70 allowance for system losses and orientation, the available electrical power is about 238 watts. Put the same system at ten astronomical units from a star as dim as Proxima, with roughly 0.0015 times the Sun's luminosity, and the arithmetic gives about 0.36 watt. Spectral response can further change the latter result (Kopp and Lean 2011, Ribas et al. 2017).

The square of the distance determines when this increase is useful. A ship covering 127 kilometers per second takes about 13.6 days to travel one astronomical unit if it maintains its speed. At ten astronomical units, light may be sufficient for a limited revival. The vessel still has months rather than centuries in the immediate approach, and less if it waits for more power. Before relying on a large array, it must point the first collector, regulate its output, check instruments, and deploy equipment that has waited since departure. The first watt has a job even if it is too small to move the ship appreciably.

Where the braking comes from

Heller and Hippke modeled a path that uses starlight pressure and gravity to capture a very light photon sail in the Alpha Centauri system. Gros modeled momentum exchange with ionized interstellar material using a magnetic sail, including a slow mission whose travel time is of order ten thousand years. These are proposals with different assumptions about area, mass, materials, plasma, and deployment. They establish reasons to calculate possible capture paths for specific designs (Heller and Hippke 2017, Gros 2017).

One option for the book's reference mission is a distinct initial brake that acts before or near approach. The ship would wake on starlight in time to establish navigation and confirm the maneuver. Once bound, it could use solar-electric propulsion to change its orbit gradually, with light converted to electricity and electricity used to accelerate a reaction mass. NASA's Dawn used low thrust to enter, change, and depart orbits in the solar system, providing an operating precedent for this latter phase (NASA n.d.). A mission can use the same star for revival, orbital work, and eventually the power to make things. These uses have different power, thermal, and timing budgets.

A scale example shows what the spiral can reasonably be asked to do. Suppose a ten-tonne ship has 100 kilowatts for propulsion, an effective exhaust speed of 40 kilometers per second, and converts 65 percent of that input into exhaust kinetic power. The ideal thrust is twice that useful power divided by exhaust speed, or 3.25 newtons. At the starting mass that produces 1 kilometer per second of velocity change in roughly 36 days if the thrust is continuously useful. For a total velocity change of 10 kilometers per second the ideal rocket equation assigns about 22 percent of starting mass to propellant. At 92 kilometers per second it assigns about 90 percent. Constant input power, fully available propellant, and ideal thrust direction are optimistic assumptions. A practical capture calculation would integrate changing mass, light, trajectory, and thermal limits together.

This is the sequence each workshop has to qualify for its own route. Launch near a star. Let the protected machine cross the dark interval in a dormant state. Find a measurable way to dispose of enough incoming energy to stay. Build up power and move to a working orbit. Survey accessible material before committing to the workshop. The motion has an endpoint more precise than arrival. The receiver must be able to stay, see, and work long enough to send back evidence of what it has done.

Chapter 30 carries one Alpha Centauri A craft through the full sequence using a common mass, power, and time budget. The calculation puts the first battery charge, the actual removal of orbital energy, and the later electric spiral on separate accounts. It gives each proposed mechanism a threshold it would have to meet in a qualification test.

8 A Message and a Receiver

Consider a transmitter on Earth that sends a short build file once toward Proxima on the day a receiver ship departs. The first bits reach the destination roughly 4.25 years later. At that moment the receiver coasting at one tenth the speed of light is only a tenth of the way there. A receiver in the ten-thousand-year case is only about 0.043 percent of the way there. The beam continues through the destination and the ship arrives long after it. The light crossing time and the time to deliver a usable receiver are both real. An interstellar plan must arrange for a receiver to be present when the message passes it.

If the hardware is already in place and listening, the first part of a newly sent message can arrive after 4.25 years. If the hardware leaves today and is expected to be ready 42.5 years from now, we can schedule transmission to reach it at commissioning, or we can send an archive with the hardware and use the light link for updates. In the ideal coasting example, the schedule for a message to reach the star at the 42.5-year mark begins 38.25 years after departure. An actual schedule has to include the full duration of the transmission and uncertainty about when the receiver will be ready. A missed window can be costly if the only copy of the program was sent once. The millennial ship therefore departs with a complete baseline archive. Updates are an added capability for an institution that can still find, address, and authenticate the ship.

A simple delay account helps. Imagine the receiver completes its local survey and reports success. Its signal takes 4.25 years to reach Earth. A reply composed at once takes another 4.25 years to go back. The earliest decision that depends on that survey returns roughly 8.5 years after the original report leaves the destination. A mechanism that cannot handle an ordinary local failure without contacting Earth would have a very long time to wait. The receiver can still report and be guided on long intervals. Its early choices must be legible and safe within the authority assigned in advance.

What the hardware carries

An instruction to build a cell can look quite small in bytes if it names components and functions that the receiver already knows how to make. A short call such as *prepare the reader* can stand for a deep stack of physical operations. Imagine opening the receiver on arrival and asking where each such operation resides. Some operations are represented in the delivered file, some in the machine's software and

physical structure, and some in an organized material it carries. A meaningful account of the message length includes this division of work.

The receiver would need to recognize what local material is available and whether it meets its assumptions. It would need a source of usable energy and a place to send the heat. It would need instruments to distinguish a product that meets its specification from an apparent success. A receiver intended to build biology would also need a contained place to do the work and a way to observe whether the products renew themselves. None of these requirements disappears when the genome is digitized. They are the reason the word *receiver* has a physical meaning in the book.

The bookkeeping is most transparent if each part of the mission has a declared arrival state. A conventional biological carrier arrives with a preserved sample and whatever is required to wake or sustain it. A machine carrying a digital archive arrives with a readable record and equipment for acting on it. A two-stage design can combine the two. The archive can carry expected properties and independent versions of a component, while a protected sample supplies a comparison or a limited part of the starting state. Comparing these designs calls for the same endpoint. Each has to give the specified life an independent future in a specified destination environment.

A hardware mission can benefit from waiting to make the full habitat. If the destination supplies water and suitable elements, the receiver might carry the equipment needed to process them and let the habitat mass accumulate there. It has to be able to do that processing under observed conditions. A package that can fabricate with laboratory feedstocks could fail where the material occurs in a different mineral or where the accessible energy source produces a different thermal cycle. The declaration of the feedstocks belongs beside the declaration of the receiver's capabilities. The message is conditional on both.

What the signal can keep intact

A digital file has a useful property for a long crossing. If the channel adds errors, the receiver can often discover or correct them using extra symbols transmitted for that purpose. The possible correction depends on the error pattern and the redundancy chosen. A common failure in a decoder, a permanently misaligned telescope, or a loss of power can defeat every redundant copy. It is worth giving redundancy an explicit location in the design. Several copies on one device share

the fate of that device. Copies in distinct machines can still share a faulty interpretation of the file format.

Suppose a file has a trillion bits and a simplified channel independently flips one bit in a billion. The expected count is a thousand errors before correction. A seemingly low per-bit rate gives a poor chance of an exact file when so many decisions have to be made. Checks on blocks, error-correcting codes, and retries scheduled within the light time are tools for that problem. They also add to the communication budget. Even a perfect final copy leaves the question of whether the file describes a build that works.

There is a similar distinction between light time and data rate. Sending a trillion bits at an illustrative delivered rate of ten thousand useful bits per second takes one hundred million seconds, about 3.2 years. The first bit travels at the speed of light and the last bit leaves the transmitter years later. The example assumes the channel can maintain that rate and already accounts for whatever overhead was used to define the useful rate. The capacity to deliver a bit at an interstellar distance is a separate engineering question, requiring a link budget, power, aperture, pointing, and a listening strategy.

The ability to send updates nevertheless matters. Imagine that a crucial improvement to the program is discovered during the crossing. As long as the basic equipment can run it and receive the signal, the improved program can be sent on a schedule that intercepts the ship or reaches it at the destination. The state of a living cargo over that interval is a history the sender cannot retransmit. It can send advice about recovery, and the ship might carry a way to use it. Some missions will make excellent use of both approaches.

What waits after the crossing

The difference between the two plans appears most clearly in what must wait. In the biological plan, living material waits in a constrained state until it can continue its history. In the build plan, the receiver can wait for sufficient power and material, test its working space, and choose when to begin that history. If the survey finds that the environment would undermine the intended build, the program can remain archived while the receiver reports or changes its plan. Timing becomes a decision in the instruction.

In either case the habitat is an independent test. An organism arriving alive can die for lack of a renewable bath, and an organism made successfully can do the same. A

receiver that constructs one cell has demonstrated a narrow capability. A receiver that constructs an environment in which a community can persist has demonstrated a wider one. A destination intended to support a person has to provide flows of food, water, and air over many cycles. The burden of those flows is a large part of the reason to build near the place where they will be used.

A comparison on Earth can preserve the logic of the mission. Put two teams on opposite sides of a boundary. The sender chooses what to specify and transmits it under declared limits. The receiver reports what it already possessed, what it derived, and what it had to obtain physically. A test then asks whether the specified endpoint follows after the sender's intervention has stopped. That experiment begins with an ordinary cell. The next question is what the cell's description must carry beyond its sequence.

9 The Genome and the Working Cell

The receiver has passed its first serious test. A packet has arrived, it agrees with the checks sent from Earth, and the machine can read every base of a bacterial chromosome. On the display is the sequence of a laboratory strain of *Escherichia coli*. In the chamber is a set of prepared materials. The receiver can assign a checksum to the chromosome and verify the result of a synthesis step. What experiment would show that it has crossed from making a chromosome to making the kind of thing that can make another one?

The raw sequence fits in a surprisingly ordinary file. The K-12 MG1655 reference chromosome, GenBank accession U00096.3, has 4,641,652 base pairs. Give each base two bits and send one strand, with its complement recoverable by rule. The payload is 9,283,304 bits, or 1,160,413 bytes. The starting coordinate, order, and circular topology require a convention. The number says exactly how much a certain representation of the sequence costs. It leaves the chromosome's physical state, the material that surrounds it, and the route to a first division for the construction account (NCBI, U00096.3).

A ready recipient can supply much of that account. In 2010, Daniel Gibson and colleagues synthesized a bacterial genome and transplanted it into a recipient of another *Mycoplasma* species. The resulting lineage was controlled by the synthesized chromosome. At the start of the experiment a cell was already there, with a membrane and working molecular machinery. The achievement tells us what happens when a new genome takes over a cellular system. Our distant receiver could use this principle if it had suitable living material. Its independence from a living payload depends on where the initial cellular machinery comes from (Gibson et al. 2010).

Inheriting a place to begin

A daughter cell inherits a chemical environment as well as a sequence. It starts from the products and organization of a previous cell. Ribosomes translate the messages that code for ribosomal proteins. A membrane maintains differences between inside and outside, and transport processes depend on the membrane and those differences. The question for a receiver with no living material is how to begin these mutually dependent processes. The genome has tremendous power over the subsequent life of

the cell once there is a system able to express it. The early conditions can decide whether that expression ever occurs.

One of the clearest demonstrations that spatial structure can be passed on came from a very different organism. Janine Beisson and Tracy Sonneborn showed that altered organization of the cell cortex in *Paramecium* could persist through many divisions. The organization of an existing surface helped guide the organization of the new surface. Franklin Harold later assembled cases in which landmarks, gradients, local interactions, and continuity of structure contribute to the placement of molecules. The implication for a constructor is practical. When a previous structure normally supplies a location cue, a machine starting from other material must supply that cue, provide a route by which it arises, or demonstrate that the chosen system can function without it (Beisson and Sonneborn 1965, Harold 2005).

I use *bath* for the working material and arrangement around the sequence. For a cell that includes the boundary and compartments, the relevant chemical activities, and the molecular assemblies present when the genome begins to operate. On a machine that already understands a catalogue of molecules and knows how to make and combine them, we can give the bath file an illustrative size. Allow 10,000 abundance fields at 16 bits each. Provide 4,500 records of 48 bits each that point to sequence intervals and carry a little context. Set aside 120,000 bits for physical settings and references to assemblies. The result is 496,000 bits, or 62,000 bytes. The total sequence plus coarse bath packet is about 9.78 megabits. The counts and widths describe the example format. Proteomic measurements of more than 2,300 proteins across 22 conditions give a reason to expect a substantial catalogue and to specify which condition we want (Schmidt et al. 2016).

There is an immediate way to see what the half-megabit figure buys. Replace the receiver by a machine that can position atoms very precisely and understands none of the catalogue. Give it the same file. A reference to a membrane assembly is now a short integer whose interpretation has yet to be supplied. A good mission budget includes the catalogue, the shared interpreter, the machine that acts on the interpretation, and whatever the machine has to learn before it can execute a build. A message can be short because the destination already embodies a long history of design.

Where the work goes

The cost of a snapshot depends just as strongly on the task we define. Take a nominal bacterial volume of one femtoliter and use a spatial grid of one tenth of a

nanometer. A trillion sites fit into that volume. The approximate 0.30-molar measured metabolite pool would represent about 180 million molecules in the chosen volume. As a simple exercise, place their centers in separate grid sites. A list of positions would take around 40 bits per center, or more than seven billion bits for the pooled metabolites alone. If the centers are indistinguishable and their number is known, the file can identify a set of sites instead. The base-two logarithm of the corresponding binomial count is about 2.5 billion bits. Chemical identities, orientations, correlations, and motions would change the task again. This exercise counts the output of a specified recording convention for positions (Bennett et al. 2009).

An idealized atomic printer makes the other end of the comparison concrete. A femtoliter of water weighs about a picogram and contains about 3.3×10^{10} molecules, or 10^{11} atoms. Use the latter figure as an order-of-magnitude proxy for a wet bacterial cell. A serial machine handling one million atoms a second would take about 28 hours to place that many atoms. At a billion a second it would take about 100 seconds. That calculation gives the time to execute the assumed number of placements. A real build would also have to account for chemical bonds, parallelism, assembly errors, heat, and changes in the unfinished structure while the machine works. A recipe that delegates construction to working molecular assemblies has a different bill.

A build contract has a different output. Ask for a cell that can sustain a growing lineage under declared conditions. Individual molecular positions can vary between trials. If concentrations, spatial cues, and assemblies suffice to start a cell, the printer can let thermal motion and other local processes determine the exact arrangement. This shift from a snapshot to a family of working outcomes could greatly shorten a program. The size of the family has to be discovered. Trial runs can test whether omitting a field changes the frequency of success or creates a dependence on some unrecorded feature of the machine.

An interstellar version of that test needs another account, expressed in time, matter, and energy. A machine takes in usable energy, moves material, expels heat, and can start a system only when the outside conditions permit it to keep running. The cell itself maintains differences of concentration and electric potential through ongoing exchanges. A bit count specifies the length of a message, and the construction record has to supply the work required to implement it. The receiver gains leverage when a small number of settings place matter in conditions under which many

relevant events happen on their own. It pays for that leverage in known physical processes, feedstock, time, and the risk of a failed run.

The same distinction grows sharper at human scale. The gapless T2T-CHM13 human reference covers 3.055 billion base pairs and includes the chromosomes other than Y. One such set takes about 0.76 gigabytes in a raw two-bit encoding. A comparable pair of distinct nuclear sets would cost roughly twice that. Copying the sequence across a body is a construction quantity. The condition of the tissues, the relations among them, and the history by which those relations are reached define separate targets. A young system that can grow and an adult to be recovered in a particular state pose different information problems. Both need an outside environment. The body supplies the working milieu for many genomes, and a habitat in turn supplies the conditions the body can use (Nurk et al. 2022).

The first experiment worth sending is therefore modest in scale and exact about success. Give independent receivers the same materials, the same declared starting knowledge, and the same packet. Ask each one for an observed lineage that continues after the machine withdraws. Record what the packet calls, what the receiver had to know in advance, and what the chamber supplied as it ran. The difference between a receiver's initial state and that first growing lineage will tell us where to look for instructions that are still missing. A large project can begin with a result small enough to fit under a microscope.

10 The First Cell After Handoff

Imagine two chambers. In one, a machine has assembled a small compartment containing a genome. It makes a little protein, grows, and breaks into two. In the other, a similar compartment keeps producing descendants after the machine stops touching it. The descendants acquire matter from their surroundings and replace the molecular machinery that their founder was given. Both chambers would give an observer something worth watching. The second is the one I would ask the interstellar receiver to demonstrate.

The distinction turns on what the machine was allowed to supply, and when. A receiver can keep a cell-like object going with regular deliveries of molecular machines. It can divide the object itself when a sensor tells it the object has grown large enough. Those operations are part of the construction, even when the software calls them feeding and division. To find out whether the construction has taken over, choose a time at which the receiver gives up those operations. Leave it the job of observing and the ordinary work of maintaining the declared surroundings. From that time on, the descendants must carry out the cellular work under test.

I call this the first-cell handoff. The goal is a documented transition from a stated nonliving inventory to a renewing lineage. The experiment has to establish both ends of the transition. A good pedigree from an uncertain start leaves us guessing what entered the chamber. A well-audited start followed by a single interesting event leaves us guessing what the event could sustain. The evidence that connects the two is the central product of the experiment.

Declare what comes in

Start with a design and a chamber manifest. The design identifies the intended genetic material and the traits by which a descendant can be recognized. The manifest identifies the physical inventory at the beginning, the operations available to the machine, the intended surroundings after handoff, and the instruments used to check them. A receiver that starts with a physical DNA template has a different inventory from one that starts with its digital representation. This difference can be tested in a separate build. In a comparison of baths, keep the target and the template convention fixed so that a change in the result can be assigned to the change in the starting kit.

The word *nonliving* still leaves much to declare. A purified ribosome cannot found a lineage on its own. It can nevertheless carry out a step that a chamber of simple chemicals cannot yet carry out. Its origin also matters. A kit can contain parts

isolated from organisms while admitting no intact living cell. That kit supports a meaningful bottom-up experiment. A claim that the receiver can work at a lifeless destination asks where it obtained those same parts. If they crossed the distance in a package, enter them in the freight account. If the machine made them, enter the synthesis capability and evidence that it worked. Treat purification from a donor as a physical source of machinery, even if that donor never enters the chamber.

One way to make these distinctions useful is to compare graded kits. Begin with a rich, defined kit of preassembled molecular machinery, with a declared source for every item. In a second kit replace selected working complexes with their separate constituents, leaving their assembly to the receiver or to the developing system. In a third, ask which of those constituents the receiver can make from simpler material given the sequence and its allowed equipment. At the far end is a kit of explicitly specified precursors and an abiotic apparatus, with the target supplied digitally. The machine then has to make whatever templates and machinery its build requires. These are useful test conditions because each assigns a different share of the starting work to the receiver.

The comparison is strongest when one dependency changes at a time. Suppose a design succeeds with an assembled complex and repeatedly fails when given only its constituents. Put the complex back in a separate comparison. A recovered lineage would implicate the missing assembly under the tested conditions. The constituent kit might still work with a different construction path or a longer time allowance. The result is a boundary for this receiver, in this chamber, within the stated run time. Even a failure can reveal whether the machine was unable to make a component, unable to arrange it, or unable to find a working state often enough.

Give the machine a stopping time

Write down the handoff before beginning the build. The machine has a window in which it can inspect its work and make the adjustments authorized by the design. At handoff those permissions expire. The receiving chamber can still supply the allowed sources of matter and energy and keep its physical conditions within the declared range. A continuing nutrient flow belongs in that account. A dose of ready-made translation machinery, an operation that splits compartments, or a correction of an individual daughter belongs in a different account. The experiment can compare the two accounts by keeping assistance in one arm and withdrawing it in another.

For the duration of the post-handoff test, record every exchange across the boundary. A resource stream can contain highly organized material even when it has a modest mass. A chamber may also do some of the ordering work through a

moving surface, a timed change in its surroundings, or a sequence of signals. If such operations are part of the continuing environment, specify their schedule and count their contribution. A step timed by a particular cell's measured state is especially important. The response to that cell may be a step in the cellular cycle implemented by the apparatus. Once the experiment has transferred that responsibility, a request to repair the process ends the test of that handoff. It can begin a new attempt under a revised contract.

At the receiving star, this condition also makes the result easier to interpret. Earth can learn which actions the receiver took and which changes came from inside the chamber. Its delay makes direct supervision impractical. The same observational distinction lets a team on Earth tell a second laboratory what to repeat.

Watch the descendants do the work

The first observation is ancestry. Follow a founder into daughters, then follow daughters that themselves grow and divide. Measure the number of viable descendants associated with each tracked founder, including attempts that produce none. If the genetic material comes in several pieces, measure whether the necessary set passes together to descendants. Detecting every piece somewhere in the chamber gives a census of the chamber. The test needs to know which pieces travel together through division.

The second observation is renewal. A daughter may inherit a finite stock of functional molecules and use it for a while. A lineage that keeps going has to replenish the parts it consumes, damages, or distributes as it grows. Track enough post-handoff descendants to put the startup inventory under strain. Compare the amount of working material in the lineage with the maximum supplied at the beginning and any material admitted across the boundary. Measurements that distinguish the original stock from parts made from later feedstock strengthen that account. Ask whether newly made components enter working assemblies and serve the next round of growth. The birth of a new protein and the birth of a functioning machine assembled from it are separate observations.

The third observation is how often the transition occurs. A population can show a promising average while most founders stall. A rare founder can be scientifically valuable, and the cost of finding it belongs in the result. Count independently started trials, the time until a documented handoff, and failures through the declared horizon. Report the spread of outcomes between machines and batches. A receiver that achieves one success in a large search has a different contract from one that can start a lineage predictably. For a finite experiment, say how long the lineage renewed its machinery and how closely that renewal was measured. Its record can

support a longer trial without asking the reader to treat a finite observation as a guarantee of endless growth.

The controls protect the connection between these observations. A chamber run with the target instructions absent tests whether the measurement or the inventory can produce a misleading positive. A chamber with a defined omission and a matched restoration tests whether the apparent dependence is specific to that entry. A continuing-assistance arm shows how far the apparatus alone can sustain events that resemble the proposed handoff. A separate living reference tests whether the observations can actually resolve a growing pedigree under the relevant conditions. The reference belongs in separate custody so that the test for contamination does not introduce the very source it is meant to exclude. Independent provenance checks, blank runs, and repeated builds with separately obtained material make an unexplained success easier to investigate. Agree in advance on which observations count, so that a conspicuous image or a late favorable signal cannot silently change the result.

The work SpudCell actually handed off

The July 2026 SpudCell preprint offers a way to see these distinctions in a real experiment. Nathaniel Gaut and colleagues began with lipid compartments and a defined cell-free expression mixture. The system carried roughly 90,000 base pairs of DNA across seven plasmids. The authors reported genome replication, growth through uptake of feeder compartments, and effects of genetic changes on competition. A separate division experiment linked the genome to a protein that helped drive division in the presence of externally supplied components. The project makes it possible to ask which part of an integrated cycle comes from the template and which part comes from the supplied chemistry (Gaut et al. 2026).

The comparison depends on reading the multigeneration and division experiments separately. In the five-generation assay, the researchers added feeder compartments and physically divided the growing population at each cycle. In the separate genetically influenced division assay, they supplied agents outside the compartment that acted with its expressed product. About 30 percent of analyzed compartments had the complete seven-plasmid set after five cycles of the multigeneration assay. That fraction is an endpoint measurement. It matters for the engineering of inheritance, and it leaves a further question about individual lineages after assistance is withdrawn. The feeding arrangements also delivered working molecular machinery, including ribosomes. The stated inventory makes those contributions visible (Gaut et al. 2026, Biotic 2026).

SpudCell therefore supplies an example of how to divide the construction problem into experiments. It establishes that a system assembled from specified parts can link several activities usually seen together in a cell. It also makes clear what a stricter handoff would ask of the next design. The lineage would have to inherit the working genome, rebuild its expression machinery, and continue the cycle on the inputs permitted after handoff. Moving one of these tasks inside the compartment can be a substantive advance before the full test succeeds. It also changes what the receiver has to carry, make, or supply on each cycle.

What the experiment puts on the ledger

The receiver's report should make it possible to repeat the accounting as well as the construction. Record the starting kit and its provenance, the message and its length, the software and chemistry the machine already understood, and the construction work the machine did. Record energy, consumables, elapsed time, the number of attempts, and what crossed the chamber boundary after handoff. Finally, preserve the evidence for the pedigree and for the replacement of the founder's machinery. A short packet that calls an elaborate procedure may be an efficient design. Its reported size becomes meaningful once the procedure has a place on the ledger.

Return to the half-megabit coarse bath file. That example pays for fields and references under a specified decoder. A successful trial would tell us which records actually had to be present and which the machine could reconstruct or omit under the stated conditions. It would also reveal operations that the file had borrowed from the receiver's repertoire. If a shorter file gives a much lower yield, record both numbers. If it works only when a more elaborate kit arrives aboard the spacecraft, put that trade in the mission estimate. Bits, physical inventory, and search time enter the same design decision without requiring conversion into a common unit.

On a successful run, a receiver could send back a compact statement with a large experimental meaning. It began with this kit, used this packet and this amount of work, stopped making cellular decisions at this point, and observed descendants making and replacing the parts they needed from the declared surroundings. Another receiver could try the same transition. That is a practical milestone for the archive and a reason to ask what a more capable receiver could build next.

Chapter 28 follows this handoff in more detail. It asks how much of the founder's original equipment could still be doing the work, and what a blinded trial could learn from an apparent success in a blank.

11 What an Unknown Receiver Can Infer

Imagine that someone intercepts a complete packet during this mission. The power history and the tracking data reveal that it was sent on purpose. The intercepting receiver gets no header and knows nothing of the intended payload. It sees 9,283,304 bits. Give it a few ways to ask questions of the sequence and we can see how much of a message is accessible through structure alone. It is an Earth experiment, with a known answer hidden from the analysis. I generated the stream from the K-12 MG1655 genome by writing A as 00, C as 01, G as 10, and T as 11. The tests begin with the raw bits (NCBI, U00096.3).

What would you look for? Every long stream contains coincidences. A pattern becomes useful when it gives the receiver a prediction that an alternative would miss. If the receiver first tests a hundred periods and then reports only its favorite, the favorite deserves another measurement. For this experiment, the first half of the windows chose the candidate periods. A second half had to establish whether those candidates continued to matter. No gene annotations or known codon boundaries entered that calculation.

An interval worth asking about

In windows of 2,048 bits, a search among periods from two to 128 brought out a component with a period of six bits. On the held-out windows, the corresponding power was about 5.13 times as large as in a shuffle that preserves the separate numbers of ones at even and odd positions. A three-bit component also stood out. A two-bit component survived even this shuffle at the same strength, as expected from a control that keeps the alternation in composition. In the actual encoding C contributes a one to one slot of its two-bit word and G contributes a one to the other. A difference in their counts can make a two-bit rhythm without an additional biological pattern. A six-bit period is consistent with three bases per codon, and periodicity at three bases has long been used in sequence analysis. Period detection alone leaves the location and identity of coding regions open (Tiwari et al. 1997).

The controls allow the receiver to learn different things in sequence. It can detect a rhythm. It can ask how much survives under a change that preserves composition. It can try another alignment on independent data. It could then be told the packing and ask whether the rhythm tracks plausible biological features. Each step reduces some

uncertainty and names what evidence would resolve the next. By using that discipline on a biological transmission whose source we know, we can see the burden placed on a receiver asked to infer an alien format entirely from its contents. A frame that both ends agreed on in advance would save work at the destination.

There is also repetition of a different kind. Two positions 82,804 bits apart begin an exact match of 5,630 consecutive bits. The search used 64-bit seeds and extended matches in both directions. Nearby 96-bit chunks were shuffled in three controls. The rearrangement retains the contents of each chunk and breaks much of the long-range continuity. None of those controls yielded a match as long as 256 bits. At the actual two repeat boundaries, the first mismatch makes the exact search stop. Extend the alignment by 256 bits on each side and 253 positions agree on one side, with 250 on the other. The resemblance is broader than the exact repeat.

A receiver can exploit the resemblance without knowing its biological origin. If it can copy a previously decoded region, a starting address in a stream of this length takes 24 bits. A full copy command also needs a length, a choice of operation, and enough information to identify any changes. Those charges are small next to a few thousand bits copied exactly. The receiver can use that observation to design a compression format. It still needs other evidence to decide whether the repetition arose through the history of the genome, the design of the encoder, or another process. The description can become shorter before the explanation becomes more complete.

A prediction can carry its own bill

Another way to test structure is to estimate the probability of each next bit from its immediate context. I fit a model on the first half of the sequence, used the next quarter to choose a context length, and refit it before evaluating on the last quarter. Thirteen preceding bits gave the best validation result among the lengths tried. On roughly 2.32 million held-out bits the charge was 0.98047 bits per received bit, compared with about one bit per received bit for an independent frequency model. The saving on that portion is about 45,300 bits before the cost of describing the predictor. Its practical value depends on the code, the model's cost, and what parts of the model both ends already share.

There is a temptation to read a regular stream as the output of a succession of logic gates. We gave that idea a chance to earn a prediction. Divide the stream into nonoverlapping records, each containing an input word and an adjacent output word of equal width. Learn an alignment and a one-to-one map from the first half, then

evaluate on the second. For two-bit words, the best map swaps the two bits. It matched 30.39 percent of the held-out output words. Shuffling output words locally gives an expected match of 25.09 percent, while the identity map at the selected alignment matched 26.22 percent. The test recovers a dependence that appears in fresh data. It also shows that most individual words miss the putative gate prediction.

When we give the gate claim a probabilistic form, there is another useful comparison. A model that treats the SWAP result as one contribution to the output distribution can reduce the ideal coding charge relative to an unconditional distribution. Four output distributions learned separately, one for each input word, do better still. Their held-out cost is about 1.982 bits per two-bit output. The improvement is small and real in the test. Claims about a physical circuit or a quantum gate would need evidence about how the data were made. Four classical symbols written as bit pairs supply a classical truth table for this analysis. They supply no measurements of quantum phase.

You can make the design decision with this exercise in mind. Suppose the packet contains instructions for assessing material on a newly reached world. You can ask the receiver to infer the file's rhythm, locate the unit it is supposed to decode, assign meaning to its fields, and judge when a model is good enough to act on. Or you can supply physical conventions, framing, test records, and a common interpreter as part of the mission. Some structure could be rediscovered by a good analyst, as we did here with a deliberately simple biological file. Years may elapse while an Earth station waits for the first evidence of what the analyst got right. A usable archive can say where a decision came from and what observation would let the destination decline to act.

There is an open scientific question even for a perfectly legible packet. Can we find a pattern that helps predict what a cell can do? Prediction of neighboring bits is a beginning. Prediction of a molecular property under a declared interpretation is another step. An observed change in construction success when the receiver uses that prediction would connect the two. What travels well as a record may be quite different from what lets a machine act well when it is alone.

12 The Address and the Search

Imagine 10^{84} indexed positions, each holding a ball. One ball is white. Every other one is black. Someone gives you the location of the white ball. A binary address has to distinguish roughly 10^{84} possibilities, so the ideal length is about 279 bits and the fixed-width field needs 280. Now give a different person access to the same positions through an instrument that shows only the color at the one position selected. They have to find the white ball by trial. With no other information and no repetitions, a uniformly hidden target is found on average about halfway through the space. The address could fit into a single line of text. The expected search is of order 10^{84} inspections.

This is an engineering distinction worth carrying to the stars. A modest file can specify a solution that took enormous labor to find. A short sentence such as *find a working configuration* can hide an equally enormous amount of labor at the destination. The usefulness of either sentence depends on the receiver's access to candidates and to a test. A checksum can recognize one known file. A test of viability has to identify when the construction can continue growing. Even on Earth the possible states far outnumber the trials a laboratory can run. The target should be a class of outcomes that accomplish the mission, with a criterion that tells the machine when it has reached one.

An address for a translation table

The abstract genetic code makes the arithmetic concrete. Four bases yield 64 possible triplets. If each triplet can independently be assigned any one of 20 amino acids or a stop, the permissive set contains 21^{64} tables, approximately 4.2×10^{84}. An ideal index requires about 281.1 bits, and a fixed-width index takes 282. Storing each of the 64 labels in a five-bit field takes 320 bits. This convention gives one label to each triplet. Initiation, contextual interpretation, and the properties of the reader call for other information. Many of the counted tables would omit a required amino acid or lack a stop. The exercise gives an address space under simple rules, with neither a probability for a viable translation system nor a historical reconstruction of how a code arose (NCBI, translation table 11).

Now tell the receiver which triplets belong to each of the 20 synonymous classes and fix the stops. If all that remains is to permute the amino acid names among the given classes, there are 20! possibilities, or about 2.43×10^{18}. This version needs about 61.1 bits of ideal index length. We have made the search smaller by supplying

structure before it begins. The question shifts to how much structure the receiver was given and how it knew which structures were promising. An encoding can count the new bits explicitly. An account of a mission also has to acknowledge the prior information stored in its format.

There is a distinction between an address and a physical implementation at this scale too. The table is easy to send. It can tell the receiver what sequences mean and let it find regions in the reference. The machine still has to provide a working translation system at the destination, whether by building one or by starting with something that already does the work. The logical map tells the design what relationship has to hold. The molecular arrangement and the conditions under which the relationship holds determine how that instruction becomes operational.

A property worth measuring

Mary Yang, William Coleman, and I studied the relation between codon letters and a physical property of their assigned amino acids. Hydropathy puts amino acids on a numerical scale relevant to their relation to water. We expressed the scores using separate contributions from the three nucleotide positions and used the resulting representation to predict structural aspects of a photosynthetic reaction-center protein from its DNA. The exercise makes a useful intermediate goal for a mission archive. A receiver may know enough about a sequence to infer a broad physical property before it can make the folded protein or verify where the protein should go (Yang, Coleman, and Youvan 1990, Kyte and Doolittle 1982).

The comparison can be made explicit. Use the 64 triplets with equal weight. Assign the Kyte–Doolittle hydropathy value to each amino-acid codon, and give the three stops the placeholder value zero. At each position, take the mean of values for codons that have a given letter there. For each codon, the additive fit sums its three letter means and subtracts twice the overall mean. On the full grid the standard table gives an explained fraction of 0.81688, of which the second position accounts for about 0.719 of the total variation. This is a calculation about a named table, a weighting, and a chosen physical property. A different scale or convention for stops can change the answer.

A second calculation asks how often that score arises in a specified comparison family. Keep the stop locations and the synonymous classes fixed. Permute amino-acid labels among classes only when they have the same number of codons. Values tied on the numerical hydropathy scale lead to equivalent score inputs, with equal multiplicities in the labeled enumeration. Of 21,772,800 distinct numerical

assignments, exactly 1,920 meet or exceed the standard score. The tail fraction is about 0.0000882. Its negative logarithm in base two is 13.47 bits. This quantifies surprise under the stated shuffle. It speaks to the placement of hydropathy on the triplet grid, while the original 282-bit field identifies one table in a different and vastly larger space. Each result names a different task.

We can use the result to discipline the language of the mission. A rare score can suggest a good direction in which to search. An address can let the receiver recover an exact choice. A program can be evaluated by the time and materials it uses to cause a specified outcome. These measurements are compatible even when their magnitudes differ by many orders. A natural selection process, a laboratory optimization, and an independent draw from our enumeration each have their own access to intermediate outcomes. An exact enumeration of our shuffle says nothing by itself about how many historical trials happened.

Three budgets for one bacterium

Quantity	Example size or cost	What the number assumes
E. coli reference chromosome	9,283,304 bits	A known two-bit alphabet and one reference strand.
Coarse bath file	496,000 bits	A receiver that already knows the molecular catalogue and assembly routines.
Snapshot of metabolite centers	About 2.5 billion bits	Indistinguishable centers on distinct sites in the illustrative spatial grid.
Genetic code address	282 bits	One table chosen from 21^{64} abstract mappings under a shared enumeration.
Blind search for that table	About 2.1×10^{84} trials on average	One target, a perfect verifier, and no repeated trials or clues.

The same receiver can use a short code address and still need an elaborate translation system to implement the table. If the receiver already shares that table, a particular genome transmission may need only a short identifier. A bath instruction adds information about how to start or make the machinery. The entries in the table address different tasks, so adding their figures would mix encodings, physical assumptions, and search work.

Give the receiver a problem it can finish

The mission would benefit from shortening an instruction when the shorter version still works. Give an Earth demonstrator a complete recipe and a publicly declared set of tools. Run the recipe under the intended range of conditions. Remove one assumption and measure what happens. If a receiver can recover an omitted feature quickly from measurements it can take, the mission can give that receiver less data. If a change produces frequent failures, the full recipe, more equipment, or a slower and more reliable search might be required. The trade is between time, the cost of the machine, the length of the archive, and the chance of a good outcome. No single count speaks for all of them.

I would ask every proposed reconstruction for three entries in its account. What was known before reception? What new choices did the message fix? What work did the machine have to do after it opened the file? Make the test small enough to try again with the same declared ingredients. An experiment on one microbial lineage can find dependencies the same way a chemist can find what matters in a reaction by holding other things fixed. Once the target becomes a living community the acceptable outcomes also include time, renewal, and recovery from disturbance. The archive must say enough for the receiver to choose a workable path through those possibilities, then recognize the difference between reaching the destination and merely completing a list of commands.

13 The Human Bath Begins Before Birth

Imagine that the printer has produced an adult. The figure has the expected anatomy, and a sample of its DNA agrees with the archive. At the first breath, blood must reach a surface across which gases can move. Pressure differences must open and close valves in the right order. The cells lining the lungs must already present different surfaces toward the air and the tissue beneath them. Every one of these relations has a history in an ordinary human life. Our thought experiment asks for their result at the first instant.

In the bacterium, the daughter receives a portion of a working cell. A membrane and the machinery of expression are already there, and their presence lets the daughter renew what it inherited. A human starts with a living cell too. The mature egg supplies a boundary, mitochondria, molecular machinery, and a store of RNA and protein assembled during its own development. Fertilization brings the two parental chromosome sets together in this setting. As the embryo begins to express its genome, the inherited cytoplasm changes and parts of it are replaced. Studies of RNA in individual early human embryos can follow the transition from contributions present in the egg to transcripts produced by the embryo (Xue et al. 2013).

One cell makes a setting for the next

The egg supplies a place to start. The next stages generate new starting conditions. Cells divide and change their positions relative to their neighbors. Their contacts and the signals they receive influence which genes they express. In turn, what they express changes the signals they send, their ability to move, and the material they deposit around themselves. During gastrulation these interactions accompany the emergence of the major embryonic layers. A study of one gastrulating human embryo identified distinct populations and stages of differentiation through their patterns of gene expression (Tyser et al. 2021).

The sequence matters at every stage. It specifies proteins that recognize signals, modify other proteins, and bind particular regions of DNA. The action of any one of them depends on which molecules are present together and where they are. A cell can have the sequence for a receptor and respond only where the receptor is expressed, reaches a surface, and encounters its signal. A small difference between

neighboring cells can therefore be carried forward as a persistent difference in their surroundings and in the work they perform.

Suppose a receiver has learned to recognize a cell of the intestinal lining. A census of those cells would be a useful measurement, and their arrangement would still matter. The side exposed to the intestine does different work from the side in contact with underlying tissue. Junctions between cells contribute to a continuous boundary. An inventory can record the presence of its parts, while a measurement across the lining can test whether the parts have acquired a common orientation and form a working barrier. We have moved from a cell's differentiation to a property of a tissue.

In an ordinary pregnancy, the exchange of gases and nutrients through the placenta supports development before the lungs and digestive tract take on those jobs. Birth changes the conditions under which the body operates. An adult has had decades in which food, infection, injury, use, and experience can change tissues. If we specify a particular person at a particular moment, we are asking for some part of that acquired history too. A genome read from a blood sample could help identify a person, while telling us little about what the person remembers.

The price of a description

Our two-bit convention gives a human sequence a definite size. The T2T-CHM13 reference contains about 3.055 billion base pairs, with a gapless set of chromosomes other than Y. It takes about 6.11 billion bits, or 0.764 decimal gigabytes, to pack one such set. Two parental sets of comparable length take roughly 1.5 gigabytes. That charge is for recording two sequences once. A constructor that needs many physical copies must account for the materials and time to make them (Nurk et al. 2022).

Put beside that file a record of all the cells in a reference adult. Hatton and colleagues estimated about 36 trillion cells in a 70-kilogram adult male, with about seven trillion nucleated cells. Give the receiver an agreed dictionary of up to 1,024 cell classes and one 10-bit field per cell. The list takes 360 trillion bits, or 45 decimal terabytes. Its entries refer to classes known to the receiver. A large population of red blood cells repeats the same entry, so a list organized by counts and regions could be much shorter (Hatton et al. 2023).

The type label gives a place in a classification. It leaves open the cell's physiological state. For a different recording exercise, let each nucleated cell have

20,000 transcript abundance fields and give each field a 10-bit code. At seven trillion cells that is 1.4×10^{18} bits, or 175 petabytes, before compression. The fields could describe a range of abundance levels if both ends already shared a table defining what each level meant. A constructor may be able to start many cells with the same settings and let their conditions produce the local variations. How much it must specify would depend on which variations change the outcome.

Now assign a center to each cell. A cube just over two meters across can be indexed at one-micrometer intervals using 21 bits along each axis. The coordinates cost 63 bits per cell, about 2.27 quadrillion bits, or 284 terabytes, for the reference count. A separate record could give a cell's shape and orientation, and a map of neural connections would need to describe relations between cells whose centers can lie far apart. The coordinate list answers a particular question at a particular resolution. Its spatial regularities would be available to a compressor, and its omission of connections would matter to a constructor.

At molecular resolution, take ten kilograms of protein and use 40 kilodaltons as an illustrative average molecular mass. The result is about 1.5×10^{23} protein molecules. A scheme that divides the body into local cubes ten micrometers on a side and records a center to the nearest nanometer needs 14 bits per local coordinate. With the cube assignment taken as given, the local position list alone would use roughly 6×10^{24} bits. This example assigns a record to every molecule. Interchangeable copies and common structures allow shorter descriptions, and much of the location detail can change during ordinary motion without changing the operation of a tissue. The model uses a measured scale for total protein mass, with an assumed molecular mass and a chosen recording convention (Wang et al. 2003).

The same protein can appear in the sequence account, a tissue inventory, and a molecular list. The three descriptions answer different questions. Its amino acid sequence may follow from the genome, its abundance from a particular developmental and physiological state, and its location at an instant from the recent motions of molecules. Counting all three records as separate instructions would count some of the same structure twice. To attach a cost to what remains to be sent, declare what the receiver already knows, what physical state it can start from, and which features it must reproduce. The experiments would determine how far that account can be compressed.

Giving the archive time

An anatomical atlas can tell the receiver what to look for when a tissue matures. A developmental record can show which structures appear before others and what conditions accompanied the transitions. A build would require a starting system capable of making those transitions and an environment that continues to support them. All the human examples here begin with existing living cells. The figures specify what particular descriptions cost and give no tested route from the reader's archive to an adult.

The distinction carries over to a biosphere. A catalogue of organisms and a diagram of a mature community can describe its members and where they live. A founder would also need the conditions under which each population can grow, what material the others supply, and the order in which those exchanges become available. If a receiver has to wait for a crop or a microbial process to establish itself, it needs a reserve while that part of the living system develops. The account that follows begins by measuring what must pass through a habitat each day to sustain one adult.

14 The Person in the Balance Sheet

The receiver has finished assembling a chamber. Its gauges show a roomful of gas, a tank of water, and a harvest in storage. All three measurements could be exactly as the instruction required. A month later the tank could be empty, the gas composition different, and the next crop weeks from harvest. We can tell much more about the chamber by opening an account for each material and taking the same measurements a second time.

I want to start the account with a single adult. A destination habitat large enough for one person is an experimental unit. It gives us a load for which we can make numerical estimates and a scale on which to watch physical changes. The destination could operate the experiment with an apparatus that draws and returns materials on a prescribed schedule. That arrangement would let us compare measurements with the specification before anyone depended on the result.

The figure at the center of this account is a flow over time. In a NASA analysis of a deep-space life-support system, Harry W. Jones used daily planning rates of 0.84 kilograms of oxygen consumed, 1.00 kilogram of carbon dioxide returned, and 0.62 kilograms of food solids supplied per crew member. The same analysis allotted 3.52 kilograms per day to water for drinking and food. It specified separate water entries for washing and for flushing (Jones 2016). For our calculation I will use a round 3.00 kilograms per day for one narrow water stream, then a 30-kilogram-per-day example to see what happens as more uses enter the circuit. These are declared modeling choices, and the distinction between the two streams will turn out to matter.

What passes through one person

At the nominal daily rates, one adult accounts for about 307 kilograms of oxygen, 365 kilograms of carbon dioxide, and 226 kilograms of dry food in a 365-day year. The three-kilogram water stream carries 1,095 kilograms in that year. The larger water loop would circulate 10,950 kilograms. Anyone designing a tank, a separator, or a gas exchange system needs to know the flow it has to handle and the amount it can hold at one time. The annual number answers a question about throughput. The capacity number answers a question about what the equipment can do during a disturbance.

For a living-cargo comparison, put that person on the faster 4.25-light-year route. At a constant one-tenth the speed of light, ignoring acceleration and braking, the crossing takes 42.5 years in the departure frame. Using 365.25 days per year gives about 15,523 days. I use this as a common accounting interval. The elapsed time in the moving frame is about two and a half months shorter at this speed, a small correction compared with the uncertainty in a lifetime of consumption.

The accumulated throughput over the crossing is about 13.0 metric tonnes of oxygen, 15.5 tonnes of carbon dioxide, 9.6 tonnes of food solids, and 46.6 tonnes of the narrow water stream. Oxygen and dry food together represent about 22.7 tonnes moving toward one person. Carbon dioxide represents a separate stream moving away. If the wider water circuit were counted at 30 kilograms per day, its cumulative throughput would be about 466 tonnes.

Those numbers tell us what a recycling system would have to accomplish. A store that supplies a substance once and never receives it back must start with the whole amount or obtain more along the way. A system that returns the substance to service can circulate a far smaller inventory through the same total throughput. The habitat still needs to transform the return streams, and the equipment that does the work must last long enough to do it. Counting the same water every time it circulates is appropriate for assessing lifetime work on a separator. It would inflate an estimate of the water that actually has to be launched.

At the destination we can impose the same rates and observe the same balances. The starship calculation makes the stakes visible, and the destination experiment gives us somewhere to test an alternative. We could wait for the receiver to build and grow a productive system before the first human-scale load is imposed. This trades a sustained burden in transit for a founding problem after arrival. Whether that trade is favorable depends on what can be obtained at the site and how long the system takes to become productive.

Carbon moves even when the atmosphere looks steady

Carbon dioxide contains one carbon atom per molecule. Using approximate molecular masses of 12 for carbon and 44 for carbon dioxide, the return of one kilogram of carbon dioxide carries about 0.273 kilograms of carbon per day. That is almost 100 kilograms per year and about 4.2 tonnes over the 42.5-year interval. The amount of carbon in the air at any one time might stay small because the carbon goes somewhere else. It can become new plant tissue, dissolved carbon in water, or a solid deposit in material that is slow to exchange with the living system.

The oxygen and carbon dioxide rates also show why a pair of valves cannot stand in for the accounting. A daily consumption of 840 grams of oxygen is about 26.3 moles. A daily return of 1,000 grams of carbon dioxide is about 22.7 moles. The ratio is about 0.87 molecules of carbon dioxide returned per oxygen molecule consumed. For a simplified photosynthetic conversion that fixes one molecule of carbon dioxide per oxygen molecule released, fixing those 22.7 moles would return about 0.73 kilograms of oxygen. That leaves a difference of about 0.11 kilograms relative to the NASA oxygen rate in this narrowly drawn calculation.

The difference is a reason to count the other paths. The person oxidizes a mixture of food components. Carbon also travels through solid and liquid wastes. Plants respire while they grow, and other organisms process material that the crops leave behind. The oxidized and reduced products, including what is retained and what is removed, belong in the full balance. A reading from a crop enclosure under illumination gives one term. We need the atmosphere and all the compartments over a whole cycle to judge how the term contributes to a working habitat.

Three questions for every entry

Start with the amount of a substance in a named reservoir. Ask how much enters, how much leaves, and what transforms it into or out of the form we are counting. If a kilogram of carbon dioxide reacts with a wall, carbon has moved from one reservoir to another even when nothing crossed the outer boundary. If drinking water returns in a contaminated stream, the water remains present and becomes useful again only when the processing stage makes it available. The outside boundary, the interior reservoirs, and the form in which a substance can be used each get an entry.

We can say the same thing about stored food. The inventory is a set of dated harvests and preserved rations. The demand is daily. Production is intermittent, and some of the crop is unusable as human food. A new crop at the beginning of its growth cycle can look productive in a photograph while contributing very little to the inventory that week. A habitat may post a good annual production figure and still go through a gap when one crop is finished and the next has yet to mature. A serious instruction needs to arrange the timing of the transformation as well as the size of the average.

The age of the habitat complicates the first observation. The earliest years may include the filling of pools and stores, the consumption of material brought from Earth, and growth in the amount of tissue. A destination can use a large beginning

inventory to get through these stages. We can give the receiver enough time to establish a crop, use stockpiles openly, and ask how the rate of withdrawal from those stockpiles changes. Once the chamber is meant to operate as a renewable system, a stable reading at the point of use is more informative when it accompanies stable or deliberately replenished reserves.

The time window changes the meaning of a small error. Imagine that the site's gas account has an unexplained oxygen deficit of one percent of the ordinary human-scale draw. At 0.84 kilograms per day that is only 8.4 grams per day, easily obscured by larger transfers. Over a year it amounts to about 3.1 kilograms. Over the 42.5-year interval it reaches about 130 kilograms. Whether the deficit represents a leak, chemical uptake, or bias in a sensor determines how the project should respond. An experiment that resolves the discrepancy after thirty days has established a method for finding losses before the cumulative size becomes substantial.

At an interstellar site, someone on Earth may wait more than eight years for a round trip through that conversation. The local receiver has to distinguish a noisy measurement from a steadily shrinking reserve without waiting for a reply. The construction message therefore needs to carry measurement procedures and an account of where to look when the measurements disagree. Even a broad rule that keeps a second record in a different physical quantity can make the returned data easier to interpret.

The central result of this small calculation is a list of jobs for the destination, each with units. It has to carry a daily human-scale load, process the return streams, furnish enough energy to keep the transformations going, and preserve enough reserve to ride out uneven production. The first estimate of the printer's work ends at the point where a fabricated object exists. The account for the biosphere begins there and stays open for the life of the experiment.

15 A Floor Area for Food

The oxygen number is usually the first one people notice. I find it more revealing to start with a meal. Take a planning intake of 2,500 dietary kilocalories per day. Multiplying by 4,184 joules per kilocalorie gives 10.46 megajoules per day, an average rate of about 121 watts of chemical energy. Over a year the food supplies about 3.82 gigajoules. This quantity gives us one way to estimate how much light a destination farm has to intercept.

The diet and the rate of 0.62 kilograms of food solids introduced in the last chapter describe different aspects of the load. If both applied to the same day, the mean food energy per gram of dry food would be about 4.0 kilocalories. A change of menu can shift that ratio. A complete diet also has requirements for amino acids, fats, vitamins, and other nutrients. Calories give us a starting scale for the growing area and leave the nutritional problem visible as another specification.

Light divided by yield

Suppose that a crop surface receives 20 megajoules of total sunlight per square meter per day. This is a declared example of an Earth-like illumination budget at the canopy. At an actual destination we would have to measure the spectrum, angle, exposure, and amount of usable light after whatever passes through or around the shielding. Let the net edible conversion efficiency be the fraction of that incident energy recovered as available food energy, averaged over full growing and harvest cycles. This definition covers the fact that plants make roots and stems, use energy themselves, and lose some of the harvest along the way.

At one-half percent net conversion, each square meter supplies 0.10 megajoules per day. The required growing area is 10.46 divided by 0.10, or about 105 square meters. At one percent the output is 0.20 megajoules per day per square meter and the result is about 52 square meters. At two percent the result is about 26 square meters. These are three cases in the same arithmetic, with the efficiency assumed in each. They show how much a change in the realized conversion affects the layout. At half the assumed illumination and the same conversion fraction, the area doubles.

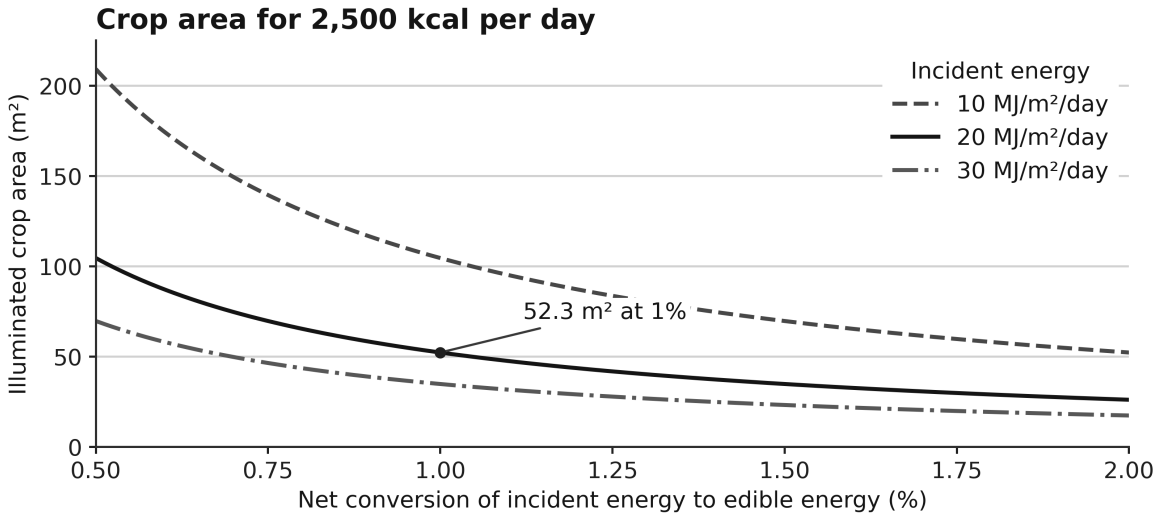


Figure 2. Illuminated crop area from the assumed incident energy and edible conversion.

These areas refer to crop canopy. A habitat also needs space for water circulation and storage, waste treatment, access, servicing, and buffering. Stacking crops can shorten the footprint of a growing enclosure, with separate lighting or distribution of light to each layer. The demand for illuminated area remains. On a site with seasons, a whole-year mean can also obscure the weeks when the light is scarce and storage has to carry the deficit. The layout and its reserves have to work together.

There is an experimental scale against which to check the estimate. In a NASA technical memorandum, Raymond Wheeler summarized studies with about 40 moles of photosynthetically active photons per square meter per day. He reported that the results suggested roughly 20 to 25 square meters of crops might supply oxygen for one human and about 50 square meters might supply the dietary calories (Wheeler 2020). This compares different objectives under the tested conditions. On the scale of those estimates, meeting the calorie target requires more growing area than meeting the oxygen target.

Our 52-square-meter calculation happens to sit near Wheeler's food estimate. The first comes from assuming a one-percent net conversion for a particular daily sunlight budget. The second draws on controlled crop experiments. Their proximity makes this a plausible order of magnitude for an example. Complete-diet production and atmospheric control would each require measurements across the crop cycle. A canopy that performs one task can contribute to the other, and its net effects depend on the rest of the living and chemical system.

Suppose the one-percent farm has 52 square meters. For a thirty-day interval it receives 10 rather than 20 megajoules of sunlight per square meter per day, and for

simplicity the conversion fraction stays the same. The model produces about 5.2 megajoules of food energy per day during that interval and falls short by roughly 158 megajoules in all. At the energy density of the illustrative dry ration, that is about 9.4 kilograms of food solids. We could add area, hold a larger reserve, substitute artificial light, or postpone the demand. The choice changes what the receiver has to build and stock, and it depends on the duration of the weak season as much as the mean intensity.

What electric light costs

At a location with poor or inconvenient natural light, a receiver could turn electricity into the photons the crops use. This adds a measurable load to the destination. Take 60 square meters of canopy and an illustrative daily light integral of 25 moles of photosynthetically active photons per square meter. That is 1,500 moles of photons per day at the canopy. A fixture efficacy of 2.5 micromoles per joule delivers one mole for each 400,000 joules of electricity, under the ideal assumption that every emitted photon reaches the canopy. Kusuma, Pattison, and Bugbee report that commercial horticultural fixtures using white and red LEDs had reached this scale of efficacy (2020).

Divide 1,500 moles by 2.5 millionths of a mole per joule. The result is 600 megajoules of electricity per day, or about 167 kilowatt-hours. Averaged across the entire day, that is 6.94 kilowatts of electrical input just for the lamps. If the lamps give the whole daily amount during sixteen hours, their mean demand during those hours is about 10.4 kilowatts. If the crop intercepts 80 percent of the emitted light in this model, meeting the same target takes about 8.7 kilowatts on average across the day. The losses have a location, and much of their energy still enters the habitat as heat.

The 25-mole case gives a lighting bill. Wheeler's benchmark uses a higher light level, and we would need evidence to connect the lower level to the same food yield. Repeating the calculation for 30 and 40 moles at the same fixture efficacy and perfect capture gives 8.33 and 11.11 kilowatts of average electrical input. These estimates leave out pump power, air circulation, nutrient preparation, waste processing, and the work of keeping the growing environment in range. The 25-mole ideal case alone consumes about 61 megawatt-hours of electricity in a year.

The average chemical energy delivered in a 2,500-kilocalorie diet is about 121 watts. Comparing that with 6.94 kilowatts gives about 1.7 percent as a ratio of dietary output to electrical lighting input if the same example achieved both goals.

This is a system-level ratio spanning the fixture, the crop, and the harvest. A useful measurement would report the photon input, the edible output, and the uses of the rest of the biomass separately. We could then see whether a different crop or light schedule reduced the electricity needed for the same harvest.

Where the energy goes after the harvest

If the lighting and crops are both inside the habitat boundary, most of the input energy ultimately has to leave as heat when the stock of chemical energy is no longer growing. During plant growth some energy is stored in biomass. Part of it later passes through the person or the waste processor. Any energy exported with products leaves by a different route. An accounting interval that covers planting, harvest, and use lets us see the net accumulation rather than count every intermediate transfer as a fresh demand.

For a physical scale, imagine an ideal one-sided black radiator at 300 kelvins looking at a much colder background. The Stefan–Boltzmann relation gives about 459 watts emitted per square meter. Disposing of 6.94 kilowatts under those ideal conditions would take about 15 square meters of radiating surface. A real calculation would use the emissivity, the temperature of the sink, view factors, and the power and temperature costs of transporting heat to the surface. This is a comparison of scales. It tells us why a receiver's claim to have built a farm includes a claim about the physical capacity to drive it and remove the resulting heat.

There is a second way to read the growing-area estimate. Suppose the 60-square-meter farm were to produce the 226 kilograms of dry food solids in our annual plan. It would have to deliver an average of about 3.8 kilograms of usable dry food per square meter per year. Planting failures and early growth reduce the production available in that year, so reporting the yield only for occupied and successful trays can distort the comparison. The benchmark belongs to the entire canopy we commit to the task and the entire period over which someone depends on the food. What looks like a single number for area is the output of several accounts that the receiver would have to keep together.

16 When Air and Water Loops Fail

A habitat can look successful whenever its pumps and lamps are running. Its weakness may be easier to see during the interval after a load changes. For an example, take a sealed 100-cubic-meter gas volume, initially at about 101.3 kilopascals and 293 kelvins, with an oxygen mole fraction of 0.21. These are model conditions for an unoccupied test volume. I will let a simulated demand use oxygen and return carbon dioxide at the daily planning rates while turning off regeneration for five days. We can work out the scale of the resulting stock changes with the ideal gas law.

Multiply pressure by volume and divide by the gas constant and temperature. The initial gas inventory is about 4,160 moles. Oxygen is 21 percent of that amount, or about 874 moles. At approximately 32 grams per mole, that represents 28 kilograms of oxygen in the gas at the beginning. Near that starting state, a change of one percentage point of oxygen fraction at fixed pressure and temperature corresponds to about 1.33 kilograms of oxygen in a 100-cubic-meter volume. This conversion gives us a quick scale, while the detailed calculation can allow total moles to change.

The simulated draw is 0.84 kilograms per day, so after five days it removes 4.2 kilograms of oxygen, or 131 moles. The simulated return of one kilogram of carbon dioxide per day adds five kilograms, about 114 moles. With other exchanges omitted, the number of gas moles falls only about 17, or about 0.4 percent. The oxygen fraction falls from 21 percent to about 18 percent. If the starting carbon dioxide fraction were 0.04 percent and none were removed or absorbed, the model would reach about 2.8 percent carbon dioxide. All these figures are inventory results for the stated artificial disturbance. The simulation leaves out respiration by other organisms, gas dissolution, changes in vapor, temperature variation, and the chemical effects of surfaces.

The small pressure change is an instructive result. In this example a gauge reporting total pressure moves little while the mixture changes markedly. We could make a pressure controller that compensates for the change and still have to ask what mixture the controller is producing. In an actual biosphere a single pressure reading tells us little about how much of the required material is available in a useful form. The measurements have to follow the species that make up the atmosphere and the reservoirs with which they exchange.

A buffer has to refill

At the end of the five-day disturbance the oxygen inventory in this model is 4.2 kilograms below its initial value. Suppose the regenerating system resumes production at exactly 0.84 kilograms per day while the simulated draw continues. The inventory stops falling and stays 4.2 kilograms below the original level. In the model, surplus production of 0.168 kilograms per day, a margin of 20 percent over the ordinary load, would refill the oxygen deficit in about 25 days. A 50 percent margin would take ten days. There would also have to be a separate account for removing the accumulated carbon dioxide. A successful restart and a completed recovery can occur weeks apart.

The same test has a food side. Over five days, the model uses 3.1 kilograms of dry food and 12,500 kilocalories. Suppose the farm normally produces just as much as the person eats. A crop failure followed by average production at the previous level leaves the reserve depleted. Extra production can replenish it, with a delay for new growth and a requirement that the recovered crop actually be usable. Two failures separated by less than the recovery interval can therefore exceed the reserve even when either failure alone appears manageable. A demonstration that stops at the end of the first disturbance will miss this dependence on timing.

The experiment can also ask where buffers come from. The chamber has walls and tanks with capacities, and the receiver can choose which stores are accessible to which parts of the system. Extra gas volume dilutes composition changes for a fixed simulated load. A physically separate reserve can limit the effect of a leak in one compartment. An extra harvest kept for later can reduce dependence on the growth of a particular crop. Every buffer has an initial material cost and some means of replenishment. Reporting how the reserve is charged after a test brings both costs into view.

Recovery rates need a denominator

NASA reported in 2023 that the International Space Station had demonstrated total water recovery of about 98 percent with the addition of a brine processor. The result concerns a system for collected water and its processing streams. The station operates with supporting infrastructure and resupply, and the water figure describes a particular recovery achievement (Gaskill 2023). It gives us a concrete fraction to use in a sensitivity calculation for our own model.

At 98 percent recovery, processing 3 kilograms per day returns 2.94 kilograms and leaves a 0.06-kilogram daily gap. If the unrecovered portion were permanently

unavailable, makeup would be 21.9 kilograms per 365-day year and about 0.93 tonnes over 42.5 years. Run the same fraction on a 30-kilogram-per-day loop and the gap becomes 0.60 kilograms per day, 219 kilograms per year, and about 9.3 tonnes over the longer interval. The percentage is identical in both cases. The effect depends on what enters the denominator.

Even that arithmetic leaves a question about where the remainder goes. Some water might leave with a waste product. Some could collect in material that is processed by another loop. Some could remain present in the apparatus in a form the crew cannot use. If another system recovers it, the two loops must have separate balances and their common transfers must be matched. An inventory loss across the outer boundary and a transfer to a temporarily inaccessible internal stock can have the same short-term effect at a tap. The long-term implications are different.

An ecosystem can add much larger flows without a proportional demand for new water. Consider a hypothetical 60-square-meter growing area transpiring two kilograms of water per square meter per day. That would put about 120 kilograms per day into the air as vapor. This rate is chosen to illustrate the physics. If the vapor is condensed and returned, it can take repeated trips through the same habitat. At a round latent heat of 2.45 megajoules per kilogram, condensation handles about 294 megajoules per day, a thermal duty of roughly 3.4 kilowatts. The heat comes from evaporation and can be moved or partly recovered. Where the lamps supplied the heat of evaporation, counting that heat as new energy when the vapor condenses would count it twice. The example shows that water handling can set the size of a heat exchange system even where fresh water demand stays small.

A second treatment stage changes the percentage by acting on the remainder. Return to the 30-kilogram water loop and suppose a different process recovers four-fifths of the 0.60 kilograms that the first one leaves behind. Then only 0.12 kilograms per day would remain unavailable, corresponding to 99.6 percent combined recovery and about 44 kilograms of annual makeup under the same assumptions. The smaller loss has to be weighed against the second process's power, replacement parts, and new waste stream. The number would be useful to the receiving system only with a description of which compounds accumulate in that stream and which consumables the process uses. Across forty years a missing filter or exhausted reactant could be as important as the water percentage printed in the first report.

What the disturbance reveals

At the destination the first test could impose a controlled reduction in one conversion process while recording inventories across the system. It could compare successive five-day intervals under the same simulated demand, following each one long enough to see whether the buffers refill. I would want to see traces for accessible water, atmospheric composition, food reserves, and the stocks that the processing system tends to leave behind. The time at which a reading returns to normal matters less by itself than the relation between that reading and the resources consumed in getting there.

The important result might be a slow change rather than an immediate loss. A chamber can remain within its chosen test bounds while its reserves decline, and the decline can be detected by comparing equivalent points in repeated cycles. The trajectory also tells us what needs a larger margin. The choice could be to store more, process a waste stream differently, keep redundant crops, or allow the system to establish itself longer before applying the load. A receiver that can make these distinctions has learned something a single successful day could never reveal. It has found out where the habitat spends its capacity to survive the next interruption.

17 A Year Inside the Boundary

Put the one-person experiment on a clock. Before starting it, the receiver weighs the water in its working tank, takes samples from the nutrient system, measures the air, and records how much edible material is in store. Then it starts the pumps, imposes the simulated human load, and waits through the harvest schedule. At the end of the year it adds up the work that has been done. It also compares inventories at corresponding phases of successive cycles, using separate dates if the harvest schedule and the year end at different points. The work done and the stocks left tell the receiver how long it could keep operating.

Use a 365-day trial and the one-adult planning rates introduced earlier. The simulated load draws 0.84 kilogram of oxygen and 0.62 kilogram of dry food per day, and returns one kilogram of carbon dioxide per day. The dietary energy target is 10.46 megajoules per day. A separate water branch processes 30 kilograms per day in this example. The reference rates come from different planning choices, and a real experiment would specify a consistent diet and a measured composition of the return streams. Here the oxygen, food, and carbon dioxide rates set the scale of the jobs, while the water branch shows how a particular recovery process affects reserves. Water transpired by crops belongs in its own account unless it passes through the same measured branch.

In the year, the oxygen draw amounts to about 307 kilograms, the carbon dioxide return to 365 kilograms, and the food draw to 226 kilograms of dry matter or 3.82 gigajoules of dietary energy. The water branch sees 10,950 kilograms. Those totals are throughput. The receiver could circulate an inventory smaller than any of the totals if its return paths were fast enough. It would still have to measure how much remained in service after each pass and what happened to the rest.

A return path that misses a little

Let U be the mass of water the system can draw from its working tank, Q the mass processed per day, and r the fraction of that flow returned to the useful stock. If daily throughput and recovery stay fixed, and there is no addition from another stock, the simplest balance is

U(t) = U(0) − Q × (1 − r) × t.

Take a working tank with 100 kilograms initially and an experimental minimum of 50 kilograms. These amounts are chosen to show how the calculation works. The

minimum is a stipulated test boundary, independent of any human occupancy requirement. The unused margin at the start is 50 kilograms. At a rate of 30 kilograms per day, the time to the stipulated minimum is 50 divided by 30 × (1 − r), with that whole product in the denominator.

Recovery per pass	Water drawn from the reserve in 365 days	Time to the stipulated minimum
98 percent	219 kg	83 days
99.6 percent	43.8 kg	417 days
99.9 percent	11.0 kg	1,667 days

The first line gives the demand for reserve water if the assumed load were maintained for the full year. The trial would actually cross its minimum much earlier. The middle line uses the two-stage calculation from the earlier chapter. A first stage returns 98 percent, and a second returns four-fifths of what is left. The first year's ending stock is 56.2 kilograms. A test that reports only that the water remained above its minimum for a year can record a pass. At the same rate it would take less than two additional months to reach that minimum. The trial has demonstrated endurance from an initial inventory. A second year under the same contract would expose the difference between that endurance and renewal.

The inverse calculation makes a useful requirement. Given a duration T and an expendable margin M, the maximum average unreturned flow is M divided by T. At constant Q, the required recovery fraction is at least 1 − M/(Q × T). In the example it is about 99.54 percent for one year and 99.77 percent for two years. Sustaining the process indefinitely with the same initial stock requires an additional return or a make-up path that balances the continuing loss. If the material is still inside the chamber, there is something to retrieve. How it can be retrieved, how fast, and at what cost become new questions for the experiment.

Where a closed total can hide an open account

Suppose the 43.8 kilograms unavailable to the working tank in the middle case collects in a brine reservoir. If the reservoir is assayed independently, the water balance over the full chamber can close to the accuracy of measurement. At the same time, the working tank's balance is declining at 0.12 kilogram per day. Both statements can be true. One describes the location of matter, and the other describes whether the apparatus can use it. The brine becomes part of the available reserve when a demonstrated process can deliver its water at a suitable quality and rate.

Now measure phosphorus at the same harvest phase as the working water. With five kilograms of accessible phosphorus, a recovery step processing two kilograms per year and returning 98 percent would transfer 0.04 kilogram out of the accessible pool each year under the assumptions of the slow-sink example. A test floor of 3.5 kilograms leaves 1.5 kilograms of margin, equivalent to 37.5 years at that constant transfer rate. The first year could therefore look nearly unchanged even when the nutrient has a longer-run vulnerability. Assaying the suspected sink can be more revealing than waiting to detect a small difference between two large readings. Carbon has the same accounting problem. If one percent of the carbon in the nominal annual respiratory return moved into a persistently inaccessible reservoir, about one kilogram of carbon would accumulate there in a year. A stable atmospheric reading would leave that movement invisible if replenishment kept pace.

For an element whose chemistry stays inside the boundary, the sum of changes in every measured reservoir should equal imports minus exports. A useful test therefore records both the residual of that sum and the individual inventory changes. A small residual says the inventory is consistent with the measured transfers within their uncertainty. It leaves room for a harmful shift from one reservoir to another. The records also need to say which material is in transit, which is ready for use, and which will return only after a delay. A nursery holding nutrient in living tissue and a sealed precipitate can contain the same element with very different consequences for next month's harvest.

The allowed residual should follow from the experiment's decision. If the test must distinguish a loss of 0.12 kilogram per day from zero, a calibration offset of 0.05 kilogram per day already corresponds to 18.25 kilograms over the year. Repeating readings from that instrument does little to reduce the offset. The receiver can compare a tank inventory with a separately measured processing flow, take samples on a stated mixing schedule, and retain the raw readings used to reconcile the accounts. It can then ask whether an upper estimate of the true loss, including measurement uncertainty, can consume the available margin before the test ends.

Carry the power through the same year

In the illustrative electric farm, the canopy covers 60 square meters and receives 25 moles of photons per square meter per day. The emitted-light efficacy is 2.5 micromoles per joule. Let c be the fraction of emitted photons intercepted by the canopy. Lighting takes 6.94/c kilowatts averaged over the day. For c between 0.8

and 1.0, it takes about 6.94 to 8.68 kilowatts. Give the rest of the apparatus an assumed continuous two-kilowatt load to see the scale of an integrated bill. The resulting average is about 8.94 to 10.68 kilowatts, or 78 to 94 megawatt-hours over 365 days. The yield belongs in a harvest record made under the same lighting and operating conditions.

If the photons arrive during a 16-hour lighting interval and the auxiliary draw stays constant, the electrical demand during that interval is approximately 12.4 to 15.0 kilowatts. The facility must furnish the interval demand as well as the daily total. A supply that can deliver the daily energy only at a lower rate would need storage or a different lighting schedule. If the modeled electrical input ultimately leaves as heat, an ideal radiator at 293 kelvins with emissivity 0.9 and an unobstructed view of cold space would need about 24 to 28 square meters of emitting area on the average-power calculation. The peak demand, the thermal store, absorbed radiation, and the connection from the heat sources to the radiator would have to be checked separately. Condensing transpired water can impose a large internal heat-transfer duty. The same energy should enter the outer-boundary account only once.

A year of operation can therefore have several distinct results. The food store can meet the scheduled load, the water system can remain above its test minimum, and the nutrient reservoir can already show a continuing transfer into a sink. The power account can balance on an average day while a particular interval exceeds the available supply. The receiver has to reproduce the time series as well as the totals. Comparable checkpoints reveal what the system spent, and the intervening readings reveal when it nearly ran out.

18 The Average Year and the Failed Harvest

Put a person and a garden inside an imaginary habitat. At the end of a year the garden has supplied exactly the energy that the person consumed. That statistic allows two very different histories. In one, harvests arrive close to the times when food is eaten. In the other, an early harvest builds a reserve that is nearly exhausted before the next one arrives. The average for the year records the same outcome either way. We need to know the lowest level in the reserve to tell whether the histories were equally viable.

We can calculate the difference before choosing a crop. Use the book's one-adult reference demand of 2,500 kilocalories a day. At 4,184 joules per kilocalorie, this is 10.46 megajoules a day. Give the garden an average incident solar energy of 20 megajoules per square meter per day and an illustrative net conversion of one percent from incident energy to edible energy. Its average production is 0.20 megajoule per square meter per day. A growing area of 52.3 square meters balances the annual arithmetic. This is an energy-equivalent yield calculation. The area describes illuminated crop surface under those assumptions, separate from walkways, plumbing, service spaces, and the land or structure that supports the rest of the habitat.

A reserve for the year we already expect

Keep the average yield and let production vary smoothly over a 365-day year. Start with an amplitude of 35 percent of average production. The high point gives 14.12 megajoules per day and the low point gives 6.80. When production falls below demand, the reserve falls. When production exceeds demand, the reserve rises. A convenient model is a sinusoidal curve around the mean. Real crops grow and are harvested on discrete schedules, so this curve is a deliberate simplification for measuring the size of a seasonal effect.

The difference between the maximum and minimum stock is the integral of the surplus over the bright half of the year. It equals the daily amplitude of 3.661 megajoules multiplied by 365 days and divided by π. The answer is about 425 megajoules. Expressed in the same units as the adult's demand, that is about 41 days of food. It is the ideal swing in storage between two dates. The phase at startup determines how much of the stock must already be present when the printer hands

over. If we start during a drawdown, we have to supply the remainder of that drawdown before the next sustained surplus begins.

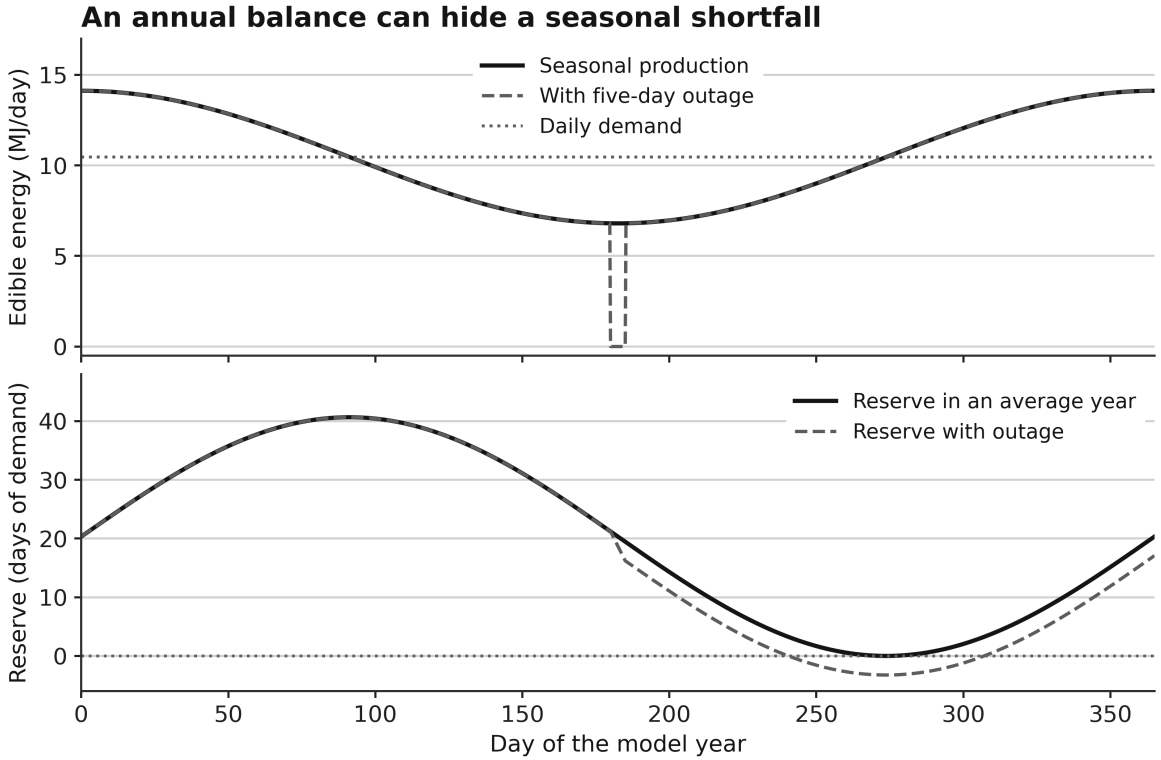


Figure 3. The seasonal food flow and its cumulative reserve. A negative reserve represents a shortfall.

The annual mean leaves the 41-day swing invisible. That is the practical difference between supplying a rate and supplying a system. If the reserve is spent on food before the garden begins producing, the build instruction has to provide that initial food. If the founding habitat is allowed to grow for a year before a person arrives, the plants and the storage system may build the reserve first. The description can then specify a readiness test and wait for it. Either program still owes us the full account of what entered the chamber and when.

The reserve rises in proportion to the size of the seasonal variation. A ten percent swing requires the equivalent of about twelve days of food in this model. A fifty percent swing requires about fifty-eight days. Each number still assumes that the mean production exactly meets mean demand and that the store has no losses. A real farm might reduce seasonal amplitude through a controlled light supply, planting schedules, or production modules whose harvests peak at different times. We would count the lights, other equipment, and changes to the operating program as part of the intervention. The improvement would show up in the lower part of the measured stock trajectory.

The shock that changes the next year

Suppose five production days near the seasonal low are lost altogether. The occupant draws 52.3 megajoules from the store during those days. The planned seasonal ledger already expected to supply part of that withdrawal, because production was low. The extra loss is the food that would have been harvested. At the bottom of the illustrative seasonal curve it is about 34 megajoules over five days, a little over three days of demand. The same outage at the production peak removes about 71 megajoules, nearly seven days of demand. The timing of the outage also determines whether the missing production falls just before the reserve would normally reach its minimum.

This way of charging a disturbance to the budget can be used for a longer crop failure. We need the counterfactual trajectory, the difference made by the failure, and the subsequent time available for recovery. A season with zero net production surplus has no mechanism in this calculation to refill an extra loss. It can follow the same annual cycle next year at a lower level and still report an average production equal to demand. A small positive surplus changes the outcome. A surplus of five percent of demand would replace 34 megajoules in about sixty-five days if the surplus were available steadily. The seasonal model will move the date at which the reserve actually recovers. We should specify the recovery window in the test.

In an inhabited world, this loss might coincide with other deficits. The source of the light could also power the water loop. A disease in a crop could affect food while changing gas exchange and the amount of dead material that decomposers receive. The calculation tells us where to put our attention. We can connect the same disturbance to several stocks and see which one crosses its operating range first. A count of spare components tells us less than a test in which the common power source or a common service is interrupted.

Power can substitute for time

The same garden can be represented by a lighting bill. The book's illustrative electric case uses 60 square meters, a daily photon delivery of 25 moles per square meter, and 2.5 micromoles delivered per joule of electrical energy. That requires about 167 kilowatt-hours a day, or 6.94 kilowatts averaged over a full day. The edible energy going to one person is about 2.91 kilowatt-hours per day. The two numbers have different roles. The lighting number represents an input to the habitat, and the food number represents energy in one of its outputs. The lighting

assumption leaves room for losses in the equipment and for energy that goes to plant maintenance, inedible biomass, and heat.

Suppose a solar collector receives the same average incident energy per area as the garden and converts a quarter of that energy to electricity. Its nominal output is 5 megajoules per square meter per day. Supplying the 600-megajoule daily lighting budget then takes about 120 square meters of collecting area. That is a calculation under a stated average, with electrical storage and distribution still to account for. The direct-sunlight case uses a different crop-yield assumption. A comparison on an actual site would have to use the local spectrum and timing, the expected edible yield, and the architecture of the habitat. Here the calculation shows what changes when an instruction replaces a variable natural input with a controlled input.

There is also an output on the power ledger. Heat must be rejected over time. In a simple cold-space limit a radiating surface at 293 kelvin, with emissivity 0.9 and a clear view to cold surroundings, radiates about 380 watts per square meter. If the lighting load is 6.94 kilowatts, a corresponding radiating area is about 18 square meters. The calculation divides the assumed load by the Stefan–Boltzmann flux and omits sunlight on the radiator, intervening structures, and heat transport to it. On a planetary surface the surrounding temperature and atmosphere change that accounting. A large enough electrical supply answers only the input side of the question.

The habitat whose average looks favorable still has to carry itself through the interval when its provisions are least favorable. The builder can pay for that interval with stocks, spare productive area, steadier energy, or a recovery program. Their comparative cost is something the Earth experiment can record. It will give a future receiver a better instruction than a single annual yield number.

19 The Slow Sinks

In the Arizona desert, a glass roof enclosed a forest, cropland, water, soil, concrete, and people. The designers of Biosphere 2 meant to watch an ecosystem live through its cycles inside a measurable boundary. During the first sixteen months of closure, the oxygen fraction in the air fell from about 21 percent to about 14 percent. A person looking for the carbon dioxide counterpart of that oxygen loss had to look beyond the air. Microbes were consuming organic matter in the soils. Carbon dioxide produced by respiration was reacting with the concrete, storing carbon in solid calcium carbonate. Investigators traced a change in the atmosphere to a change in the building (Severinghaus et al. 1994).

The connection is easy to grasp on paper. In a simplified model, oxidation of one mole of organic carbon consumes about one mole of molecular oxygen and produces one mole of carbon dioxide. If the carbon dioxide enters a mineral, it disappears from the air record while the oxygen stays consumed. Real respiration varies with the fuel, and real concrete contains a mix of reactive phases. The simplified reaction still gives us a useful accounting scale. One mole of carbon weighs 12 grams. One mole of oxygen gas weighs 32 grams. An unrecorded transfer on the order of 120 kilograms of carbon could therefore accompany roughly 320 kilograms of oxygen consumption in the simple oxidation example. This is a scale illustration. Calculating the oxygen mass from the reported change in Biosphere 2 would require its air inventory and the details of the atmospheric balance.

At the far end of an interstellar message, no one would be available to deliver a fresh ecosystem. A receiver with a catalog of species and a complete floor plan could still have the Biosphere 2 problem. It would need to discover what the walls and soils actually do. A wall that had been treated as part of the machine might belong in the chemical inventory. A patch of soil that looked like an inert place to hold roots might be a respiratory load. The physical construction and the ecological construction would have to be tested together.

Where the phosphorus went

Imagine a small unoccupied research habitat. Its working inventory includes five kilograms of phosphorus in forms that can reach the living community through the

system's ordinary recycling paths. In a trial year, two kilograms of phosphorus pass through a particular recovery step. The step returns 98 percent of what enters it to a form the system can use again. This rate looks impressive on a dashboard. The unreturned fraction is still 0.04 kilogram per year. Suppose the missing material stays inside the boundary as precipitate or adsorbed material, and suppose the working system has no routine for retrieving it. The gross inventory can remain unchanged while the useful inventory falls.

If the same load and loss continued for 42.5 years, this toy model would transfer 1.7 kilograms out of the useful pool. The accessible five kilograms would become 3.3 kilograms. A second model might say that the loss is two percent of whatever stock is available each year. That model gives about 2.1 kilograms remaining after 42.5 years. The two calculations use the same percentage and describe different processes. The first percentage applies to the flow through one step. The second applies to the whole current inventory. A receiver reading the specification needs the denominator before it can interpret the number.

The years in this example borrow the ideal coasting time from the faster living-cargo comparison. The destination experiment would set its own duration. The comparison makes the scale of a slow loss visible. A resource moving into a sink at the rate of a few grams per month can matter in a project whose scientific horizon is longer than a human grant or a growing season. The result changes if the habitat starts replacing the transferred material from local rock or an external stockpile. The useful question then becomes how much make-up material it needs and how reliably the receiver can obtain it. Physical closure is a measurable relationship between the boundary, the reservoir, and the duration of the test.

The movement into a sink can be part of a reversible process that happens to be slow on the period of interest. Phosphorus inside a mineral can become a future resource. It belongs in the accessible account only if the system has an effective route to bring it back at the rate required. A plan that promises to reclaim it eventually owes us a time scale and an energy cost. The available stock is a practical category. It depends on the operating system and the interval over which it must perform.

Nitrogen can leave the useful account in the air

The same vocabulary applies to nitrogen. Suppose the nominal food account includes 75 grams of protein per person-day. For a planning illustration, assume that 16 percent of that protein mass is nitrogen. The product is 12 grams of nitrogen per

day, about 4.4 kilograms per year for one person. This is an assumed food composition, chosen to give the order of magnitude of the flow. The system has other nitrogen stocks in roots, growing tissue, organisms, dissolved compounds, and gas. It can circulate the same atoms repeatedly, so annual throughput can be large relative to the standing stock of one compartment.

Microbes can move nitrogen among chemical forms, including gaseous forms, and that activity is central to biogeochemical cycling (Falkowski, Fenchel, and DeLong 2008). If a part of the habitat converts reactive nitrogen to nitrogen gas faster than other parts restore it to usable compounds, the air can hold ample nitrogen while the growing system has less available nitrogen. The nitrogen has moved within the chamber from one reservoir to another. Imagine a net transfer out of the useful pool of just one percent of the illustrative dietary flow. That works out to roughly 44 grams per person per year. In a habitat for a thousand people the corresponding transfer is about 44 kilograms per year. The meaning of one percent changes when the flow is large and a test runs long enough.

This is one reason a construction plan needs chemical categories with measurements behind them. A reading for total nitrogen can agree with a reading taken last year while plant-accessible nitrogen falls. A balance over the whole chamber and a balance over the root zone can each be correct. Their difference tells us which transformations matter. When a sealed instrument reports that its total stock has stayed constant, it has answered a question about the physical boundary. The receiver still has to tell how the stock is distributed and how rapidly it can move through useful pathways.

How a sink first appears

Suppose the first crop cycles end with declining yields, and the measured soluble phosphorus falls each time. There are several explanations. A growing organism might temporarily hold more phosphorus in its body. A poorly mixed reservoir might make a sample unrepresentative. A material in the vessel might retain the phosphate. A sensor might drift. We can improve the experiment by deciding which observation would separate these possibilities. A material assay that finds a rising amount of bound phosphorus, coupled with a declining soluble stock and a known input, would be evidence of a transfer. The time series across both reservoirs gives a much better account than a low reading at one port.

The same logic applies to the carbon budget. One nominal person exhaling one kilogram of carbon dioxide a day moves about 0.273 kilogram of carbon per day

through that path. A thousand nominal people would move about 273 kilograms of carbon a day. An unreturned one percent of that flow would be roughly 2.7 kilograms of carbon a day, approaching a tonne in one year. This example refers to the selected respiratory path and holds the nominal flow fixed. It illustrates how the bookkeeping error from omitting a small persistent sink grows with the experiment. A measured percentage for one recycling process cannot stand in for every branch of the carbon cycle.

A receiver could handle a sink in several ways. It might change a material whose chemistry it can characterize. It might include a return path that was missing from the initial design. It might make the quantity of make-up material part of the experiment's permitted inputs. Each decision gives a different account of what the habitat must inherit. We can ask how much of the response can be expressed in a general rule and how much has to be particular to this material and this community. The answer matters to the size and portability of the instruction we send.

This is how a slow sink enters the book's information problem. A reservoir that is easy to omit can change the outcome of a program. No one needs a file describing the position of every carbon atom in a concrete wall. A successful specification might have to bound the wall's reaction with the atmosphere and name the measurements that verify the bound. The shorter instruction would remain credible because its omitted details had been tested. The constraint is learned by following the missing material to where it went.

20 The Second Harvest

A message from the far end shows the first crop in storage. The plants grew, the harvest passed its test for edible material, and the machine knows how to open the store. A second crop now has to start. Its nutrients may still be in a waste tank, its planting material may be in a batch with poor viability, and the machinery that supported the first crop may have spent some of its reserve. It is possible to report the first harvest accurately while leaving its replacement uncertain. The receiver needs a calculation whose test continues through the next deadline.

Three dates on the calendar

Use the reference demand D of 10.46 megajoules a day. Suppose a small facility delivers batches of usable food every 30 days. This is an illustrative schedule for a production system that could include staggered plantings with crops that mature at different times. Each normal batch supplies 33 days of demand, or 345.2 megajoules. Begin at day zero with 40 days of food in store, or 418.4 megajoules. At the test scale the stock is treated as interchangeable edible energy, with no spoilage or extra imports. A real menu requires several separate stock accounts because a shortage in one component can persist even when total energy is ample.

Let R be the stock immediately after a batch arrives. If the next delivery takes Δ days and has usable energy H, the stock just before it arrives is $R - D \times \Delta$, and the stock immediately after is $R - D \times \Delta + H$. Thus a batch can look good when it arrives while leaving too little for the following wait. A margin of three demand-days per delivery is enough to repair a small shortage, but it arrives in discrete increments. The store has to carry the load until the increments arrive.

Give the chamber one normal batch on day 30. On day 60, suppose a fault leaves only half of the planned edible energy available. This is a stipulated loss in usable yield, independent of what the plants weighed. Assume no further delivery until day 90. In units of days of demand, the arithmetic is as follows.

Date	Stock just before delivery	Stock just after delivery
Day 30	10 days	43 days
Day 60	13 days	29.5 days
Day 90	0.5 day of demand unmet	Test has already failed

At day 60 there is a harvest to display and food in the store. At day 89.5 the store is empty. Record half a day of unmet demand separately from the zero stock. A batch arriving on day 90 can start meeting demand from then on. The earlier shortfall stays in the record.

The same calculation gives a design requirement for the second batch. Immediately before day 60 the chamber has 13 days in stock. To last until day 90 it needs at least 17 more days of food at that point, equivalent to 178 megajoules. If the test requires a continuing reserve of ten days, the second batch has to add 27 days of demand, about 282 megajoules or 82 percent of the normal batch. That floor would be crossed at day 79.5 after the half-yield batch. An alternate supply preserving the floor has to arrive by then. If it comes as one delivery, the necessary additional stock is 10.5 demand-days, about 110 megajoules. In the illustrative dry ration that represents about 6.5 kilograms. The amount is modest only if a substitute can actually reach the chamber by the required date with the necessary composition.

Those figures expose a choice a receiver can make before it starts. It can fill an extra reserve, arrange for an independent producer to supply the shortfall, or change the harvest spacing. A stored reserve carries a holding cost. An earlier harvest has to mature in time. An independent producer has to retain power, water, and processing when the first production path fails. The schedule puts those alternatives in common units and makes their assumptions inspectable.

The stock we put back into the ground

Food is one product of the first harvest. Planting material for the next crop is another. Imagine a trial with 100 usable units of planting stock, each unit defined by a predeclared germination or propagation test. Sowing the planned canopy takes 80 units. The first batch of edible material reaches its target, and the propagation branch yields 64 usable units. Twenty of the original units remained after sowing, so the stock at that checkpoint is 84. It can sow the planned canopy once more. If the next crop also returns 64 units, the remaining stock is 68 after the second

sowing and harvest. Two food targets could have been met while the capacity to plant a third time fell below the level the experiment requires.

The model can be written with the same reserve equation. After each crop, the change in viable planting stock is the usable output from propagation minus the stock committed to sowing. A deficit of 16 units each time is visible only if the receiver keeps that account. It may change its propagation method or maintain a separate nursery that restores the stock. Either response comes with its own space, light, time, and material requirements. The failure in this example concerns the stipulated planting schedule. Predicting food yield from the remaining stock would require further measurements.

Other stocks have the same property. A residue may be rich in an element while the treatment system takes too long to return it to plant-available form. A recovery reagent can disappear while output water looks normal. A filter can support the first two cycles and reach the end of its useful life before the third. Each test must place its checkpoint where the next cycle needs the material, and it must count the work needed to put the material there. This is where a list of starting materials becomes a description of an operating system.

A recovery that can be checked

Run the experiment first without an occupant. Give the simulated load its schedule, declare the reserve floors and permissible imports, and log the location of stocks. Compare post-harvest inventories across cycles and check the intervening minima. After a disturbance, the evidence should include the date a substitute became available, the amount of reserve used to bridge the interruption, and the date the reserve was restored. If output returns to exactly the ordinary draw, the reserve stays depleted. A successful restart alone cannot close its account.

For a simple recovery scale, suppose a separate repair leaves a ten-day food deficit and the chamber can provide its normal 33-day batch every 30 days thereafter. It can add three days of stock at each delivery while continuing to meet demand. Four uninterrupted deliveries are then needed to restore at least ten days of reserve. This arithmetic presumes the store can carry the load between deliveries and has space to keep the surplus. A stockout before the first of those deliveries would change the result. An acceptance test could deliberately introduce a second interruption before the fourth delivery and compare the observed margin with the predicted one.

The most revealing disturbance may be a common one. Two crop rooms that use the same water-treatment branch can both lose their next harvest when that branch stops. A second room has demonstrated independent support only if the tested fault leaves it able to produce usable food in time. A schedule that repairs the treatment branch and restarts both rooms gives evidence about repair. A schedule that preserves the load during the outage gives evidence about continuity. The receiver needs to know which result it has earned.

The second harvest thus asks for evidence about a set of repeatable transformations. The machine has to put edible energy in the store, return nutrients to service, maintain usable planting stock, and preserve the reserves that cover the next interruption. It has to do so with a recorded dependence on power, local material, and replaceable parts. Each successful cycle narrows the uncertainty about what was established at the destination. A measured deficit tells the next experiment where to look.

21 A Biosphere for a Community

A one-person account makes the quantities visible. A hundred-person account makes the connections visible. Picture a receiver that has reached the point of growing a crop and sustaining a research load equivalent to one adult. It could try multiplying every structure in the chamber by one hundred. A century's worth of operating experience in the original chamber would still leave open what happens when the new crop rooms share pumps, electrical supplies, and an atmosphere. The first multiplication is valuable because it gives us a reference from which to measure everything that changes.

Use the nominal adult in the previous chapters as a unit of demand. At 0.84 kilogram of oxygen per person-day, one hundred people account for 84 kilograms per day and one thousand for 840. The carbon dioxide flow grows from one kilogram per day to 100 and then to 1,000 kilograms. Dry food rises from 0.62 kilogram a day to 62 and then 620. On a 365.25-day convention, the thousand-person food account reaches about 226 metric tonnes per year. With the illustrative three kilograms of potable-water flow per day, the thousand-person water account is about three tonnes a day. These are rates of use or exchange. The values leave open how much material the habitat makes available again, which material arrives from outside, and what must be held in reserve.

In the earlier food-area model, 60 square meters of crop canopy is one illustrative planning case per nominal adult. That scales to 6,000 square meters at a hundred adults and 60,000 at a thousand, or six hectares of active canopy. A three-tier building could arrange six hectares of canopy over two hectares of occupied cultivation footprint before making room for circulation, machinery, and storage. The stacked area would have the same requirement for light at the canopy under the same growth assumptions. Artificial lighting at roughly 6.94 kilowatts of average input per person in the earlier example would grow to 694 kilowatts and 6.94 megawatts. The numbers concern lights. Processing food, moving water and air, manufacturing replacement parts, and moving heat out of the facility would add to the power account.

Suppose the destination were the site of the thousand-person calculation and the lighting branch ran at that average for a year. It would take about 61 gigawatt-hours of electricity, using 8,766 hours for the year. Six hectares can sound modest to a farmer. A load measured in megawatts feels different to an arrival probe. It creates

an order of construction. The receiver has to establish the power and heat-handling functions before it can learn much from turning on all the lights. It might eventually use sunlight or a different means of cultivation. Each choice has to be compared at the same canopy, crop schedule, and definition of edible yield. A surface figure stated without its assumed input can hide the machine the colony needs.

What gets cheaper when it is shared

Some needs can be shared. One instrument can serve many compartments. A common structure can enclose several rooms. The ratio of surface area to volume can improve in a larger vessel of similar shape. In a geometric example, increase the volume of a cube by a factor of one hundred. Its surface area grows by one hundred to the two-thirds power, or about 21.5. If the envelope per square meter cost the same, that is a large material saving relative to building one hundred separate cubes. A real chamber has pressure loads, doors, thermal requirements, and internal partitions. Its shell cost need not follow this ideal scaling. The example gives the reason to ask which costs grow with area and which grow with the number of inhabitants.

Shared functions also give failures a wider reach. A single nutrient distributor can reduce the amount of plumbing and work. A problem in that distributor can then stop crops in every room. Independence matters because the goal is to keep useful functions available while some components are out of service. Dividing a chamber into separately observable modules could let a receiver test and compare them. If they all rely on a common power supply that fails, the separation has protected little of the immediate output. A system drawing shows which dependencies are shared. Data from deliberate interruptions show which shared dependencies are consequential.

The economic question is therefore about structure as well as total material. A receiver could group the rooms that tolerate a common interruption. It might give some critical functions a separate physical supply and keep enough stock to wait while a second source starts. These choices change the capital material, energy, and information accounts. They also change the test of reliability. Comparing systems at the same population makes those costs visible. A design that uses the fewest parts on its best day can be more demanding to restore on a bad day.

Four rooms and a common disturbance

Consider four equal crop rooms. In a toy statistical model each produces the same average yield, and the standard deviation of one room's yield is 20 percent of its own mean. Independent fluctuations would put the relative standard deviation of total yield at 10 percent. If every pair of rooms has a correlation of one-half, it rises to about 16 percent. If they move together perfectly, it is 20 percent. With four equal variances and pairwise correlation ρ, the relative spread of the total is $0.20 \times \sqrt{[(1 + 3\rho)/4]}$. Correlation can erase much of the apparent benefit of multiplying copies.

Diversity may help if alternative producers contribute to the same useful function under different circumstances. It can also add relationships whose effects have to be understood. In an Earth grassland experiment, plots with higher plant diversity maintained more stable total production over time. Tilman, Reich, and Knops attributed that result to averaging and higher productivity of mixtures. Their data gave no evidence of a covariance effect (2006). That experiment establishes what happened in those plots. A receiver should treat the pattern of correlation across the actual crop rooms as a new measurement. Even two different species can share a vulnerability to a common lighting loss or a disruption in water delivery.

The same principle applies to less visible services. A habitat needs a route to return waste to useful material, a route to keep gases in balance, and a means to move water between compartments. Counting the organisms assigned to those services gives a starting description. A deeper design account follows each process under a perturbation. Does a second community perform the same conversion when the first is limited? How quickly can it take on the load? Does it use an independent resource? The answer may be easier to express as a measured capacity at several conditions than as a list of species. The functions and their variation are part of the payload if the receiver cannot learn them reliably on site.

A harvest that arrives late

Take a hundred nominal adults consuming the illustrative 62 kilograms of dry food each day. A 35 percent decline in edible output for twenty days would use 434 kilograms of stock, if production had previously matched use. Suppose output returns with a sustained 10 percent surplus after the disturbance. At 6.2 kilograms of surplus per day, seventy further days would be needed to replace what was drawn down. The first day of recovered production is only the first point of that return

trajectory. A second disturbance during the seventy days would find the habitat with less reserve than the first.

This calculation assumes a stock of interchangeable dry food and a uniform shortfall. Actual losses can involve one food component or one harvest date. A reserve has different value when it fills a gap in a particular component. In a model, the response can be described as a vector of stocks and a schedule of harvests. In an experiment, measured outflows and recovered reserves give the vector a physical meaning. A receiver that reports total biomass while leaving edible stocks unspecified can miss what caused the deficit.

One hundred and one thousand are arithmetic examples. A permanent human community would bring separate questions about health, population continuity, repair, and governance. The figures let us compare a small demonstrator with the facility a future society would need, and they point to the functions whose failure becomes more costly with scale. A colony of Earth life needs a habitat that can keep its common infrastructure working across interruptions and replace the stocks on which a generation depends. A successful small build would be a basis for measuring the next step.

22 Deciding When to Hand Off

Imagine that the receiver sends back an image of green plants. The image takes years to reach Earth. It tells us that a particular camera saw leaves at a particular instant. The message we actually need is different in kind. We want a record of how much edible material was made, where its elements came from, what the production cost, what happened after a loss, and whether the materials needed for the next harvest were still available. Photographs would be valuable evidence of a build. They would be part of a longer experiment.

The first-cell handoff in Chapter 10 asked whether a cell would keep functioning after the printer stopped arranging its components. For a habitat the same test has to name the machinery that continues to run. A pump, a valve controller, and the supply of light may all be declared parts of the environment. Their energy, upkeep, and replenishment belong in the record. One version of the test would allow ordinary operation of the apparatus while removing the interventions that brought the system into its founding state. Another version would ask whether the receiver can maintain and repair the apparatus over a specified period. Those are related capabilities. Each can be tested under a contract that says what crosses the boundary.

An early experiment belongs in an unoccupied enclosure on Earth. A simulated load can stand in for the chosen material flows of a person. The experiment can then challenge the pathways without treating a participant as the instrument that discovers a failure. The burden, timing, and permitted make-up inputs would be specified in advance. A second laboratory could set up the same account using a distinct chamber. The exercise would teach us how much of the construction instruction is portable before we trust the greater independence required by distance.

Compare corresponding points in the cycle

A cycle can include planting, growth, harvest, processing, and a pause. At harvest a large fraction of a nutrient may be in tissue. Later it may be in storage or in the recycling path. An investigator who compares those readings directly could mistake a change of phase for a long-term change of inventory. Instead, choose comparable points in successive cycles. Measure the useful stocks available there, including forms that can be returned within the required time. Record the net movement of

material across the test boundary between the two measurements. The remaining change is evidence about the system itself.

Suppose two cycles begin with twenty kilograms of a particular useful stock. The first run ends at nineteen kilograms after an external addition of two. In the budget, the internally generated change is a loss of three kilograms. The second run ends at nineteen with no additions, for an internally generated loss of one kilogram. The same final reading conceals a difference in the rate at which the system can carry itself. If the first run counted the addition as part of ordinary practice, that entry belongs in the field report. A future receiver needs to know which supplies have been built into the stated result.

Some of the strongest evidence comes from a perturbation that affects one function enough to identify it. Suppose part of a crop becomes unavailable in an unoccupied trial. Record the decrease in available food and the dates when substitute production appears. Measure the withdrawals from reserve and check when the reserve recovers. Track energy and nutrient demand during the response. A substitute that can supply biomass while consuming the stocks needed for the following cycle has solved only the immediate deficit. Its trajectory through recovery tells us how much help the replacement actually provided.

The significance of a trend

The amount we can measure and the amount we need to know are separate. Imagine an accessible stock whose readings fluctuate with mixing and sampling by roughly two kilograms. A steady loss of 0.02 kilogram per day would be hard to distinguish from the noise over a few days. Over a year it would add up to about 7.3 kilograms. A long experiment can reveal a trend that a short demonstration misses. It can also accumulate errors in calibration. An independent assay and a second material balance would help distinguish a loss of useful stock from a change in the measurement.

For a simpler calculation, imagine that one unbiased estimate of stock change has a two-kilogram standard deviation. Six independent estimates with the same variation would give a mean with a standard error of about 0.82 kilogram. That calculation describes the uncertainty of the mean when the errors are independent. A single drifting scale used in all six trials can preserve the same bias in every estimate. Independent tests should therefore mean more than pressing the start button again. A second apparatus, a different measurement principle, or a separately calibrated instrument can make the result informative about a fault shared across repeats.

Length matters because biological systems can develop a slow variable. An apparent steady regime might coexist with a mineral accumulating in a pipe, a change in the seed stock, or a decline in one organism's contribution to recycling. The conditions for a good result should give those effects an opportunity to appear. We can also use separate observations to understand how quickly they are likely to matter. The oxygen decline in Biosphere 2 took months to reveal the full reach of a connection between soil and concrete. A project that ended after one productive harvest could have told a different story about the same chamber.

How much evidence does a series of successes buy

Suppose twenty independent trials meet their stated criteria. In a simple model with a constant failure probability, a 14 percent chance of failure per trial would still give about a five percent chance of observing no failures in twenty runs. The figure follows by raising the probability of success to the twentieth power. At a hundred trials the corresponding failure probability is about three percent. Such arithmetic is helpful because it sets a scale for the evidence. It depends on how a trial was defined, the independence of the results, and whether the distribution of conditions changes. Running a hundred copies of one favorable season leaves the design's response to a different season largely open.

A reporting plan should keep incomplete and aborted trials visible. Imagine that an observer disconnects a failing branch to protect expensive equipment. The reserve record after that event tells us what the apparatus can do with intervention. The disconnected branch tells us something about the behavior before intervention. A count of successful demonstrations has a clear interpretation only if readers can see what an attempt includes. These observations matter when the method is part of the argument for sending an instruction across a distance that precludes quick supervision.

Handoff would come in stages. First, establish in an unoccupied test that material cycles operate and that the intended response to disturbances is observable. Then test successive starts and increasingly demanding sequences of inputs, with explicit limits on external help. Compare independent builds against the same metrics. A future proposal involving people would bring additional criteria and independent judgments far beyond these experiments. A device that grows food for a simulated load has established a laboratory capability. The work of demonstrating a human habitat would still be in front of it.

The remote version makes the specification sharper. At a distance of 4.25 light-years, a measurement and reply would take at least 8.5 years in the Earth frame if the endpoints were treated as fixed. A cycle that fails in a week requires a local response. An anomaly that becomes apparent over a year still requires the receiver to recognize what it means and preserve the conditions for a second measurement. Earth can provide science and eventual advice. The short-timescale competence has to be present at the far end. The autonomy needed there belongs in the same ledger as the radiation the instruction was meant to avoid.

A useful closing record would show the largest deficits observed, the paths that replaced them, the cost of the recovery, and the stocks left at comparable points through the trial. It would also show where the evidence ends. The result could be a design that keeps certain cycles operating with a quantified dependence on material from outside. That would be a success with a readable boundary. If we later broaden the boundary to include a larger community, the earlier record would show which capabilities have to be repeated, which can be shared, and which still belong to the world on which we tested them.

23 A Workshop That Can Arrive Alone

Send eleven workshops. The first to wake may find a quiet route to a working orbit. Another may discover that the intended brake can no longer deploy. A third may arrive with its radio intact and the main archive unreadable. The fleet makes these events useful if each craft can do something with the circumstances it finds. Its smallest complete mission is the work of one receiver, with its own way to wake, find its bearings, attempt capture, recover its instructions, and report what happened. The rest of the fleet may be years away or already gone.

An illustrative allocation gives the three nearby groups three workshops apiece. The two farther groups get one each. Within the Alpha Centauri group, a craft's plan specifies an actual approach and the body or star relevant to its capture. A shared label on the table gives the craft a common scientific program. It carries no promise that three arrivals can reach the same orbit or even observe the same site.

This allocation is a way to ask what a fleet could learn under a fixed count of vehicles. The more distant singletons can return valuable site records even if no second instrument can confirm them in place. A three-craft group can compare several independent arrivals and give separate reports about a shared question. All five routes have to earn their places with their own propulsion, capture, power, protection, and communications estimates. The number of craft fixes none of those quantities.

Give every craft a complete first job

A workshop travels cold in the sense established earlier. Its major fabrication and computing systems spend most of the crossing dormant. The restart chain, surviving archives, tools, and authorized nonliving starting inventory all go aboard the same craft. Once awake it can check its power, choose among its authorized actions, verify a nonliving reference build, and measure whether local material is recoverable. Its own clocks, instruments, and controller handle decisions that must be made during approach and commissioning. It can report to Earth while the record is still useful to itself. A model that puts the receiver on one craft, the archive on a second, and the machinery for reading it on a third would make rendezvous or delivery between vehicles a prerequisite for the first result.

Within each group, the three designs can emphasize different questions. The first might carry more survey instruments and qualify a route to a stable working site. The second might carry a separately built startup chain and repeat the core measurements on another route or with a different set of instruments. The third

might put more of its mass into contained construction and measurements of renewal. Each begins with the same authority to survey, to build a nonliving reference object, and to establish whether its own experiment can proceed. The emphasis determines which optional work it can do in greater depth. When a workshop goes missing, the others still have their initial jobs and their local standards of evidence.

Capture and revival offer genuine reasons to vary the design. One candidate might depend on braking that begins well before approach and use starlight for final revival. Another might carry more stored capacity for the first switches and reserve an independently qualified propulsion stage. The useful comparison is between complete arrival sequences, each with a budget for the distance and time at which its equipment must begin to work. A favorable estimate for an already bound orbit says little about a vehicle that still has to become bound. A candidate whose brake cannot be qualified may still be worth considering as a survey flyby if its passing trajectory and observing window justify that narrower assignment.

The archive has a similar choice. All eleven can carry a self-contained core with the reader, units, reference tests, permissions, and instructions for the craft's first job. One design can devote more protected storage to a full preloaded edition. Another can carry a smaller edition that remains useful while it listens for optional additions from Earth. Readers that fail in different ways can give more assurance than many copies readable by one mechanism. The craft logs the edition it actually opened and which blocks survived. A later transmission can extend the work once the receiver has verified it. The first report has its own route to completion from material already aboard.

Stagger the waves and keep the evidence separate

Waves spread the fleet across launch opportunities and manufacturing histories. In one schedule, each wave sends one workshop toward each of the three nearer groups. The two singletons depart when their own route and vehicle are ready. The first departure can teach later builders about problems observed near home. A report from another star will arrive on a different schedule. A later workshop can carry a revised edition when it is ready to leave, and it still gets a plan that makes sense if the earlier workshop sends nothing. Launch dates, expected arrivals, and the earliest plausible return reports belong on separate lines of the schedule. A stagger can improve the experiment even when the first news of capture comes long after all eleven vehicles have departed.

The value of three arrivals depends on how they can fail together. Three versions built around one vulnerable reader can lose the same record. Three teams drawing from one contaminated source can obtain the same misleading result. Independently

qualified instruments and starting kits with separate custody give the comparison more force. Use a common definition of the test, maintain records of the differences, and preserve original observations before the teams exchange interpretations. The craft exchange measurements as signals. Their physical inventories and contained lineages stay at their own sites. A success elsewhere can change what a receiver looks for or which instructions it studies. It supplies no living material for a local build.

Carry the result as far as the site permits

At the first-cell stage, each qualifying workshop declares what came from Earth, what it made, and what it obtained at its site. It attempts the handoff in its own enclosure. The receiver can then show whether a pedigree continued after it withdrew organized assistance, whether descendants renewed their machinery, and which inputs crossed the boundary. A second workshop can repeat that experiment with a separate kit. Similar records would establish more than a photograph of division from either site. Different records can reveal how the result depends on the equipment, the local conditions, or the starting inventory.

A workshop that reaches the next stage faces the second harvest. It must show where the usable food, replacement planting stock, recovered inputs, and restored reserves came from. A second chamber's good output cannot refill this chamber's balance sheet. The fleet's value is that it can compare outcomes from independent balances, including the one that runs down after a convincing first harvest. The local chamber remains the unit of proof. Authorization for a contained build depends on the evidence at that site and the permissions already in force. The receiver can preserve its archive and keep observing when those conditions are unmet.

Each route belongs in the final record. Give the date and geometry of the attempted encounter, the last established instrument state, the archive edition, and the evidence for each completed test. A failed capture can leave a flyby record. A working orbit can yield a materials survey even when the printer cannot qualify its chamber. A missing report leaves an unresolved route, with the last verified event and a stated limit on what can be inferred from silence. Keep every launch in the denominator. Eleven attempts may leave one full account of a workshop that woke, stayed, built, and watched a local system take over. That one account can be read on its own, beside the ten records that explain how much of the fleet's work remains open.

24 Booting the Receiver

Picture the first apparatus in the receiving system. It may be in orbit. It has arrived alone, out of contact with the people who last checked its wiring. Its first work is to find its reference stars, measure what power it can sustain, and tell Earth that a receiver exists at the far end of the link. A message that arrives before this work succeeds can cross the right patch of space at the right wavelength and accomplish nothing. The receiving apparatus is part of the proposal from the first day of the voyage.

If one uncrewed payload is preferable to a ship carrying a population, we should be able to say what was taken off the ship and what went aboard in its place. The receivers, power equipment, sensors, and tools still make the crossing. They need structures that keep them intact and some means of protecting and recovering their records. The habitat shell, much of its working inventory, and any later population could come from local resources only if the arriving tools can turn those resources into usable material. This is the premise that earns the architecture a smaller freight bill. The premise belongs in the list of things to test.

Mass illustrates the value of the distinction. One tonne traveling at 127 kilometers per second has about 8.1×10^{12} joules of kinetic energy in the departure frame. A second tonne at the same speed adds the same amount. These figures describe the moving payload. A propulsion estimate also has to account for fuel, exhaust, beaming losses or whatever carries away momentum, and a means of slowing down. There is a real incentive to substitute local bulk for imported bulk. There is also a real price for every tonne of tools that makes the substitution possible.

The first work is physical

A remote receiver needs a contract in the same sense that the cell printer needs one. An operation in the build language has to correspond to a machine behavior with a known domain. Suppose an instruction requests a pressure boundary of stated dimensions and tolerances. The arriving hardware needs a way to obtain the material, shape it, join it, inspect the joins, and recognize an unusable part. Some of these steps can be supplied by local processes and some by imported mechanisms. Their dependency remains even if the instruction takes one line.

The Mars Oxygen In-Situ Resource Utilization Experiment, or MOXIE, is a useful precedent at this scale of argument. A machine brought to Mars used a gas present

there and made a measured product. NASA reported 122 grams of oxygen over sixteen runs, with production reaching 12 grams per hour under favorable conditions. Its design and manufacture were done on Earth. The experiment established a particular transformation under Martian conditions and gave designers measurements they could use for a larger machine (NASA JPL 2023).

For scale, the one-person oxygen flow used in this book, 0.84 kilograms per day, averages 35 grams per hour. MOXIE's best reported hourly production was about a third of that numerical rate. Its complete reported production was about one seventh of the adult's nominal daily flow. The comparison is between quantities, in settings with different goals and duty cycles. A receiver for a habitat would need production, storage, distribution, and verification that could keep working across many cycles. A small, measured transformation earns one entry in the capability list.

Picture the fabrication chain as a set of promises, each with a sample it can make. A tool that can test an Earth-supplied reference sample has cleared a different hurdle from a tool that can retrieve and refine its own sample. A machine that fabricates the enclosure of a replacement part has cleared a different hurdle from one that can build its sensor, controller, and bearings. A proposed receiver should name which specialized parts it imports, which spares it carries, and which replacements its local workshop has actually demonstrated. Otherwise a closed factory can appear on the drawing while the machine still depends on a long, invisible supply chain.

The site gets a vote

The receiver needs to measure the conditions under which its tools can work. A remote spectrum can suggest a target worth visiting. A material sample and the evidence of repeated local tests provide a different basis for the contract. Assume a site appears to contain water. The workshop still has to find it, extract it in useful form, separate unwanted constituents, and find an energy source that can pay for the operation. The first useful result is a measured cost per unit of recovered material, together with the range of conditions over which the cost was measured.

Heat is part of the same account. Suppose a later test uses seven kilowatts of electric power, roughly the canopy-lighting example in this book, and suppose essentially all of that power eventually appears as heat to reject. An ideal surface near 300 kelvin radiating into a cold sky at emissivity 0.9 emits about 413 watts per square meter. Seven kilowatts would then require about seventeen square meters of effective emitting area under those assumptions. A real radiator depends on its

temperature, its view of the surroundings, solar loading, and how the heat gets to it. Even the ideal estimate puts a new object on the freight and fabrication list. A heater can be bought at the price of its input power. A habitat also buys the path by which that energy leaves.

The site affects communications as well as fabrication. The transmitter needs an orientation reference, enough power for a scheduled contact, and a clear interval in which to point. Stored records need a path out when the preferred instrument fails. If the first receiver is a single machine with a single optical head, losing the head could turn a successful crossing into a lost mission. Several observing and communication paths cost mass. Their benefit is an experiment that can still be interpreted after a component fails.

Autonomy belongs in the contract for the same reason. At 4.25 light-years, asking Earth whether to make an emergency correction and receiving a reply would take at least 8.5 years, even before operators deliberate. A receiver can send its evidence home and wait before attempting an irreversible action. During the wait it needs local behavior that conserves energy, preserves data, and keeps the enclosure in a known state. The machine can be designed with a wide range of authorized routine actions and a narrow range of authority to alter its surroundings. Long delay makes that distinction measurable. The safe state itself has to be sustainable for the time the mission asks it to wait.

Making the first report useful

A useful report from the receiver would give the date and location of the measurements, the instrument states, the calibration results, and the extent of local verification. It would tell Earth which materials were found in quantity and which remained hypotheses. It would say how much useful power had actually been delivered across periods of shade, poor pointing, and scheduled maintenance. Such a report can be compared with the assumptions made at launch. A list of successful operations without the conditions that enabled them would be hard to use in the next design.

The order of the mission follows from this report. First establish a functioning apparatus and a basis for deciding where its work may occur. Then verify power and material flows and build nonliving structures that can be tested and remade. Establish the archive as a readable and recoverable record. A later stage might try to assemble a contained working cell or a material cycle whose outputs can be

accounted for. The evidence from each stage can narrow the choice of what comes next. The sequence is a way to learn how far the receiver's reach extends.

A printer on Earth might rely on a technician to notice that its optics have shifted or that a nominally pure input behaves differently in this batch. At the other star the equivalent observation has to come from an instrument, a local rule, or a report that arrives years later. We should therefore account for the instructions that make the receiver an adequate witness to its own actions. Those instructions matter alongside the commands that move matter. They let us tell the difference between an expected outcome and something the machine has actually achieved.

25 Archive of a Living World

Suppose the receiver has established contact and reports enough capacity to maintain an archive. The temptation is to transmit a catalogue of genomes. A catalogue can be a fine first volume. A living world has a much longer table of contents. A seed's genome, the structures in the seed, the environment in which it develops, and its relationships with surrounding organisms each carry different parts of the work. A useful archive records which of those parts it can supply, and in what form.

We can keep three descriptions in view. One is a readable record, which might be a genome sequence, an anatomical atlas, or a map of observed interactions. Another is a generative instruction, which a declared machine can use to bring about a specified outcome. A third is a set of tests for deciding whether the outcome was achieved. The distinction shows up in small examples. A file can correctly name the molecules a cell contains. A build has to put them in suitable states and relationships. A test has to establish that their joint operation continues after the assembler withdraws. Each document has a different job.

The Earth BioGenome Project gives a genuine precedent for a large biological catalogue. Its published proposal describes sequencing and characterizing the genomes of Earth's eukaryotic biodiversity for a wide range of scientific and conservation purposes (Lewin et al. 2018). An interstellar build archive could use records of that kind. It would add provenance, conditions, and evidence about which records a recipient could turn into functioning systems. Genome cataloguing makes a powerful body of evidence available for that next question.

From a sequence to a transfer time

An uncompressed four-letter sequence takes two bits per base in a simple encoding. The *E. coli* chromosome used earlier takes 9,283,304 bits. At one thousand useful bits per second the bare sequence alone would take about 9,283 seconds, or two hours and thirty-five minutes, to send. The T2T-CHM13 human reference reported by Nurk and colleagues is about 3.055 billion base pairs. A simple encoding of one reference of that size would take about 6.11 billion bits, or about seventy-one days at the same useful rate. The published reference represents a specific assembly, with its own chromosome coverage. An archive intended to represent variation across individuals needs additional records (Nurk et al. 2022).

Now imagine an archive of one thousand reference genomes averaging one hundred million bases each. This is an arithmetic example, with no claim about how many or which organisms would found an ecosystem. It contains one hundred billion bases, taking two hundred billion bits, or twenty-five decimal gigabytes. At one megabit

per second it would take about 2.31 days to transmit the raw sequence payload. At one kilobit per second the same payload takes about 6.34 years of continuous service. An archive of a million references averaging a billion bases would be two quadrillion bits, or 250 terabytes. At a useful megabit per second its raw transmission would take about sixty-three years.

Useful rate and wire rate need separate labels. If the wire carries a thousand bits per second and a hypothetical error-correcting code spends four of every five bits on payload, the usable rate falls to eight hundred bits per second. The twenty-five-gigabyte example then takes about 7.92 years before scheduling gaps and additional retries. If the link is usable for a quarter of the calendar time, the same example stretches to almost thirty-two years. A data sheet that names only the fastest demonstrated burst can conceal the time required to deliver a usable library.

The NASA Deep Space Optical Communications demonstration sent a video at 267 megabits per second from about 31 million kilometers away in December 2023. That is an actual performance point for a particular transmitter and receiver. A distance of 4.25 light-years is about forty trillion kilometers, more than a million times that demonstrated separation. With the same beam pattern, collected power at a fixed receiver falls with the square of distance in the spreading regime. A link budget for the new distance needs its own apertures, pointing, background, light source, and coding. The rates in our calculation are planning cases whose feasibility would have to be earned by a separately specified link (NASA JPL 2024).

The archive has to survive reading

The sequence must arrive in a form the receiver can check. Suppose errors remaining after a communication step occur at an average of one in a billion bits. For the example archive of two hundred billion bits, that expectation is two hundred errors. Even an apparently small residual rate can leave the completed file with mistakes. Which mistakes matter depends on where they occur. A damaged annotation and a damaged instruction that controls a moving part demand different consequences. In both cases we should be able to detect damage before treating the item as authentic.

A stored block can include a short integrity check and be compared with an expected value. Such a check gives evidence about whether the block survived. Correction needs redundancy. Here is a simple case. Store two binary blocks and a third formed by their bitwise exclusive OR. If one data block goes missing and the other two are intact, the missing block can be recovered by applying exclusive OR to the two remaining blocks. The cost is a third block for every two data blocks. A practical archive would choose a scheme for patterns of error and failure expected in the link and storage system. Independent copies can help when a single hardware fault removes a group of blocks together.

The mission can use correction locally without waiting for a report from Earth. It can also receive repeated blocks at different times or from separate transmitters. The return link will still be valuable for identifying missing material and commissioning later versions. Because one request and answer cycle takes at least 8.5 years at the example distance, a design that depends on a fresh request for each bad block wastes a large part of the calendar. Receiving a self-checking edition, with redundancy that permits local repair, is a different service from a conversation that assumes a quick acknowledgment.

The archive also needs to say what its symbols mean. In the blind bitstream exercise, periodicity and repetition survived after we withheld the genome's annotation. Such patterns can suggest a reading frame or a copied passage. They do little by themselves to say which physical operation a block authorizes. The same ambiguity would confront a receiver working through an undocumented library. Formats, units, version identifiers, agreed reference tests, and executable checks are part of the transmitted knowledge. Many can be shared across packages. Their shared use is a genuine saving only when the receiving copy is intact and understood.

What belongs at the front

A first installment could contain a complete small example, including what the receiver must already possess. It could show a format decoded by the actual remote machine and a physical object tested with the available instruments. The next installment could add a larger catalogue organized around measurable functions. The complete lossless record of a specimen, the simplified instruction for producing a functional class of objects, and the measurements that justified the simplification could all be stored. Later researchers might care about different features and choose a different compression.

One thousand reference genomes represent one thousand lists of bases, under stated conventions. A viable community represents linked populations with particular starting relationships and a history. A set of photographs of a watershed and a collection of all its species would give a future engineer useful evidence. The engineer would still need to know where water goes, when floods occur, what persists through a dry season, and which changes are reversible. An archive that serves the mission should therefore be organized around both records and operations. It can tell the next reader what was observed, what was built, and what that build depended on.

26 Where a Seed Is Welcome

A receiver arrives, sees a planet at the right distance from its star, and receives a library intended to begin a new living system. What should it do? The strongest answer may be to observe the planet for a long time and keep the library in storage. A proposal to extend life's reach gets its purpose from the value of living systems. That value should also govern what the proposal does when it encounters a system it did not make.

Proxima b is a useful case because it is nearby in interstellar terms and its orbit makes it a target of scientific interest. The label *habitable zone* describes where a planet receives an amount of stellar energy compatible with surface liquid water under appropriate atmospheric conditions. An actual world has an atmosphere, surface, history, and pattern of exposure. Those determine whether it has that water or can support any proposed environment. NASA's overview of Proxima b makes the distinction explicitly (NASA Science 2019).

Observations of the host star also give a reason to separate a mean from an operating condition. MacGregor and colleagues reported a short, intense flare on Proxima Centauri observed across a range of wavelengths (MacGregor et al. 2021). One flare cannot tell us everything about surface conditions on a nearby planet. It shows why a stellar average is an incomplete description for a build that has to survive changes in radiation and energy input. A remote probe can measure how conditions vary while a decision about biological activity remains open.

Observation has to precede intervention

Consider a world with no signal we can confidently call biological. The absence of a signal may reflect an absence of life, a low abundance, a subsurface habitat, a chemistry beyond our detectors, or simply a short observing campaign. An instrument reports the range of things it could have detected under stated conditions. It is hard to turn such a statement into a proof that nothing of interest exists anywhere on a planet. A receiver able to make matter more complex must also have a way to preserve uncertainty in its records and decisions.

NASA's Ladder of Life Detection describes how different lines of evidence can contribute to a case for life. The framework asks what a measurement could mean in its environmental context and how alternative origins might be evaluated. The same discipline helps here because a seeding proposal has two reasons to learn a

site's natural history. It needs to know which conditions can support the apparatus. It also needs to preserve the evidence that allows anyone to decide whether they found life there before the apparatus arrived (NASA Science, n.d.).

A seed release changes both inquiries. A new population could leave chemistry, structures, and genetic signatures that persist after the first observation. Future instruments would have to distinguish local phenomena from descendants of the introduced material or its products. If the new system changes the environment itself, even the interactions we hoped to study could disappear. These are reasons to begin with contained experiments and with a recorded separation between observation and introduction. A launch manifest and a dated set of pre-intervention measurements are scientific evidence as well as project records.

The planetary protection principle already has a place in spaceflight. Article IX of the Outer Space Treaty calls for avoiding harmful contamination in the exploration of celestial bodies and for consultation in specified cases of potentially harmful interference. Applying the treaty to a future interstellar project would require a much more detailed institutional judgment. Its underlying concern is directly relevant. A decision made by one group can alter what other investigators, and later generations, have a chance to encounter (United Nations 1967).

An enclosure can change the question

An orbital test or a separate enclosed habitat offers a different target from changing a planetary atmosphere. The intended outcome is a bounded system whose materials can be counted and whose failures can be identified. The design still has to account for leakage, disposal, physical breakup, and where its equipment goes when it ceases operating. A contained experiment can earn confidence by showing how its barrier performs through simulated failures. Its scale and position should be chosen so that accidental dispersal is part of the evaluation from the start.

For a planet that may already harbor life, the mission can decline to release biological material there. If uncertainty remains high, it can also decline to build a biologically active system in a place where escape could affect that planet. It can receive and preserve a verified archive, make observations, and construct nonliving facilities. These are substantial outcomes for a receiver that has reached another star. They preserve the ability to change course when evidence changes.

The difference between a contingency and an intention matters here. An archive can contain instructions for building a system under declared conditions. Its presence

does not have to grant the receiver authority to execute them. Authorization can be a separate, verifiable input, with a requirement that the receiver has met its current limits and that a review has taken place. This adds years to some decisions. The reason for giving local machinery broad discretion about routine repairs is the same reason for constraining its decisions about release. An action can be prudent at one scale and carry consequences too large to entrust to the same rule at another.

If it were possible to demonstrate a sterile construction site with tolerable risk to surrounding environments, we would still have to decide what belongs in the founding community and for whose benefit. The easiest collection for the printer to reproduce could differ from a collection that is robust across change. A habitat designed for a person may depend on organisms whose persistence is valuable even during periods when the person is absent. This is another reason to characterize systems by trajectories and material exchanges. It gives the destination a place in the evaluation as a system that can be affected, with a history that continues beyond the immediate test.

What can be left for the future

The desired endpoint may change during the voyage. A receiver can arrive after the civilization that commissioned it has learned new things about life, about the target, or about the risks it was willing to impose. It can also arrive before it can receive revised instructions. A mission with such delays needs versions of both its archive and its permissions. It should be possible to read later which permissions were active when an action was taken, and which observations supported it.

One can imagine future custodians looking at a remote habitat and asking whether its founders gave the new system room to develop. They could also ask whether the founders left room for a different system to exist nearby. The first question concerns the design of the habitat. The second concerns its boundary and the site chosen for it. Both are part of the mission's success. A seed is welcome where the people responsible can make a defensible case about what they are preserving, what they are introducing, and what the consequences of a mistake would be.

27 The First Delivery

The first delivery should be small enough that everyone can agree on what was sent and what arrived. Start with several receiving laboratories. Give their declared capabilities to a separate team that writes and sends the instruction. Keep the reference culture out of the receiving laboratories during the test. Send a digital package for a bounded cell-building demonstration and specify which organized materials may already be present. Ask an independent group to measure which physical steps actually took place and what the resulting system can sustain. The laboratory is on Earth. The division of information and capability is the same one that would matter across an interstellar link.

A series of tests can start with the parts we know how to verify. Can an independently assembled, nonliving reference object be made within its tolerances using the transmitted program and allowed machine operations? Can a material transformation be made to recur in a sealed apparatus, measured by observing the inputs and outputs independently? What organization had to arrive embodied in a reagent, a supplied assembly, or the tools? These questions are achievable milestones even while a complete cell build remains a research goal.

This staged experiment also distinguishes several ambitions often put under the word *panspermia*. A physical seed carries working material, perhaps dormant or protected. It has to preserve enough of that organization throughout travel and startup. A transmitted record can be checked and replicated. It requires a receiver able to make the working material elsewhere. A mixed strategy could send some durable organization and transmit later instructions that let it do more. Each strategy leaves a different ledger of transit dose, protection, storage, hardware, and arriving capability. We can ask which ledger gives a working result under the same assumptions about the destination.

Send one version of the instructions to independent receivers and another version that uses a different starting inventory. Give them a common way to report what happened. A machine fault that spoils one run can then be separated from a requirement missing from both recipes. If one team reports a lineage, a separate build can test whether that result depends on an undocumented material or practice. If no team succeeds, record the stage at which each stopped. The distribution of failures is evidence about what to change before the next attempt.

The habitat evidence must come first

Next, send a habitat instruction to a receiving team that builds an unoccupied test enclosure. The receiving team can account for air, water, chemical stocks, energy, and heat. It can substitute a known test load for the person whose requirements

motivated the design. The one-person calculation provides a starting scale for measuring those flows. The experiment can run through more than one cycle, watch for accumulated deficits, and tell the sending team which assumptions failed. A clean test of one particular cycle would be a result to reproduce. It would also suggest where to look for the first slow instability.

The early enclosure can use Earth-grown cultures and plants, with their origins recorded. That trial measures the renewal of a supplied community. The sterile flight version has to qualify a source for each founding function separately. Making one bacterial lineage would establish a way to start that lineage. A crop, a decomposer, and any required partner bring their own construction and developmental demands. A staged program would test those dependencies in contained modules and then measure how they work together. The complete mission has to account for that additional build as well as the balance sheet.

A community changes the numerical test. The simple electric-lighting example in this book is about seven kilowatts per adult's assumed crop allocation. At the same allocation, a hundred adults represent about 0.7 megawatts of lighting and a thousand about seven megawatts, before other uses of power. The first multiplication describes a shared supply burden. It gives us a reason to ask which costs scale with the number of people, which infrastructure can be shared, and which reserves have to remain independent. A hundred trays sharing one controller may still have one failure point. Counting them as a hundred safeguards would make the arithmetic look better while leaving the system as exposed as before.

A founding population raises additional questions. A community has an age structure, knowledge, relationships, institutions, and demands for care. It needs a way to keep skills available and recover from events that affect many people at once. A transmitted record for an eventual human community would raise questions about the conditions under which people could develop and the authority anyone has to decide their future. A serious mission design makes a place for these questions before it claims a destination is ready. The unoccupied cycle is an engineering test with a deliberately bounded claim.

The Earth demonstration can yield a first publishable result in a form other researchers can use. It can name which organisms and parts were imported already organized, how the system was started, the allowed external inputs, the duration and quality of the measured flows, and which cycle ended the run. An unsuccessful result can be equally precise. It might show that the storage reservoir was drawing down while measured output held steady, or that the system recovered only after an external material replacement. These outcomes tell us what the receiver could build with its current instructions and what the next instruction must address.

What would cross the distance

An actual interstellar demonstration would start with an instrument that can return an interpretable signal from a new environment. It would follow with a tested receiver whose reported operations can be compared with the announced design. A recorded sequence and a nonliving verification object would give early opportunities to test data transfer and local manufacturing separately. A complete archive might take years to send at the rates we have considered, so there is value in making early packages small and complete.

Imagine one reference receiver commissioned at year 10,000 at a fixed distance of 4.25 light-years. An immediate report reaches Earth at year 10,004.25. A reply sent at once reaches the receiver at year 10,008.50. Confirmation reaches Earth at year 10,012.75. The fractional dates describe only signal travel. Braking, deployment, commissioning, data transmission, and observation add their own time. The receiver carries its initial instructions and has to be able to conserve its capabilities if no reply comes. Any later Earth institution that chooses to reply will need enough records to understand what it is being asked.

At the other destinations those dates spread over tens of thousands of years. An early report may reach a surviving fleet archive long before any long-shot craft can return a measurement. A probe already on a different course can use an authenticated update if its geometry and listening schedule allow it. Its baseline authority and ability to work have to be present at launch. The timetable separates the dates of local success, reception elsewhere, and confirmation by an independent site. Those events can fall in different eras.

This is one reason to design the Earth experiment now in a form an independent team can inherit. The account should preserve the build, the tests, and the justification for acting, rather than assuming the launch team will still be there to explain them. A long project needs a history that can be questioned. The result will deserve trust if someone else can reproduce it or locate the uncertainty that prevents them from doing so.

A result worth the waiting

The original question asked what it takes to make one living cell begin to divide. In the interstellar case the first genuinely new result might be modest. A remote machine receives an edition of a small package and returns evidence that a declared material circuit now closes within stated limits. A later package gives a system a chance to maintain more of its own organization. It might be a contained cell test. Eventually the experiments could show a larger system renewing the air or nutrient stocks the next interval requires. Each result specifies what was added to the receiver and what the receiving environment supplied.

Life can travel as matter already organized to persist, and a mission may find cases where that is the best way to carry it. The alternative in this book gives us another claim to test. Given the right receiver and environment, an instruction might carry enough of life's starting organization to let matter at a different star do the work. A useful test of that claim is a comparison with living cargo under a specified radiation exposure and a specified time. It puts the origin of our question back into the design of the experiment.

We can imagine the report from the other star with unusual specificity. Its clock would show the day the last assembly command ran. Its records would show what was used and what was returned to the cycle. The enclosed system would be at the beginning of an interval we can compare with the interval before it. A later report would arrive with more of the same measurements. The report from a second workshop might arrive centuries later. We could compare what each machine began with, which processes the descendants renewed, and which material stocks the enclosures could put back into service. A fleet earns its value when two histories can tell us something that either history by itself would leave open.

The tests that follow give the receiver and its observers a more specific job. The first begins with a count of the equipment a founder might inherit. The last asks whether a later observer can recover an experimental result after the people who planned the flight are gone.

28 What Counts as the First Cell

A result a second observer can recognize

An image arrives from the receiver. It shows a small compartment narrowing in the middle. In the next frame there are two compartments where there was one. The image is exactly what the builders hoped to see. The workshop sends something else with it. There is a record of what entered the chamber, a trace of the events before the image, and measurements of what the two compartments did afterward. One of the two grew. The other remained visible but did nothing we could measure. We have a promising event and the beginning of an explanation.

Now imagine the same report coming from a second workshop. Its images are sharper and its descendants more numerous. The two reports would be hard to compare if one team counted each visible compartment as an offspring and the other required evidence that a daughter could make a daughter of its own. Before the first run, the teams need to agree on the unit they are trying to build and on the observations that distinguish a build from an object that merely resembles one. They can use different instruments. Their instruments must answer the same questions.

The first question concerns ancestry. The record needs a founder, a time at which it passed into the conditions defined for handoff, and a way to connect the later cells with that founder. Where observation becomes discontinuous, the record should show the gap. A sequence found in the chamber after a gap can identify material that was there. It has a weaker claim to identity with a particular object seen before the gap. Track the ancestry of failed branches too. If one founder produces a thriving line while ninety-nine never divide, the positive result describes a possibility under the tested conditions. The other ninety-nine describe a problem the receiver has to solve before it can count on the possibility.

What could still be left from the founder

Suppose the receiver provides a founder with working molecular assemblies. We want to know how many later assemblies can be explained by carrying the original ones along. For a hypothetical count, set a conservative upper limit of 10,000 complete copies at handoff, including any intact copies that could become active later. Follow the founder through six rounds of division. In this deliberately simple

example, all its branches survive and there are sixty-four descendants. If a validated measurement puts a lower bound of 400 active copies in each descendant, there are at least 25,600 copies in that cohort. At least 15,600 working assemblies must have appeared after handoff, provided none entered from outside. Unequal partitioning can change which cell gets a founder copy. It cannot add a copy to the sum.

This arithmetic asks for a count of assemblies at one checkpoint. Adding the measurements from successive generations would count some objects more than once. Counting parts that are free in the cell would substitute an easier measurement for the one that matters. If a particular assembly turns over quickly, the number of assemblies made between checkpoints can exceed the net increase by a great deal. The estimate from the stock change is a lower bound. The assay and its error determine how useful that bound is.

A second measurement can ask whether newly made parts are in those assemblies. In ordinary bacterial cells, Stephen Chen and colleagues used stable isotope labeling and mass spectrometry to measure the incorporation of newly synthesized proteins into assembled ribosomes (Chen et al. 2012). That work supplies a precedent for the measurement. In a founder test, the same general distinction would help separate old parts from parts made after handoff. A label detected somewhere in the descendants gives evidence of uptake. A label in a newly made component of a functional assembly gives evidence about the contribution to renewal. The record would still need to allow for reuse of older parts and for any labeled material admitted in a form the lineage could use directly.

The simplest successful result would name one important machine and show how it was renewed. A broader claim about the working cell asks for several linked demonstrations. The descendants must keep the necessary genetic material together through division, maintain their boundaries, and obtain enough matter and energy to build again. Different founders may fail at different points. A useful test lets the failure keep its location.

How a plausible intruder gets into the account

A successful pedigree could have started in the wrong place. If a trace of an existing culture enters the experiment, it may be better at growing than anything the apparatus has just assembled. A digital checksum can show which file the receiver opened. A genetic match can show that the cells carried the expected sequence. If an existing organism with that sequence had access to the apparatus, the match would be compatible with either history.

The provenance record has to follow the physical materials. Give each input lot its own identity and retain evidence about its source. Associate tools, chambers, and instruments with the runs in which they were used. Keep the living reference in separate custody and use it to test the ability to observe the relevant outcome in an independent setting. A negative observation in a blank has more meaning when a separate check shows that the assay would have found a small amount of the kind of material under suspicion. No one instrument has to do every job. Observation of the living history and measurement of the sequence are stronger together when they can fail separately.

It helps to put a number on the question asked by the blanks. Suppose the chance of a detectable false lineage were one percent in each independent run. The chance of at least one false lineage in forty-eight trials would then be about 38 percent. In twenty-four matching blank runs, the chance of seeing none would still be about 79 percent. Both figures follow from the same model. This is an illustration of what a limited series of clean blanks can tell us. It supplies no estimate of the actual contamination rate. Shared tools or input lots can make the outcomes depend on one another, and a blank with a different physical history may never encounter the source that affects the build.

One receiving team can test whether removing a needed component blocks renewal and restoring it recovers the outcome. A second team can work from separately sourced components while carrying out the same comparison. A third group can review the raw observations under labels that conceal which run was expected to succeed. Each measure has a particular job. A restored run can identify a dependency. A separate supply path can challenge a shared origin for an apparent success. Independent review can expose an observation that becomes persuasive only when the observer knows what it is supposed to show.

Keep the failed claim in the record

Suppose a workshop receives forty-eight build opportunities and four produce lineages that reach the endpoint. All four came from one input lot. The first report can list the four passes, their ancestry evidence, and their measurements of renewal. The second report can ask why that lot mattered. One possibility is that it provided an essential material in a usable state. Another is that it carried a contaminant. Both explanations give us a reason to retain the lot history and repeat the comparison with a physically independent supply.

A trial in which no observed founder divides has also delivered an answer. If the membrane formed and the first protein appeared, the receiver can preserve those facts even if the later stages failed. A chamber disconnected for an instrument fault belongs in the record with the time and reason it was disconnected. The distinction between a completed failure and a run that ended before the outcome became observable will matter to the next designer. Changing the endpoint after a run can be a reasonable way to propose the next experiment. It also changes the question that experiment will answer.

When a report has to cross four light-years, much of this discipline needs to be designed into the apparatus. The record should preserve a small set of original observations before a later model condenses them. A summary can carry the declared endpoint, the status of the controls, and the range of results. A larger follow-up can carry image sequences and calibration records for a reader who wants to test a different explanation. If only the first report gets through, it should still give the evidence for the claimed stage and the reason the receiver counted the other stages as open.

What crossed the distance was a set of instructions and a machine with a known inheritance of its own. What the first-cell experiment adds is a line of descendants whose history starts with the declared build and whose continued work can be measured. The two histories touch at handoff. A good report lets another observer trace them to that point.

29 Building the Next Living Layer

The workshop after the first lineage

The first lineage has passed its test. The receiver has watched descendants rebuild their machinery and has sent the record home. In the next chamber it has been asked to grow something that can feed another population. It has a record of a plant, including the sequence of its chromosomes. It still has to decide what living material begins the build and what makes the passage from a working cell to a plant that can yield another crop.

A bacterium that renews itself can become a useful member of a community. Its success settles neither the construction of the next member nor the exchange between them. A first task is to list the jobs the larger system must perform, then look at how many kinds of material and machinery each job needs. The workshop may discover that some of the jobs can be done by a reactor and some are better done by an organism. That is a decision about the habitat. It need not be made in advance by the list of organisms in the archive.

A cell that can start a plant

A leaf cell has several kinds of inheritance to account for. Its chloroplasts develop from earlier plastids, which carry their own genomes and receive many of their proteins from the rest of the cell. Keeping plastids available as a tissue grows involves their own division machinery and contributions encoded in the nucleus (Chen et al. 2018). The assignment to make a leaf therefore hides a step before we can ask whether the leaf fixes carbon. A receiver needs a starting condition that can yield the organelles the developing tissue will use. The nuclear sequence helps specify what many of their parts are. A demonstration has to show where the first working arrangement came from.

A plant also has to put its parts in places where they can do different things. A root and a leaf are made by descendants of the same founder. One makes contact with a water supply and the other with light and air. A seed holds material with which a plant can begin development. A cutting brings existing tissue with it. Either can be a useful starting point in an Earth experiment. Each makes a contribution to the work and should be named in the report. A machine that starts a plant from a viable cutting has passed a propagation test. A machine that starts one from a digital record and a declared nonliving inventory has a different test to pass.

The archive can help by describing what would make the two results comparable later. It can specify the growing interval, what goes into the seed or cutting stock, and which observations count as functional growth. If the eventual crop must provide the next crop's starting material, its output includes planting stock with a measured ability to grow. If the first test only needs a photosynthetic producer, a small, contained producer with a simpler life history may let the workshop demonstrate the relevant exchange first. The smaller test earns credit for the exchange it measures.

Begin with the jobs that can be separated

Suppose the next demonstration calls for a producer, a consumer load, and a material return. At the beginning the consumer can be an instrumented process representing a stated load. It takes a specified product stream and returns a known material stream to the experiment. We can ask whether the habitat can return part of that stream to a form the producer can use. No participant has to live on the output while we are finding the size of the loss.

Run the producer first with inputs from a known supply. This gives the investigators a comparison for what the producer can do under the stated light, timing, and environmental conditions. Run the return stage separately with a declared feed. Check its output for composition and for its ability to support the next growing interval. The chemical measurements can tell us that an element returned. The plant response can tell us whether the returned form and the rest of the mixture served their purpose. A stream that passes one comparison can still require further work.

A biological return stage may use organisms that transform residues into useful compounds. A chemical stage may make some of the same elements available using different machinery. Compare the two against a common output and a common time limit. Record the energy, replacement materials, and stock they leave behind. The biological branch needs a means of renewing the population responsible for the transformation. The chemical branch needs a means of restoring reagents and replacing parts. Either path can work at one scale while depending on inputs the workshop cannot make at a larger one.

Then connect the stages and allow the transfer itself to be tested. In a control enclosure, keep the producer on the known supply and run the return stage in parallel. In the coupled enclosure, let measured recovered material displace a predeclared share of the known supply. Keep the growing schedule comparable. A decline in both enclosures points to a condition they share. A decline confined to

the coupled run makes the recovered stream and its timing worth investigating. If adding back a known supply restores the outcome, record the amount and the time of the addition. The rescue can identify a bottleneck even when the run no longer qualifies at its original limit on imports.

Where the missing grams went

Take an illustrative phosphorus account in one unoccupied enclosure. Begin an interval with 100 grams in a plant-accessible reservoir. Transfer 60 grams to the growing system. At the same checkpoint next cycle, recover 52 grams in the accessible form, six grams in a measured solid residue, and two grams in a spent filter for which the apparatus has no recovery path. Assume that all other stocks are measured at comparable stages and do not change in this example. The total is accounted for. The available reservoir has fallen to 92 grams.

Now give the workshop a secondary process that restores all six grams from the solid residue before the next growing interval. The accessible stock reaches 98 grams, at a measured cost in energy and time. Suppose the initial acceptance plan called for at least 97 grams at every such checkpoint. The first cycle meets that criterion. If the same transfers repeat in the next cycle, the reservoir falls to 96 grams. The first crop, the second crop, and even the complete elemental account could each look reassuring. The downward change in accessible stock still matters to the next run.

The measurement needs resolution sufficient to distinguish that trend from sampling error and drift. A calibration sample can establish what the measuring instrument reads today. A second measurement path can help test whether it reads the same material the same way at the next checkpoint. The remaining stock also needs to be usable when the plant next takes it up. Recovering all six grams after planting starts can make a different crop even when the eventual mass balance looks the same. A workshop can report a successful delayed recovery and a failed planting window in the same record.

A temporary transfer can also deprive another member of the community. A producer that captures a nutrient rapidly may draw down a pool another producer needs earlier in its own cycle. The total can be ample while the schedule is wrong. Take the measurements at the locations where each part uses the material and over intervals short enough to see the deficit. Where one shared transfer serves several populations, a test that interrupts it can show whether either population has time to wait. The desired capability is a material path and a working schedule.

What the two versions of the test establish

An Earth trial can use organisms from a trusted stock and begin with a workable growing system. Its contribution is to test the transfers, show the effects of disturbances, and establish how many growing intervals the combined system can sustain under specified inputs. A receiver beginning with no viable material has a further obligation. Before combining the stages, it must demonstrate an origin for every population whose living work the design uses. If it can construct one producer and must import a viable partner, the report should show exactly which contribution came from each path. The mixed experiment remains informative. It describes a different starting inventory.

There are also several ways to finish. A small contained ecology can establish measured renewal under declared power, material, and maintenance inputs. A food production test can add a defined requirement for usable composition and continuity of supply. The provision of food is only one requirement for a community of people. Gestation, health, and the work of raising children introduce further obligations, with independent evidence and authority needed for each. The receiver can qualify a producer or a recycling loop without treating its success as permission to cross those boundaries.

The contained version has to remain safe while the investigators are learning how the parts behave together. Inspect the boundary and verify the path for spent material before the living trial. Decide how the system will respond to lost power or a failed transfer. A successful repair inside the enclosure can become one of the results. A compromised boundary changes the terms of the experiment and ends the biological trial. Work on the apparatus and interpretation of the record can still continue.

The next delivery from another star may show a producer restoring a usable stock and a recipient growing on the returned material. It may show the stock declining slowly, with the missing fraction found in a place the workshop can reach only later. Either result can be compared with the first experiment on Earth. The time to build a second living layer begins to have a meaning when we can say what it inherited, what it renews, and how long it can wait for the next exchange.

30 One Ship Through One System

The first craft in the fleet needs a whole journey on paper. A working transmitter is of little help if the ship never turns into an orbit. An ample power budget at the working site does little for a controller that fails four months before the site comes into reach. This case follows one ten-tonne craft to Alpha Centauri A. Its mass and energy accounts fit together under stated assumptions. Several of those assumptions are requirements a real design would have to meet. If one fails, the destination, the load or the mission itself must change.

Departure and the shape of the crossing

Use 4.35 light-years for the initial distance and 127.4 kilometers per second for the unbraked speed relative to the Alpha Centauri system. The simplest intercept gives roughly 10,240 years. The motion of the system changes the aim and the energy at departure. Its barycenter approaches the Sun at about 22 kilometers per second and moves transversely at about 23 kilometers per second. In a straight-line intercept at the chosen relative speed, add that stellar velocity vector to the craft's closing velocity. The departure speed in the heliocentric frame is then about 108 kilometers per second, corresponding to about 5.8×10^{13} joules of kinetic energy for ten tonnes, or 16 gigawatt-hours. The corresponding speed relative to the star sets the braking problem. The actual intercept must include the orbit of A about the barycenter and retarget during approach (Kervella et al. 2017).

The first line of this account begins after an external launcher has brought the craft to cruise. Sixteen gigawatt-hours is energy in the departing mass. A beamed or other external launcher would have to deliver the required momentum and absorb the costs of coupling, pointing and inefficiency. Neither the launcher nor its mass is included aboard. At cruise the main workshop, thrusters and communications terminal are dormant. The retained navigation system must keep an independently testable estimate of the target and of its own state. A one-milliwatt average allocation still corresponds to roughly 90 kilowatt-hours over the crossing. One cold-start design would hold most of that work until the final approach, after which power grows with proximity. Any in-flight wakeups need a source and a separate qualification.

The cargo volume for the comparison in the next chapter is bounded by an inner vault. A spherical shell with an inner-scale radius of 0.5 meter and a nominal areal density of 100 kilograms per square meter takes about 314 kilograms of shielding material in the thin-shell approximation. Penetrations, seams, an outer dust shield and the route-specific radiation spectrum all change the protection. The budget keeps the same vault and exposure across the three cargoes. A shield thickness alone supplies no survival probability.

Cruise allocation	Mass in kilograms	Items counted
Sail system	1,500	Reflector, deployment, orientation and reefing
Electric propellant	2,900	Working propellant, including transfer reserve
Power and propulsion hardware	1,400	Tanks, thrusters, first collector, main array, storage, radiators and converters
Structure and protection	1,000	Vault, primary structure and external impact protection
Avionics and communications	650	Navigation, independent restart chain, archive and optical terminal
Science and containment	1,850	Common test enclosure and the starting cargo
Unassigned mass	700	Subsystem growth and provisions not yet priced
Total at cruise	10,000	Mass used in the capture calculation

Wake before the brake

For Alpha Centauri A use 1.1055 solar masses and 1.521 solar luminosities. At ten astronomical units its irradiance is about 20.7 watts per square meter. A two-square-meter fixed section could yield about seven watts to trigger a low-power controller. With that section and the deployed panels totaling 30 square meters, 25 percent conversion and a 0.70 loss allowance give about 109 watts. A 50-watt survival and pointing bus then leaves roughly 59 watts for startup functions and charging at that distance. A separate 24 kilowatt-hour usable store would carry that bus for 20 days without light. Delivering those usable watt-hours after the crossing, finding a usable view around the sail and deploying the first collector are explicit requirements of the design (Salmon et al. 2021).

Ten astronomical units is the start of active braking in this example. The approach from there to 0.25 astronomical unit lasts about 138 days in the simplified trajectory

below. The startup sequence therefore has a practical order. First establish the time, attitude and power needed to use the sail. Then recover the approach plan and compare it with the observed star positions. The sail has to become useful near the planned start radius. It cannot wait until the workshop is awake and ready to run a laboratory.

Remove enough orbital energy to stay

The reference brake is a photon sail that acts on the inward leg and is feathered at the end of it. The star's radiation presses outward while the craft moves inward. If the sail kept applying the same inverse-square force on the outward leg, the craft would regain the work the sail had taken away. The cutoff matters as much as the area. At 0.25 astronomical unit, aim to leave the craft at the periapsis of a 0.25 by 1.0 astronomical unit ellipse around A, with a speed of about 79 kilometers per second. The same craft approaching with unaltered positive orbital energy would have about 155 kilometers per second there. Solar-electric propulsion enters the plan only after the sail has removed enough energy for a bound orbit.

The incoming specific orbital energy is half the square of 127.4 kilometers per second. The desired specific orbital energy is minus the star's gravitational parameter divided by twice the ellipse's semimajor axis, which is 0.625 astronomical unit. The difference is about 8.90×10^{9} joules per kilogram, or 8.90×10^{13} joules for the craft. An ideal face-on reflector receives a pressure equal to twice the incident energy flux divided by the speed of light. Integrating that outward force along the change in radius from 10 to 0.25 astronomical unit gives about 8.06 megajoules per square meter of reflecting surface. An ideal 16-square-kilometer reflector would supply about 1.29×10^{14} joules of work along that leg. The design therefore needs about 69 percent of that ideal radial work.

For the bookkeeping model, represent the sail force by 69 percent of the ideal radial pressure at every radius. Use conservation of angular momentum in this central-force model and select the incoming angular momentum for the desired cutoff speed. Integrating the resulting radial speed gives the 138 days. A real attitude law changes the direction and magnitude of the force and may change the orbit plane. A useful design must demonstrate the required work, attitude authority and cutoff together on a trajectory that includes the companion star. The suggested area is equivalent to a square four kilometers on a side. The allocated 1,500 kilograms implies about 0.094 gram per square meter for the whole deployed system, including support and control. Reefed sail mass stays aboard during the subsequent

transfer. Heller and Hippke's study shows why the photogravitational problem is worth calculating. Its much lighter reference sail and payload do not qualify this deployment (Heller and Hippke 2017).

The thermal gate is separate. Near cutoff, incident irradiance is about 33 kilowatts per square meter. Absorbing even one percent would put 330 watts of heat into each illuminated square meter of sail. The temperature then depends on infrared emissivity on both faces and how the sail and support structure see one another. The spacecraft and collector have their own exposure and shade requirements. After cutoff, residual sail forces have to stay low enough to preserve the bound orbit. Braking starts only with a tested recovery plan for a slow deployment or a low measured pressure. Where capture is no longer recoverable, the science changes to a limited flyby survey and the culture enclosure remains closed.

Report the flyby as partial evidence. The fleet’s survey milestone in Chapter 32 requires a sustained operating orbit, so a flyby cannot qualify for that count.

Earn the working orbit and the power to use it

At the 1.0 astronomical unit apocenter, the capture ellipse gives about 19.8 kilometers per second. A circular orbit there requires about 31.3. The difference, 11.5 kilometers per second, is the ideal impulsive benchmark for a transfer concentrated near apocenter. At 40 kilometers per second effective exhaust speed, the rocket equation gives about 2,500 kilograms of propellant from a ten-tonne start. The 2,900-kilogram load provides about 13.7 kilometers per second of ideal total capability. The remaining roughly 400 kilograms after the benchmark transfer is held for finite-burn penalties, navigation corrections and other assigned maneuvers. The companion's phase and orbit stability have to be checked for the actual site (Akeson et al. 2021).

A 110-square-meter array at one astronomical unit provides about 39.8 kilowatts at the same electrical conversion and allowance used for startup. Allocate 32 kilowatts to electric propulsion and five to the bus. At 65 percent conversion to exhaust kinetic power and 40 kilometers per second exhaust speed, thrust is about one newton. Spending the 2,500 kilograms takes about 1,110 powered days, so the transfer is a multiyear stage. A 30 to 50 percent thrust duty would make this roughly six to ten calendar years if the propellant requirement stayed near the benchmark. A finite-thrust trajectory has to check both assumptions. Power available during any closer pass is limited by panel articulation and thermal design.

The propulsion case leaves about 11.2 kilowatts of electrical input outside the exhaust kinetic power. Counting that conservatively as heat aboard, plus the five-kilowatt bus, gives about 16.2 kilowatts to reject. A 55-square-meter radiator at 330 kelvins, emissivity 0.85 and a 0.75 view and utilization allowance gives about 23.6 kilowatts to cold space. That allowance presumes it can be shaded and has an open view. In the working orbit the budget can support a 15-kilowatt laboratory load alongside the five-kilowatt bus within that heat ceiling. Sunlight absorbed by panels outside the bus, transient loads, loops at different temperatures and any warm view of the sail need their own accounting. The quantities above give the thermal designer a ceiling to verify.

The first laboratory work uses a declared kit aboard. Local resources may become a measured option after survey, with assays of composition, recovery rate and energy cost. The initial cell demonstration does not depend on a successful mining expedition. It depends on verified containment, a functioning blank and reference channel, and enough energy and inventory to follow the lineage to the predeclared endpoint.

Return a record that can be read

At a conservative design range of 4.35 light-years, assign 20 watts of emitted optical power at 1,550 nanometers, a 0.5-meter transmitting aperture and a 30-meter effective receiving aperture. The ideal aperture factor is the square of π times the product of the diameters divided by four times the wavelength and range. This gives about 5.3 received photons per second before losses. A 15 percent efficiency allowance after the ideal aperture factor gives about 0.8 detected signal photons per second. Background photons, mispointing and detector behavior must be measured separately. NASA's Deep Space Optical Communications experiment demonstrated beacon acquisition and photon-counting reception at solar-system ranges (NASA Jet Propulsion Laboratory n.d.).

Assign 20 detected signal photons per delivered information bit and a usable observing duty of 25 percent. These are design targets for link testing. The resulting average rate is about 0.01 bit per second. A 10,000-bit structured status report would take about 12 days to send, followed by the one-way light time. A one-megabit record takes about 3.2 years to transmit at the same average rate. The small report can include the clock, orbit, kit identity, test result and pointers to the retained evidence. The larger record carries the observations that let Earth assess the

result. A cryptographic digest helps verify that a received record matches the one the craft retained. It cannot stand in for the measurements.

At one astronomical unit from the star, the maximum projected separation from it is less than one arcsecond at the design range. The transmitter must acquire the future receiver position across light time while the Earth receiver resolves a weak signal near a bright star. The 30-meter equivalent aperture and an institution that still listens are resources assigned to Earth. Onboard data survive loss of a scheduled contact. The lowest-rate beacon should be attempted from the first useful orbit, without delaying capture.

A rough clock can carry these allowances together. Begin with about 10,240 years for the unbraked reference crossing, allow days to weeks for the approach correction in this simplified case, then add roughly six to ten years for transfer. At the working orbit a 60-day trial, about twelve calendar days to send the first status record at the assumed average rate, and up to about 4.35 years of signal travel put the first post-trial report roughly ten to fifteen years beyond the nominal crossing date. The resolution of that estimate is years. A capture beacon could arrive earlier and a complete evidence package later. Neither is a scheduled date until the trajectory and link have been qualified.

This is a case a fleet designer could reject by test. The launcher, the long-life restart chain, the sail and the return link all have measurable gates. The mission clock includes the crossing, the transfer, the test and the time to receive a result. The same ledger also provides the fixed transport case for asking whether the first cell should cross in a dormant seed, be built after arrival, or draw on both approaches.

31 Three Cargoes on the Same Route

Put three versions of the receiver on the route just described. Each has the same ten-tonne cruise mass, shield geometry, arrival orbit, electrical ceiling and communications link. In one the live cargo is dormant. In another no viable organism or cell-derived working component is admitted to the starting inventory. The third carries dormant material and a limited set of construction tools with separately protected records. The question for all three is whether a contained lineage can renew itself at the destination. A second question asks where the organization needed to start that lineage came from.

Hold the endpoint and the boundary fixed

Give the three receivers the same predefined target phenotype in a benign reference lineage and the same acceptance test. A receiver must follow a documented founder and its descendants for at least 30 days and through at least 20 measured net doublings after handoff. During that interval the chamber may supply qualified simple nutrients, water, energy and environmental control. It may neither introduce viable material nor replace the machinery under test. A second line of evidence measures renewal of a designated cell component from post-handoff inputs. Twenty doublings amount to a millionfold population increase in the ideal case. That multiplication makes an initial stock a poor explanation for the bulk of the final population, while a label or equivalent independent measurement helps establish which components were actually renewed.

A blank channel follows the same handling history. A reserved reference instrument checks that the chamber's observations would distinguish the intended lineage from a contaminant or an inherited fragment. Predeclare which failures invalidate a run and which make it a clean failure to establish a lineage. A failed blank invalidates the evidence even if the main culture grew. For all three receivers, store the measurements behind the condensed message sent to Earth.

The acceptance test begins after the working orbit and shared laboratory pass commissioning. Allow up to 60 calendar days for the first test attempt including startup, lineage observation and verification. At the 15-kilowatt laboratory ceiling, 60 days would use 21.6 megawatt-hours, plus the shared bus. A single 25-kilogram allowance of qualified test feedstock and water aboard each craft fixes the same boundary for the initial comparison. These are ceilings, not estimates of the biological load. Each team has to show that the apparatus and stocks can support its

protocol within them. An early result with a simpler apparatus counts if it reaches the same endpoint.

Spend the same mass three ways

Of the 1,850 kilograms in the science allocation, reserve 1,000 for common containment, measuring equipment, the qualified test stock and science support hardware. The other 850 kilograms may be allocated as shown below. These are design allowances. An allocation that cannot fit qualified equipment of the required capability has failed before the voyage.

Strategy	850-kilogram allocation	First route to the shared endpoint
Dormant seed	300 kg for packaging and isolated lots, 400 kg for recovery equipment and consumables, 150 kg for spares	Recover and test an Earth-origin founder
Abiotic receiver	550 kg for synthesis and assembly hardware, 200 kg for precursors and consumables, 100 kg for spares	Assemble a founder from a declared nonliving inventory
Hybrid	200 kg for protected dormant lots and their immediate supports, 430 kg for recovery and limited assembly equipment, 220 kg for consumables and spares	Attempt recovery and test the separately declared assembly path

The dormant receiver spends more of its cargo budget on protecting organization already assembled on Earth. The relevant measurement is usable founders after the route history, followed by pedigree and renewal. Spores recovered from a six-year low Earth orbit exposure show why dormancy deserves consideration. Ultraviolet protection and packaging mattered in that experiment. A millennial voyage with a different spectrum and thermal history still requires its own survival model and experiments (Horneck, Bücker, and Reitz 1994).

The abiotic receiver spends more on equipment and consumes part of its test allowance bringing that equipment into verified working order. The first result must show both that a lineage meets the shared endpoint and that the inventory met the stricter provenance requirement. A chamber that reaches the endpoint using a preassembled, biologically sourced catalyst can be a step in a ground development program. Its manifest places it outside this particular cargo class. At equal launch mass the pivotal question is whether a machine small enough for the allowance can supply the needed organization and retain that ability after storage.

The hybrid carries two ways to begin work, though the second path is deliberately limited by mass. Its small assembly bench may be able to provide a missing component for recovery before it can construct a founder from simple materials. A repair that rescues an Earth-origin lineage counts under the shared endpoint, with its provenance recorded. A separate abiotic success requires a separate inventory, chamber and run. If the two paths share a controller, a damaged archive or a seal process, count that dependency once at the level where it can disable both. Spreading the budget has value when it buys a genuinely different path through the failures that matter.

Compare damaged kits as well as fresh ones

Test the three kits on Earth under the same assay contract. Use fresh controls, independent packages and blinded status where possible. Give representative vault locations the radiation and thermal histories a transport model predicts, and measure the history actually delivered. Test component-specific alternatives when compressing exposure changes the expected damage. An accelerated radiation exposure, a long cold soak and a test of restart after deliberately damaged records answer different parts of the question. Their results belong in separate columns until a physical model supports a joint extrapolation.

Replication within one lot does little about a defect in the lot. In an illustrative model, eight independent packages each have a 12 percent chance of giving a qualified founder if they escape a shared failure. Their chance of yielding at least one is about 64 percent conditional on that gate. A 30 percent probability of shared failure brings the unconditional chance down to about 45 percent. The shared cause could be a sterilization error, a flaw in the common enclosure or a vulnerability to the same radiation field. The example is a reason to assign independence carefully. Its numbers supply no estimate for a flight.

Suppose a particular Earth test produces zero qualified lineages in 60 independent trials. Under a fixed per-trial probability model, the one-sided 95 percent upper bound is about 4.9 percent, obtained by solving the probability of sixty failures for a value of five percent. That conclusion describes the conditions and independence of those trials. If the 60 packages came from three common manufacturing lots, the trial record should carry the lot identities and model the shared risks separately. Even a well-measured result may leave the destination question open if the imposed stress is poorly related to the route.

Choose the cargo for the question

The first decision is whether any path passes the transport and laboratory gates. If the main question is the earliest qualified living lineage, select among qualifying paths using measured robustness, reserves and time to acceptance. For a test of reconstruction from abiotic starting material, the extra provenance condition determines which result answers the question. A survey-only craft can still return a useful record if the life experiment remains unqualified. Its status belongs in the fleet account as a survey result.

A comparison at Alpha Centauri A would answer a modest but specific question. It would show how much of one line of life can cross embodied in a dormant organism, how much can be restored through local work and what inputs an independent builder needs. The same table could be run on Earth long before the first launch. If the table will not close on Earth, waiting ten thousand years adds no evidence that it will close around another star.

32 Choosing What the Fleet Must Prove

The fleet can be said to have succeeded in several incompatible ways. One ship might return a good map of a system in which it could never open its biological chamber. Another might build a verified microbial lineage and find that it has no way to close the larger nutrient cycle. A third might produce useful crops while drawing down a reserve that would eventually be exhausted. These results should remain distinguishable after the fleet has been reduced to a page of numbers.

Define the result before counting ships

Use three reports. The survey report documents a sustained operating orbit, calibrated observations, and enough evidence about materials, power, and hazards to judge at least one proposed worksite. A finding that the site is unsuitable can satisfy this first criterion. The lineage report adds a verified renewing lineage in containment, with a declared starting inventory and observed renewal after the apparatus withdraws organized help. The habitat report adds a contained growing system that meets its reserve floors and replenishes its usable stocks through the declared recovery test. An unoccupied enclosure can meet that requirement. The fleet carries no baseline promise to establish or support a human population.

The lineage endpoint accepts any cargo route that meets the common test in Chapter 31. Each report gives the founder’s provenance separately. A recovered dormant founder can satisfy the lineage criterion. A strict abiotic result must also meet the stricter inventory rule. The model below gives no probability for that more specific result. A habitat report likewise records the provenance of each required population, including any viable producer, partner, or recycler brought aboard. Appendix E gives examples of both decisions.

For the calculation that follows, require each report to include the evidence for every earlier report in the sequence. This convention keeps the outcomes nested. A different program could build a crop chamber from a separately qualified seed stock while its first-cell experiment is still unresolved. Its results would require a branching account of the evidence. In either version, draw the boundary around delivery to Earth if that is where the result will be judged. A workshop might perform a successful trial and lose its transmitter before sending the observations.

The event belongs in the destination record and the delivery failure belongs in the acceptance record.

Keep the eleven reference craft. Send three each to Alpha Centauri, Barnard’s Star, and Epsilon Eridani. Send one each to 61 Virginis and TRAPPIST-1. Equip every craft to survey. At each nearby destination two craft can also attempt to establish a verified renewing lineage in containment, and one of those two carries the additional capabilities for the renewal trial. The 61 Virginis craft can attempt the same lineage test after its survey. The TRAPPIST-1 craft surveys. In this accounting, an experimental package makes its craft eligible to attempt a milestone. The orbit, equipment, and containment must still qualify on arrival.

Report	Alpha Centauri	Barnard’s Star	Epsilon Eridani	61 Virginis	TRAPPIST-1
Survey	3	3	3	1	1
Verified lineage	2	2	2	1	0
Renewing habitat	1	1	1	0	0

Treat the Alpha Centauri assignments as a single branch for the shared-site term in this illustration. Real craft in that system would have distinct approaches and worksites. The branch should eventually be split wherever those choices remove a shared constraint. Two craft can see the same star and still face different dust, maneuver, or local material conditions. They can also go to different stars and carry the same defective sensor.

A calculation with three outcomes

For each kind of report, give the model three probabilities. A fleet-wide factor g covers success at all the common requirements for that report. A site factor q covers the conditions shared by craft in one destination branch, conditional on clearing the fleet-wide requirements. A craft factor p covers the remaining chance of delivering the report once those common requirements and local conditions are met. The lineage and habitat factors are cumulative to their named endpoint. Multiplying a cumulative lineage factor by a survey factor again would charge the same work twice.

For a given endpoint, suppose that the remaining craft failures are independent within a branch, and that the branch failures are independent once the shared fleet

requirements have been met. At a site with n craft, the branch success chance is $s = q \times [1 - (1 - p)^n]$. For the whole fleet, multiply the branch failure chances and subtract the product from one, then multiply by g. This is the calculation of Chapter 3 applied to three distinct outcomes. Each assumption can be replaced separately.

Report	Shared fleet factor g	Nearby site q	Nearby craft p	Long-shot site q	Long-shot craft p
Survey	0.90	0.80	0.7000	0.65	0.600
Verified lineage	0.75	0.55	0.3850	0.30	0.300
Renewing habitat	0.60	0.35	0.1925	—	—

The values are illustrative assumptions for all five systems. For a nearby craft, p = 0.385 for the lineage report can be read as 0.70 for the craft-specific survey work followed by 0.55 for the further craft-specific work. The habitat row multiplies 0.385 by a further 0.50. The shared and site factors also decline as requirements accumulate. In one construction of the model, each factor can be generated by a sequence of conditional pass chances. For example, the shared habitat requirement passes in 0.60 ÷ 0.75 = 80 percent of the states that clear the shared lineage requirement. That interpretation keeps a failure at an earlier stage from being reversed by a favorable draw for a later stage.

In the lineage row the nearby branch has two eligible craft and $s = 0.55 \times [1 - (1 - 0.385)^2] \approx 0.3420$. The eligible long shot at 61 Virginis has $s = 0.30 \times 0.30 = 0.09$. The eventual fleet chance of a lineage report is about $0.75 \times [1 - (1 - 0.3420)^3 \times (1 - 0.09)] = 0.5555$. Applying the same calculation to the other rows gives the table below. The decimals expose the arithmetic of the stipulated inputs.

Report	After all assigned attempts	Only three nearby branches eligible
Survey	89.64%	89.02%
Verified lineage	55.55%	53.63%
Renewing habitat	11.33%	11.33%

The left column includes five survey branches, four lineage branches, and three habitat branches. Its 55.55 percent is the illustrative chance of at least one accepted lineage report across its assigned routes. Compare a strict abiotic subset only after declaring its own eligibility and its dependence on shared equipment and inventories.

The right column imagines a cutoff by which all nearby branches have had their full assigned opportunity and the long shots have had none. For a real cutoff, each craft must have time to reach its assigned worksite, make the observations, finish the trial, and transmit enough evidence for acceptance. A habitat test has a clock that a survey does not. The light time of the return path then belongs on both clocks.

Spend the next ship on a named purpose

Hold the fleet at eleven and replace the TRAPPIST-1 survey craft with one equipped for all three reports at a nearby destination. The illustrative lineup changes from three groups of three and two singletons to nearby groups of four, three, and three with one distant singleton. The changed nearby branch now has four survey craft, three lineage craft, and two habitat craft. After all assigned attempts have had their opportunity, the survey chance falls from 89.64 to 89.44 percent. The lineage chance rises from 55.55 to 57.92 percent and the habitat chance rises from 11.33 to 14.17 percent. Under the early cutoff, the additional nearby craft raises the chances by about 0.07, 2.60, and 2.84 percentage points, respectively. The example holds the number of launches fixed. The additional equipment, qualification, and operations for the fully equipped craft would have to be priced separately.

An added craft contributes $q \times (1 - p)^n \times p$ to branch success, where n counts existing craft equipped for that report. The survey branch already has three chances. The habitat branch has one, and the displaced craft carries no habitat package. The objective changes the value of the replacement.

A program comparing stars might prefer a farther branch even if another nearby ship improves the odds of one successful habitat. A second site can show whether a method travels beyond its first worksite. Track the first accepted result and cross-site replication as separate outcomes.

Measure the failures that travel together

A common dependency may lower a probability without making failure certain. Imagine two craft at one suitable site. A particular fabrication process has been qualified under two equally plausible operating regimes. Each craft has a 0.70 chance in one regime and a 0.10 chance in the other. Their remaining failures are independent within either regime, but both face the same regime. Averaging the per-craft chances gives 0.40. Treating those averages as independent predicts a two-craft chance of $1 - 0.60^2 = 0.64$. Averaging the probabilities of joint failure within the regimes gives $1 - (0.30^2 + 0.90^2) \div 2 = 0.55$.

A different process, calibration, or design can reduce dependence and add a penalty in mass, schedule, or performance. Compare alternatives under the same adverse regimes and record which fail together. A second assay can test a diagnosis about destination composition. A second destination changes the material under test. The report should identify which shared failure each change addresses.

33 How a Mission Earns Its Next Attempt

The first mission can take place on Earth. Put an archive on one side of a boundary and a workshop on the other. Give the workshop a declared inventory. Ask it to recover its instructions, test its instruments, construct a specimen, and produce a report from which another team can judge the result. Then change one condition and repeat the trial. A program earns a more demanding test when it can say why the previous one succeeded or where it failed.

Begin with separate chains of evidence

A readable archive and a viable lineage are different achievements. The workshop can verify a checksum over instructions that make a useless object. It can also produce a culture whose origin is unclear even though its records are intact. Both outcomes matter, and each calls for its own acceptance test.

For the archive, write down what a later reader needs to recover from the first working symbol to the last interpreted instruction. Use reference records with known outputs. Have a second team build a reader from the specifications and use it with test archives that contain declared errors. Test startup with the primary reader unavailable and test a valid archive accompanied by a wrong catalogue entry. A correct result includes detecting a disagreement the workshop cannot resolve and retaining evidence for the report. Copies with identical readers and identical software can survive independent media defects while sharing one failure of interpretation.

For biology, give independent workshops the same statement of the required result and different independently prepared inventories. Let each recover its record, assemble a contained test, and report the origin of material on either side of the handoff. The accepted lineage must satisfy the renewal measurements defined for that experiment. A damaged specimen, an unexplained lineage, and a loss of the reporting channel are three different failure records. Keep records of all initiated runs, including those that produced no interpretable result. A count of successful runs divided by the number of reports received will overstate reliability if loss of the reporting channel is related to the conditions that caused the other failures.

A later combined trial connects the archive, the workshop, the living system, and the transmitter. It should include at least one disturbance that reveals how

responsibility passes between them. For example, a failed assay could leave the organism growing while depriving the machine of the evidence it needs to declare a handoff. The receiver can preserve the culture, mark the result unresolved, and attempt a new observation without claiming a verified lineage. That decision exercises the part of the mission that links a living event to evidence another observer can use.

Lengthen the trial by qualifying mechanisms

A trial of extended operation tests the equipment that runs throughout the trial. A long period of dormancy followed by a restart tests a different set of mechanisms. Use separate specimens and instrumentation for both. Check the parts that must wake first and the readers, switches, joints, seals, and actuators that have to work after the power returns. Follow the order in which the actual craft would need them. A dormant actuator that a laboratory technician frees by hand has just identified an intervention that must appear in the flight account.

Accelerated tests can expose mechanisms whose rates depend on temperature, radiation, or time under load. A stronger stress may also change the mechanism. A useful extrapolation names the damage it is trying to predict, gives the range in which its model was tested, and records the effects that remain outside that range. A piece of memory media can pass an exposure and still require an untested connector to be read. The complete chain deserves a separate test under combined conditions. A cumulative radiation yardstick describes a selected exposure. It cannot stand in for that chain.

There is a simple way to see how much evidence a series of successful trials supplies. Suppose there are n independent repetitions of one specified trial, each with the same unknown probability f of failure. With zero failures observed, a one-sided 95 percent upper confidence bound is f = 1 − 0.05^(1/n). Thirty successes in thirty trials give a bound of about 9.5 percent. Three hundred give about 1.0 percent. Those are bounds under the stated model for that trial population. A hundred runs on the same instrument may say little about whether another manufacturing lot behaves similarly. A hundred short restarts say little about unobserved aging modes over a millennial crossing. Give separate entries in the record to sampling uncertainty, common design faults, and extrapolation beyond the tested conditions.

Set gates that can stop the sequence

A sequence of Earth and space demonstrations can address different uncertainties about the voyage. At each gate specify the evidence that can authorize the next trial, the failure that sends the design back for revision, and the safeguards that still apply while a decision is pending.

Demonstration	Evidence to carry forward	Reason to repeat or stop
Independent archive recovery	A second reader recovers the reference records and flags corrupted or ambiguous inputs	The fallback depends on the failed reader or gives a confident wrong output
Contained first-cell handoff	Independent runs report founder provenance, descendants, and renewal after withdrawal of organized assistance	The record cannot distinguish the claimed lineage from contamination or continued supplied organization
Unoccupied renewal trial	Recovered inputs support a subsequent cycle, floors are respected, and usable stocks recover at matched phases	Production continues by spending an unqualified reserve or importing an unrecorded input
Dormancy and integrated revival	Sealed assemblies recover power, read the archive, diagnose faults, and transmit a complete record under stated stresses	Restart needs a laboratory intervention absent from the flight design
Near-space precursor	Deployment, controlled motion, communications, and safe disposition are demonstrated at the tested scale	Evidence supports a component while the required sequence still fails
Interstellar mission review	A route-specific budget and an evidence record support the chosen experiment and its stopping rules	An unclosed capture, containment, or reporting requirement defeats the selected objective

The first three demonstrations can run independently, including an early crop facility seeded from an existing stock. The joint test matters before a flight claims it can carry out both operations from its permitted inventory. In the same way, a near-space precursor can exercise a real deployment and an optical link without proving that the distant star will supply the power or materials the workshop needs. Record what each result licenses the program to try next.

Decide what the unattended machine can do

The ship can arrive in a system that Earth understood poorly at launch. It can also arrive after the institutions that prepared its instructions have changed beyond recognition. A rule that requires a message from Earth before every action would

make routine operations depend on years of light time and on the survival of the sender. A rule that treats a missing reply as approval would turn a communications failure into permission. A workable division gives the ship authority for specified reversible surveys and contained operations, with more consequential actions held to independently satisfied criteria. The uncertainty about native biology that motivated Chapter 26 remains relevant after a survey has failed to detect it.

A workshop should have a plan for its disabled states as well as its successful ones. If a planned site cannot be verified, the machine may hold its cultures sealed, transmit the measurements, and keep the capability to examine a different site within its qualified envelope. If control of its trajectory is lost, the mission ought to know whether the resulting paths threaten a protected target. Design choices such as a qualified holding orbit and the ability to isolate the biological enclosure make that obligation physical. A trajectory whose long-term disposition depends on indefinitely available steering needs its own reliability account.

A valuable early result might be a sealed, inspectable record of a workshop that failed to build anything living. It could show that the archive survived, that a component started on available light, or that a material once thought accessible could not be recovered with the tools aboard. A valuable later result might establish renewal in a chamber without authorizing use outside it. That is how the fleet accumulates evidence at each step. Its next attempt begins from a known condition instead of from a fresh assertion about everything that came before.

34 Keeping the Record Alive

Who will be listening

The receiver has finished its work. It has stored the local measurements, the complete execution record, and enough images to see what happened when the apparatus stopped helping. It can point its transmitter at the old position of Earth after allowing for motion. Its original instructions name the receiving network. Ten thousand years is a long time for a network to keep the same name, use the same clocks, or accept a call on the same frequency. A signal can reach its intended star without reaching anyone who knows how to listen.

The transmitting problem in Chapter 30 has two ends. On the ship, a transmitter has to deliver enough energy toward the right patch of sky. At the destination of the signal, someone has to look in that direction, at the right frequencies and times, with a working detector. A design can test both ends with equipment we have now. The longer promise requires a succession of institutions and observing instruments. Name that succession as part of the proposed experiment. Its survival is a different question from the survivability of the craft.

For each route, start with an estimated arrival window in a common timescale. Carry the flight time, local qualification time, and propagation time as separate entries. An uncertainty of six months in commissioning and two years in arrival gives up to two and a half years of displacement on either side of a nominal reporting date if both numbers are symmetric hard bounds and the errors can have the same sign. An observer may need a longer window when the receiver has discretion to pause its work. A bound tells the observer where to put telescope time. A probability distribution would serve another purpose and would require evidence about how likely the delays are.

An observer also needs a place to look. Record the coordinate system, reference epoch, stellar position and motion, uncertainty, and the geometry of the intended observer. Send the absolute transmit times when the clock is trustworthy. Also send a sequence number and the interval since the event, which can remain useful when the ship loses its absolute time reference. For a craft that arrives from the wrong direction, a position reconstructed from its survey may tell a later generation why a search based on the launch plan came up empty.

Put the small report first

A full experimental report can be much larger than the first thing the receiver can reliably deliver. Suppose a conservative net rate after coding is one bit per second for the time the transmitter is actually on target. A 256-byte status record takes about 34 minutes of useful transmission. Ten copies take about 5.7 hours. A one-mebibyte evidence package takes about 97 days of continuous useful transmission, or about 2.7 years if useful contact is available ten percent of the time. These numbers are scheduling examples. They require a link that has first established the assumed rate and they leave out time spent acquiring the signal before useful transmission.

Suppose the receiver estimates that a short listening opportunity remains. It should send a record that is still interpretable if the long archive never follows. The status record names the craft and the test version, says how far it proceeded, gives the last accepted checkpoint, and identifies the evidence package by a digest. A set of flags can distinguish a broken instrument, insufficient observations, a stopped trial, and a completed trial that missed its criterion. Give each status a meaning in the archived format. The same radio silence can follow any of those events if the craft cannot transmit.

After the short record, send an evidence summary that can be read independently. This can contain the number of runs, the intervention cutoff, the observed range of outcomes, and an inventory of the samples and files that support the result. Then send the raw records in blocks with their own identifiers, lengths, and integrity checks. The evidence package has to tolerate the loss of some blocks without losing the meaning of those that remain. A receiver on Earth can report which portions it has and which portions it could independently interpret.

Send the short report again at published intervals. Copies sent in one optical pointing episode share the same alignment and the same receiving weather. A repeat from a different antenna, season, or receiving station can cover a different failure. The relevant question is the one we asked of the fleet. Which condition has to change before a repeat adds an independent chance?

Make an archive that a stranger can test

The archive at Earth has a job besides retaining bits. A future reader needs the specification that explains how the bits become integers, timestamps, spectra, and observations. The space-data archive reference model treats that interpretation as information in its own right and anticipates that a later community can need more

explanation than the community that first received the record (CCSDS 2024). Apply that lesson at the level of the experiment. A later reader needs to know how a count was made, what instrument reported it, what units and calibrations apply, and how a missing measurement is encoded.

Give each edition of the test a set of small worked records. For instance, publish the bytes for a dummy measurement, the expected decoded row, the units conversion, and the numerical result of one comparison with a reference. Include at least one deliberately damaged record and one valid record containing a measurement marked unavailable. A reader that treats those as the same event has located a problem in its own decoding before it evaluates a distant experiment. Preserve these test vectors beside the machine-readable specification and a human-readable summary.

One terrestrial trial would keep the receiving team outside the design meeting. Give it a simulated transmission from a test craft, an archived decoder description, and whatever observation logs a future receiver would actually possess. Give it a known mixture of partial reports and conflicting evidence. Ask it to recover the same accepted, failed, or indeterminate result as a separate reviewing team. Once the teams have recorded their answers, compare the assumptions they had to supply. Then repeat the trial on a different implementation and with someone responsible for a different part of the archive. The cost of the documentation becomes visible when someone tries to use it.

An authenticated command needs its own test. A key might prove that an update came from whoever controls that key, while leaving questions about whether the key was inherited legitimately, whether the command is new, and whether the current instruments and containment can support it. A ship can check the planned scope, verify the update against its allowed actions, and keep the existing observation program when those checks fail. An authorization to collect a spectrum has a different consequence from an authorization to begin a living experiment. Preserve that distinction in the command format and in the physical controls.

Test succession before launch. Give a new custodian the archive, the listening plan, and a simulated transmission without the original staff. Ask the custodian to find and interpret the record. Run a corresponding exercise at the simulated destination with a team that inherits the craft's instructions. The mission returns an experiment when each end can pass its part of that test.

Appendix A Reproducing the Crossing Calculation

I use a reference distance of 4.25 light-years and three material crossing speeds and a light path. The material journeys coast at constant speed. The light path describes a transmission to a receiver already in place. Departure and braking time, stellar or planetary motion, and settlement time are outside this simple comparison. A year is 365.25 days for the dose extrapolation.

Route	One-way time	What has to reach the destination
Light signal	4.25 years	Detectable encoded radiation and a receiver at the right place and time.
Cargo at 0.10c	42.5 years	A physical package capable of surviving flight and deceleration.
Cargo at 0.01c	425 years	A physical package with a much longer operating or storage life.
Cold receiver at 0.000425c	10,000 years	A recoverable archive, restart hardware, and a workshop.

How the dose comparison is made

For the illustrated case, the Mars Science Laboratory cruise rates are 0.48 milligray per day for tissue-like absorbed dose and 1.84 millisievert per day for the reported dose equivalent associated with galactic cosmic rays. These values describe a particular detector and shielding environment (Zeitlin et al. 2013; Hassler et al. 2014). Multiplying by 15,523 days gives 7.45 gray and 28.6 sievert for the 42.5-year yardstick. The 425-year yardstick is ten times as large. At 10,000 years the same multiplication yields 1,753 gray and 6,721 sievert. Use those products as arithmetic, with route-specific dosimetry and cargo recovery tests treated separately. The dose equivalent incorporates a weighting that reflects the measured radiation quality. A high dose equivalent warns about the demands on a crewed route. A specimen test requires the dose and particle field in the specimen's own packaging and a viability endpoint appropriate to the organism.

Putting mass and speed on the same account

An areal density of 10 grams per square centimeter equals 100 kilograms per square meter. A sphere of radius two meters has a surface area of 50.27 square meters. A thin shell at that areal density thus adds about 5.03 tonnes. Doubling the areal density doubles this estimate. The calculation assigns mass to a geometry. A

transport model would still have to compute the particle spectrum behind the shield, including secondaries and the effects of nonuniform coverage.

The kinetic energy of a kilogram coasting at a fraction b of light speed is c^2 times the quantity one over the square root of 1 minus b^2, less one. It is about 4.53×10^{14} joules per kilogram at 0.10c and 4.49×10^{12} joules per kilogram at 0.01c. This figure gives the energy change of the arriving mass. A full propulsion budget includes the efficiency, mass, and energy of the means used to accelerate it, plus the means for losing the kinetic energy at arrival. A proton that is nearly at rest in the local interstellar medium has about 4.73 megaelectronvolts of kinetic energy in the vehicle frame at 0.10c. That scale is separate from the galactic cosmic-ray dose yardstick. It helps explain why a vehicle encounters different problems at its nose and at the sides of a shield.

Each line of the account can be changed without changing the others. A different shield geometry changes mass and dose. A different speed changes the exposure interval, the relative energy of incoming gas, and the energy required for the flight. A different cargo changes what counts as viability after the crossing. That is why the architecture choice has to be tested with the cargo named.

Arrival calculations

Divide 4.25 light-years by 10,000 years for 0.000425 times light speed. With light speed at 299,792 kilometers per second, the reference speed is 127.4 kilometers per second. The low-speed kinetic energy per unit mass is half the square of speed, or about 8.1 billion joules per kilogram. The chapter on capture uses a round 127 kilometers per second.

In a point-mass model the local escape speed scales as the inverse square root of radius. At one astronomical unit from a solar-mass star it is 42.1 kilometers per second, and at a tenth of that distance it is 133.2. Combine each escape speed in quadrature with the far-field excess speed to get the ideal approach speed at periapsis. Subtract escape speed for the marginal impulsive capture threshold. An actual trajectory needs to go below that threshold and remain usable after the burn.

Electrical power in the example is incident irradiance times projected collector area times conversion efficiency and system factor. Irradiance scales with stellar luminosity and the inverse square of distance. For electric thrust, twice the useful exhaust power divided by the assumed exhaust speed gives thrust. The propellant fraction for a given velocity change is one minus the exponential of the negative ratio of velocity change to exhaust speed. Neither calculation includes tank mass, equipment degradation, or the temporal availability of the stated power.

From a distance to an intercept

Dividing a catalog distance by 0.000425 gives 2,353 years per light-year. The destination table uses this ratio as a fixed-distance yardstick for comparing routes.

For a straight, constant-speed estimate with stellar motion, write the initial position of the target relative to the Sun as a vector r and the target's heliocentric velocity as a vector u. The distance to the predicted point of encounter after time t is the length of r + u × t. An intercept solves for a positive time at which this length equals the ship's heliocentric speed multiplied by t. The ship aims at that point. The speed relevant to capture is the magnitude of the difference between its velocity vector and the star's.

For a rough Barnard example, take 5.96 light-years, a radial approach of 110 kilometers per second, and a transverse speed of 90 kilometers per second. At 127.4 kilometers per second the craft can match the transverse component and move toward the encounter along the original line of sight at about 90 kilometers per second. The two close radially at about 200 kilometers per second. The earlier straight-line intercept then comes in about 8,900 years. The relative speed is also about 200 kilometers per second, which changes the capture problem. A real ephemeris would propagate the system and its uncertainties and account for the launch, en route corrections, and approach to an actual working site.

Appendix B Reproducing the Habitat Calculation

This appendix keeps the accounting for one nominal adult in units that can be checked with a calculator. It uses 365.25 days per year, so the faster, idealized 42.5-year crossing used for the living-cargo comparison is 15,523.125 days. That duration is a comparison interval. The same equations can be applied to a person living in a destination habitat over a much shorter or longer period. A constant daily rate represents a planning average. It leaves time variation in the occupant, the crops, and the equipment to a separate calculation.

Gross flows over a long interval

Jones (2016, Table 1) gives 0.84 kg of crew oxygen, 1.00 kg of carbon dioxide, and 0.62 kg of food solids per crew member per day in a spaceflight planning case. These are source inputs to the examples here. The two water throughputs are independent illustrative cases. They represent a compact personal loop and a wider loop with additional uses. Multiplying each daily flow by 365.25 or 15,523.125 gives the corresponding gross mass.

Stream and daily rate	kg per year	Tonnes in 42.5 years
Oxygen consumed at 0.84 kg	306.81	13.039
Carbon dioxide produced at 1.00 kg	365.25	15.523
Food solids eaten at 0.62 kg	226.46	9.624
Water circulating at 3 kg	1,095.75	46.569
Water circulating at 30 kg	10,957.50	465.694

The carbon dioxide entry is a product stream. Multiply it by 12.01 divided by 44.01 to count carbon atoms by mass. About 99.7 kg of carbon passes through this stream in a year. A single stock of carbon can pass through the occupant, atmosphere, and plants repeatedly. If a fraction remains in a reservoir inaccessible to the crops, the stock available for future cycles falls by that fraction of the transfer. The oxygen and carbon dioxide rates also imply an average of about 0.87 moles of carbon dioxide produced per mole of oxygen consumed. A simplified carbohydrate-

production reaction yields one mole of oxygen per mole of carbon dioxide fixed. In that reaction, fixing the daily kilogram of carbon dioxide would return about 0.727 kg of oxygen. Closing the actual gas balances calls for accounting for the full set of reactions and stocks, including the composition of food, unharvested biomass, and decomposition.

Energy and area for the food loop

Use a daily food-energy target of 2,500 kilocalories. With 4,184 joules per kilocalorie, the target is 10.46 megajoules per day, equivalent to an average of 121 watts of dietary energy. Divide it by the daily radiant energy incident per square meter multiplied by the fraction recovered as edible chemical energy. The 20 megajoules per square meter per day and the conversion fractions below are scenario assumptions. The fractions include losses between incident energy and edible harvest, so an additional harvest fraction would need a different definition of efficiency.

Fraction recovered as food	Edible MJ per square meter per day	Calculated growing area
0.5 percent	0.10	104.6 square meters
1 percent	0.20	52.3 square meters
2 percent	0.40	26.2 square meters

This is an area exposed to the stated average light input. Occupied floor area depends on whether the crops are stacked, how much room the system needs for paths and machinery, and how much light is lost before it reaches the plants. A harvest record must also account for edible composition, crop failures, and periods when plants are growing but no food is available. The factor of four between the first and last rows shows what uncertainty in a single conversion fraction can do to the area.

The electric-light example uses a separate ledger. Take 60 square meters of illuminated crops, with 25 moles of incident photons per square meter per day and 2.5 micromoles delivered per joule of electric input. The crops receive 1,500 moles of photons per day. Divide by 0.0000025 mole per joule to obtain 600 million joules per day. That is 166.7 kilowatt-hours per day, or 6.94 kilowatts averaged over 24 hours. This efficacy is taken at the receiving plane in the example. If the stated efficacy describes emission from a fixture instead, optical delivery losses have to be included. Pumping, processing, and climate control have separate electric loads.

Over a repeating cycle, absorbed electric and radiant energy must leave as heat or in exported material and work, except for the portion still stored at the end.

Atmospheric stocks and a five-day interruption

The gas model uses a dry 100-cubic-meter chamber initially at 101.325 kilopascals and 293.15 kelvins. Take 20.95 percent oxygen and 0.04 percent carbon dioxide by mole. The ideal gas law gives total moles as pressure in pascals times volume in cubic meters divided by 8.31446 times the absolute temperature. This gives 4,157 moles of dry gas. Multiplying by the oxygen mole fraction and then by 0.032 kg per mole gives 27.87 kg of oxygen initially.

For five days, count only the assumed human gas exchange, with all plant, soil, ventilation, and processing fluxes set to zero. Subtract 4.20 kg of oxygen and add 5.00 kg of carbon dioxide. The oxygen stock falls to 23.67 kg. Converting the two transfers to moles changes the total gas count as well. In the model, at the same temperature and volume, oxygen reaches 17.87 percent and carbon dioxide about 2.785 percent by mole. These are outputs of the stipulated mass balance. The result illustrates how a change in a reservoir depends on its initial size and on all the fluxes omitted from the simplified balance. A biosphere contains many additional sources and sinks, and the calculation supplies no occupancy limits.

Recovery fractions and their denominators

In a water loop with gross throughput Q and effective whole-loop recovery r, the daily makeup burden in this model is Q multiplied by the unrecovered fraction, 1 minus r. The results below assume constant flows and no other sources of makeup. They count losses from the usable water loop, including material sequestered inside the boundary if the loop cannot recover it.

Effective recovery	Loss at 3 kg per day in 42.5 years	Loss at 30 kg per day in 42.5 years
95 percent	2.328 tonnes	23.285 tonnes
98 percent	0.931 tonne	9.314 tonnes
99.5 percent	0.233 tonne	2.328 tonnes

The International Space Station demonstration reported by Gaskill (2023) illustrates that high water recovery can be achieved in a specific engineered system. Its result applies to the streams and apparatus accounted for in that demonstration. For a habitat, collection and treatment both matter. If each stage recovers 99 percent of

what reaches it, and unrecovered water is unavailable for later recovery, the serial stages return 98.01 percent of the original flow. A processor that returns 98 percent of what it receives returns 88.2 percent of the whole flow when only 90 percent reaches it. The remainder still appears in the overall balance, whether as export or as water that no longer circulates.

Nutrient loss can be expressed against stock or against throughput. Consider an illustrative available inventory of 10 kg of nitrogen and 1 kg of phosphorus, both with a loss of 1 percent of remaining inventory per year and no replacement. After 42.5 years the available fraction is 0.99 raised to 42.5, or 0.652. About 3.48 kg of nitrogen and 0.348 kg of phosphorus have become unavailable. In a separate model, suppose 12 g of nitrogen and 1 g of phosphorus circulate per day and 0.1 percent of each daily transfer is made unavailable. Over the same period the losses are 0.186 kg and 0.0155 kg respectively. Some unavailable material may remain in sediment or on hardware, where a separate process could eventually return it to circulation. The rates are hypothetical and their denominators are essential to interpreting them.

Checks on the arithmetic

Doubling the incident light in the area equation halves the area if edible conversion stays fixed. Doubling all chamber dimensions multiplies its initial gas stock by eight at the same pressure and temperature. The same five-day oxygen draw would then remove one eighth as large a fraction of that stock. A loss-free recovery calculation gives zero makeup under the stated assumptions, and a zero-recovery calculation gives the gross flow. Two identical people double the human-flow terms when all their per-person rates are held fixed. Growing area, machinery, and reserves can scale differently once the habitat is treated as a coupled system.

The important checks compare initial and final stocks as well as flows. Food can appear during an interval while accessible soil nutrients decline. Recovered water can leave a treatment device while the chamber loses vapor somewhere else. Oxygen concentration can fall while a solid surface takes up the corresponding carbon dioxide. For each quoted fraction, record what crossed the measuring point, the duration over which it was counted, and the reservoirs still available at the end. An operating demonstration can then test whether the initial stock was spent or renewed.

Appendix C Tests on a Received Sequence

The reference is the linear presentation of the circular *Escherichia coli* K-12 MG1655 chromosome, GenBank accession U00096.3. Its 4,641,652 bases map to 9,283,304 bits under A = 00, C = 01, G = 10, T = 11. The blind tests use bit values and positions. The receiver gets no list of genes or known codon boundaries for those runs.

The spectral control

The upper panel averages spectra over 1,547 windows of 3,000 bases each. The control shuffles bases within each window and preserves their counts. The gray line averages eight such shuffles. Power at one cycle per six bits is 7.54 times the shuffled power. At one cycle per three bits it is 7.10 times the shuffled power. Power at one cycle per two bits remains the same. The lower panel shows the cumulative difference between G and C counts, with the mean trend removed, and uses a separate whole-sequence shuffle.

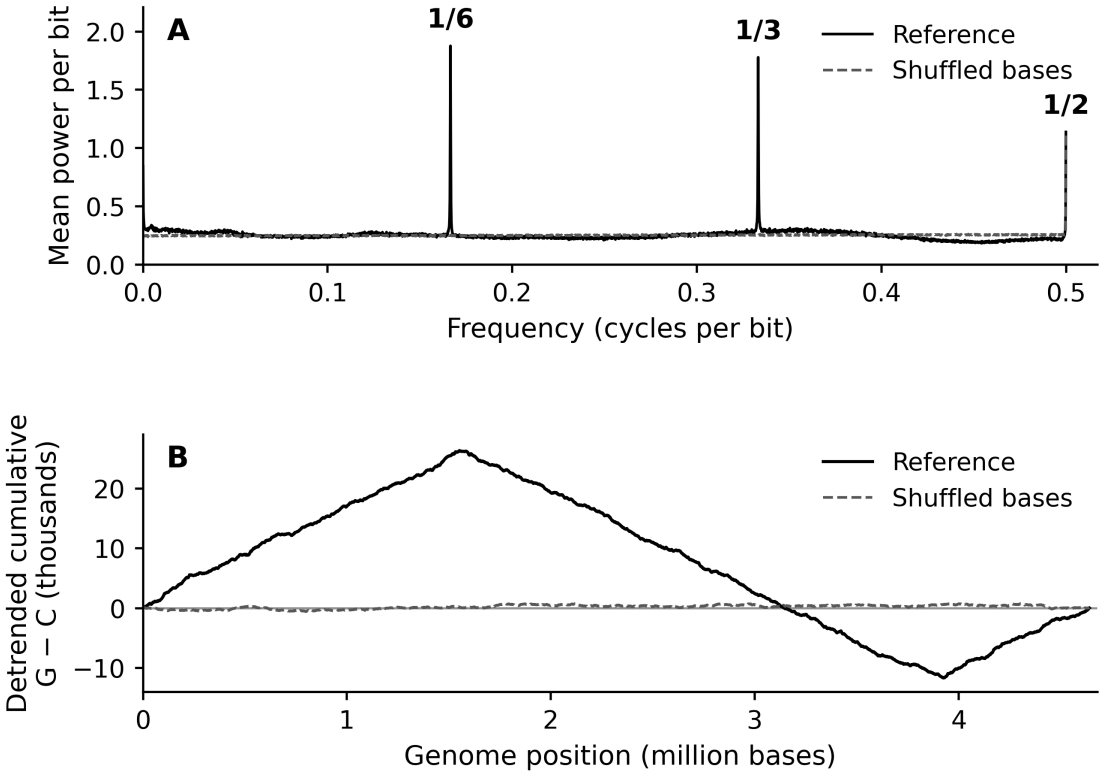


Figure 4. (A) Mean power in the MG1655 bitstream and eight base shuffles. (B) Detrended cumulative G − C along the reference and after a whole-genome shuffle.

The separate blind scan used 2,048-bit windows and ranked periods from 2 to 128 on the first half. Within the second half, the six-bit component had roughly five times the power of a shuffle that kept the odd and even positions' ones counts separately. Averaging power makes the test sensitive to local rhythmic structure even if phases differ across regions.

Prediction and recurrence

A smoothed binary model predicted the next bit from recent bits. The first half of the stream supplied the initial fit. The next quarter selected the history length, and the first three quarters supplied the final fit. A 13-bit history gave a test log loss of 0.98047 bits per received bit on the final quarter. The independent-bit model gave about one. Across the 2,320,806-bit test section the difference is approximately 45,300 bits, before the model description is charged.

An exact-repeat search compared anchors spaced 64 bits apart with seeds at every bit position and extended matches in both directions. The 5,630-bit match starts at zero-based offsets 8,337,236 and 8,420,040. Shuffling 96-bit chunks within local regions of up to 96,000 bits preserves some short-range structure. In three such controls no match reached 256 bits. The repeat can serve as evidence of common material. A receiver looking for record boundaries should also examine where similar copies continue beyond the exact match.

After the blind search was finished, I checked the reference annotations. The two starting offsets correspond to bases 4,168,619 and 4,210,021 in one-based chromosome coordinates. Each exact block starts 22 bases before an annotated 23S ribosomal RNA gene and overlaps most of it. The two genes are rrlB and rrlE. The bit-level test supplied the locations. The GenBank annotation supplied a biological interpretation for them (NCBI, U00096.3).

The reversible map experiment

Pairs of adjacent words were tested for fixed width and starting alignment. The first half selected a reversible lookup map that optimized excess agreement over local independence. The baseline holds input-word and output-word frequencies within blocks of up to 4,096 pairs and breaks their pairing. On the second half, a two-bit SWAP agreed with 30.39 percent of output words, compared with 25.09 percent in the corresponding baseline. Three-bit and four-bit reversible maps gave 17.45 and 9.48 percent, compared with 12.61 and 6.30 percent. These are modest statistical preferences under a defined framing. They show how a receiver can measure an

apparent rule against its alternatives before making it part of a construction instruction.

The reference and the computational companion

The sequence experiments use the linear GenBank presentation of the circular *Escherichia coli* K-12 MG1655 chromosome, accession U00096.3. The reference contains 4,641,652 bases. Reading in the order stored and assigning A = 00, C = 01, G = 10, and T = 11 produces 9,283,304 bits. Coordinates are zero-based unless otherwise specified. No window wraps around the end of the reference. The companion verifies the SHA-256 of the normalized sequence and prints the observed results. Its seeded controls are independently specified reruns, so values from random shuffles need not coincide with earlier values. The code uses NumPy's PCG64 generator and records a seed for each control.

Appendix F prints the complete source, dependency versions, reference checksum, and commands for computational companion version 1.0. The companion ZIP contains the same files in runnable form.

Two ways to measure periodicity

The discovery test divides the bitstream into complete 2,048-bit windows and splits these windows into two equal groups in genome order. For every integer period from two through 128, it subtracts the mean in each window, projects the result onto a complex sinusoid at the exact corresponding frequency, and averages the squared magnitudes across windows. A comparison stream shuffles bits within each window separately for the odd and even positions. With seed 2027 the first group ranks periods by the ratio of observed to shuffled power. On the second group, the two top-ranked periods, six and three bits, give ratios of 5.13 and 5.35. At two bits the ratio is 1.00. The first two peaks are compatible with a three-base rhythm. This analysis does not locate or identify coding regions.

The upper panel of Figure 4 uses a separate control. It averages Fourier power over 1,547 complete windows of 3,000 bases each, with each window encoded as 6,000 bits. It compares the observed result with the mean of eight sets of independent within-window base shuffles using seed 2026. The ratios at one cycle per six, three, and two bits are 7.54, 7.10, and 1.00. Shuffling the bases preserves both positions of each two-bit word and therefore preserves their overall imbalance. The lower panel accumulates G counts minus C counts along the chromosome and subtracts a straight line from zero to the final difference. Its control shuffles the entire sequence

of bases once using seed 2028, preserving the total G and C counts while scrambling their positions. A rhythm and a strand asymmetry are observations about the recorded sequence under this encoding. Their biological interpretation takes more evidence.

Recurrence, prediction, and word maps

The repeat search compares 64-bit anchors at multiples of 64 with 64-bit seeds at every bit position. It extends matching seeds in both directions. It finds a 5,630-bit match starting at offsets 8,337,236 and 8,420,040. In the next 256 aligned positions outside the match, 253 bits agree on the earlier side and 250 on the later side. Three controls permute 96-bit chunks within regions of up to 96,000 bits using seed 2029. Their longest detected matches have 105, 106, and 103 bits. A repeat of at least 127 bits necessarily contains a complete aligned anchor somewhere in its first copy, so the search would detect any 256-bit match in these controls.

For prediction, count how often each bit follows a given history of zero through 16 bits. Add half a count to each possible outcome. The first half supplies estimates, the third quarter selects the history length by mean base-two log loss, and the first three quarters fit the final model. The last quarter gives 0.98047 bits per received bit for the selected 13-bit history, compared with approximately 1.00000 for a history of length zero. This is an ideal coding charge before the receiver pays to obtain the model.

The reversible-map test reads successive pairs of words of a fixed width. For each candidate offset, it keeps one alignment across the entire stream and splits the complete pairs into two halves. The first half selects an offset and a one-to-one map that maximize agreement in excess of local independence. The control computes expected agreement after breaking the pairing while preserving both word histograms within blocks of up to 4,096 pairs. The selected two-bit map is SWAP. On the second half it agrees 30.39 percent of the time, compared with 25.09 percent under local independence. An ordinary count of adjacent bases in the first half already gives 30.31 percent agreement with that map. This helps explain why a simple table of conditional frequencies predicts better than a model centered on the map. It also shows how much a baseline matters when looking for a rule in a stream.

The hydropathy comparison

Chapter 12’s hydropathy calculation gives all 64 codons equal weight and assigns stops a fixed zero placeholder. For each of the three positions, group codons by the letter in that position and sum their hydropathy values. Squaring the twelve sums and adding them orders the comparison scores because the total mean and variance stay fixed across the restricted permutations. The companion represents hydropathy in integer tenths and compares scores exactly. It permutes numerical values only among synonymous classes with the same number of codons, keeping the stops fixed. The 21,772,800 distinct numerical assignments include 1,920 that meet or exceed the reference score. The fraction applies to this defined ensemble of tables. It measures how unusual the observed fit is in that comparison family.

Appendix D Recording an Attempt and Its Outcome

A record for each attempted transition

Use this as a planning sheet and as a template for reporting an actual trial. Fill the planned columns before opening the chamber, powering the receiver, or starting the simulated load. Record both the state immediately before the declared cutoff and the result after the observation interval. Preserve the raw quantities and calibration records that support every derived number.

Field	What the record should contain
Identity and authority	Craft or laboratory, test version, site, approving role, applicable containment limits, and time reference
Endpoint	Observable that qualifies, acceptable uncertainty, number of independent attempts, observation duration, and decision rule
Starting inventory	Quantities, provenance, working organization already present, spares, control references, and measurement uncertainty
Actions before cutoff	Fabrication, repairs, parameter changes, and interventions on particular cells, cultures, crops, or machines
Actions after cutoff	Allowed flows and scheduled services, every intervention, and any reason an intervention ended the test
Reserves	Initial, minimum, and final usable stock measured at comparable phases, with material temporarily unavailable tracked separately
Evidence	Data and calibration locations, controls, interruptions, comparisons with the declared criterion, and integrity digests
Disposition	What remains powered, contained, accessible to later experiments, or deliberately isolated after the outcome

A worked decision

Suppose a one-person load emulator must receive at least 10.46 megajoules of usable food-equivalent energy per day for 30 days. The illustrative rule also requires a reserve of ten demand-days throughout and a post-delivery reserve at least as large as the one at the start of the cycle. Those are three separate comparisons. Initial reserve is 42 demand-days. The tank has 12 demand-days just before delivery, and the next batch contributes 29.5 demand-days. Reserve ends at 41.5 demand-days. The daily load and the ten-day floor were met. The

replenishment requirement missed by half a demand-day, or 5.23 megajoules. Each comparison belongs in the record.

Now suppose each stock measurement has an independent standard uncertainty of 0.4 demand-day and assume normal errors for this example. The standard uncertainty of the change is about 0.57 demand-day. A conventional approximate 95 percent interval about the observed change of minus 0.5 extends from minus 1.61 to plus 0.61 demand-days. If the predeclared rule requires the lower bound of the change to be at least zero, the run fails to demonstrate renewal. The interval also leaves the sign of the actual change unresolved. A second, calibrated measurement of net input and output could narrow the interval, as could repeated comparable cycles with independently calibrated measurements. If the same scale drifts during every cycle, averaging alone leaves the common bias in the estimate.

The report can therefore preserve progress without confusing the endpoints. It can say the load was met and the floor was maintained, state the estimate and uncertainty for the stock change, and give the predetermined outcome for replenishment. A reader can then identify what would have to be improved to qualify a longer trial. Use the same decision record for a biological lineage by replacing the reserve with a named working component, declaring the post-handoff input, and recording the evidence that the component was replaced.

Appendix E Comparing Lineages and Habitats

What the fleet counts

A mission record has two questions to answer about a lineage. Did the descendants renew their working machinery under the declared conditions? Where did the organization that started the lineage come from? Answer both in the first report. A revived spore can settle the first question while tracing its founder to Earth. A receiver can build an object from an audited inventory and still owe evidence that the object can renew itself. The two questions use different parts of the same record.

Chapter 32 counts a verified renewing lineage from any declared cargo route that meets the common test in Chapter 31. The illustrative 55.55 percent is the modeled chance of receiving at least one such result after all assigned attempts have had their opportunity. The model assigns no separate chance to strict abiotic success. To calculate that chance, specify which craft can attempt the stricter build, how their starting inventories will be audited, and which failures they share. A dormant seed can contribute to the general lineage result even when every attempt at abiotic construction fails.

The same discipline applies before the chamber opens. A craft that misses capture can send a useful flyby record with its observation times, instrument settings, and uncertainties. Give that return its own entry. The qualifying survey in Chapter 32 requires sustained operation in orbit as well as calibrated observations sufficient to assess a proposed worksite. A strong flyby record can influence where a later ship aims. Its arrival on Earth still leaves the orbital survey milestone open.

Three chambers and a fourth record

Suppose three laboratories exercise the cargo alternatives with the same acceptance rule. Each has to follow a documented founder and its descendants for 30 days and at least 20 net doublings after handoff. The laboratory maintains the environment and supplies the declared simple inputs. The evidence must include renewal of a named component under those conditions. Use a qualified reference to check that the measurements can distinguish the intended result, and a blank that goes through the same handling. Keep the observations and control outcomes with each decision.

The first chamber revives a dormant founder. The pedigree is continuous, the blank stays clear, and a renewal assay shows that descendants made a designated working component from post-handoff inputs. Count the verified lineage and record the founder as Earth-origin. The result can qualify the storage and recovery route under the tested history. If the trial used a laboratory exposure in place of the crossing, record the exposure actually delivered.

The second chamber assembles a founder from an audited nonliving inventory. Its recipe starts with simple materials, and the ledger shows the source and preparation of every organized part used in the build. No viable material or cell-derived working component enters the chamber. The assay meets the same ancestry, duration, doublings, and renewal criteria. Count the verified lineage and the strict abiotic handoff. If the assembly instead started with a working complex purified from an organism, record that contribution. Such a build can provide evidence of renewal from a nonliving inventory and would fall outside the stricter cargo class.

The third chamber belongs to a hybrid kit. A component made on site helps revive a dormant founder. If its descendants pass the same test, count a verified lineage with an Earth-origin founder and document the repair. A second chamber could attempt a strict abiotic build using the same craft. To evaluate that claim, trace the second chamber's inputs, handling, and exposure to the viable stock separately. If its blank grows a culture, keep the observations and mark the trial invalid for acceptance. An invalid trial and a valid trial in which the founder stops dividing give different guidance for the next attempt.

A fourth record can then describe the dependencies among the chambers. Suppose the dormant and hybrid kits contain two separately packaged lots drawn from one preparation. Two successes offer evidence that each package can work. A preparation error could still affect both packages on the next flight. A repeat starting from a different preparation would address that dependence. A second apparatus with an independently calibrated assay could challenge a shared measurement error. A new worksite tests transfer to different conditions. Keep these distinctions when reporting replication, so a reader can see which doubt the next trial reduced.

The origin of a working habitat

A qualifying lineage in one chamber does not specify the origin of the rest of a habitat. In the habitat ledger, give each population a function and a source. A dormant producer, a microbial partner assembled in a receiver, and a recycler recovered from a transported stock can take part in the same trial. The resulting

system has mixed provenance. It can still demonstrate renewal if the report identifies what each population does and verifies the declared recovery test. If the claim is that the habitat began without viable imports, the provenance account has to cover every population whose living work the system needs. A single abiotic lineage cannot supply that evidence for its neighbors.

Consider a contained habitat with two separately measured phosphorus pools. Begin an illustrative recovery cycle with 100 grams of phosphorus in forms the next growing interval can use and 10 grams in a temporarily inaccessible pool. A planned interruption slows the return of phosphorus from residue. At the lowest observed point the usable pool is 92 grams, and the other pool has risen to 18 grams. By the corresponding phase of the next cycle, return processing leaves 102 grams usable and 8 grams inaccessible. Suppose the measurement bounds are one gram either way at the low point and half a gram either way at the two comparable checkpoints. Under those bounds the low point stays above a 90-gram floor, and the closing usable stock exceeds the opening stock by at least one gram. The point estimates sum to 110 grams at all three checkpoints. The instrument has to distinguish usable forms from the other pool for this comparison to mean anything.

That ledger supports one phosphorus decision. The same recovery report has to cover the other named reserves, viable planting stock, required outputs, containment, and the energy and maintenance supplied during the interruption and recovery. It should say whether the processing equipment retained its ability to run the next cycle. If a single-use cartridge restored the usable pool, the system spent part of its recovery capacity. Record that expenditure beside the phosphorus transfer. A second run on a fresh cartridge answers a question about another supplied kit. A run that restores or replaces the cartridge from the declared stocks addresses a further question about continuation.

Finally, give the destination trial and the Earth acceptance record separate dates. A short status message can say which endpoints passed locally and which evidence remains aboard. If the complete record arrives later, enter acceptance when the required evidence can be assessed. A second workshop may then compare both the biological result and the conditions that allowed a distant reader to verify it.

Appendix F Running the Computational Companion

The Fleet That Wakes on Starlight computational companion, version 1.0, reproduces Appendix C and the hydropathy calculation in Chapter 12. The complete Python source appears below. The same files are distributed in the companion ZIP. The only input data needed for the sequence analysis are the public MG1655 reference bases.

Reference sequence and setup

Use Python 3.12. Obtain NCBI accession U00096.3 as FASTA. The download option below obtains it from NCBI if no local copy is present. On an offline machine, copy the FASTA into the working directory and pass its path with the --fasta option. The program hashes the concatenated uppercase A, C, G, and T bases, with headers, whitespace, and trailing newline excluded.

https://www.ncbi.nlm.nih.gov/nuccore/U00096.3?report=fasta

```
e999dceccb76562d40e49a10d9e6d4cf1c05d184f4f80b92bb32d4a2420c11d0
```

Put the dependency lines in requirements.txt and each of the three listings in a separate file with the name given at its head. Preserve indentation. Run the following from the directory containing those files.

```
python -m pip install -r requirements.txt
python sequence_tests.py --download --figure-data figure_data
python hydropathy_test.py
python plot_sequence.py figure_data figure4_print.png
```

The first analysis prints the reference checksum, spectral controls, repeat coordinates, prediction loss, and reversible word-map comparisons. The hydropathy script prints the exact enumeration. The last command renders a grayscale figure from the spectrum and strand-skew data written by the first. Appendix C gives the expected numerical results and describes the control seeds.

requirements.txt

```
numpy==2.3.5
scipy==1.17.0
matplotlib==3.10.8
pillow==13.0.0
```

sequence_tests.py

```
#!/usr/bin/env python3
"""Independently specified reruns of the book's MG1655 bitstream examples."""

import argparse
import hashlib
from pathlib import Path
from urllib.request import Request, urlopen

import numpy as np
from scipy.optimize import linear_sum_assignment, minimize_scalar


ACCESSION = "U00096.3"
FASTA_URL = (
    "https://eutils.ncbi.nlm.nih.gov/entrez/eutils/efetch.fcgi"
    f"?db=nuccore&id={ACCESSION}&rettype=fasta&retmode=text"
)
SEQUENCE_SHA256 = (
    "e999dceccb76562d40e49a10d9e6d4cf1c05d184f4f80b92bb32d4a2420c11d0"
)
SEED = 2026
BIT_WINDOW = 2048
BASE_WINDOW = 3000
BASE_SHUFFLES = 8
REPEAT_SHUFFLES = 3
REPEAT_CHUNK = 96
REPEAT_REGION = 96000
PAIR_BLOCK = 4096
PSEUDOCOUNT = 0.5
HISTORY_LENGTHS = range(17)


def read_genome(path, download):
    if download and not path.exists():
        request = Request(
            FASTA_URL, headers={"User-Agent": "MG1655-reproduction/1.0"}
        )
        with urlopen(request, timeout=60) as response:
            fasta = response.read()
        path.parent.mkdir(parents=True, exist_ok=True)
        path.write_bytes(fasta)
    with path.open(encoding="ascii") as fasta:
        header = fasta.readline().strip()
        sequence = "".join(line.strip() for line in fasta).upper()

    identifier = header[1:].split(maxsplit=1)[0] if len(header) > 1 else ""
    if not header.startswith(">") or ACCESSION not in identifier.split("|"):
        raise ValueError(f"Unexpected FASTA header: {header}")
    if hashlib.sha256(sequence.encode("ascii")).hexdigest() != SEQUENCE_SHA256:
        raise ValueError("Genome sequence does not match the reference SHA-256")
    alphabet = np.zeros(128, dtype=np.uint8)
    alphabet[np.frombuffer(b"ACGT", dtype=np.uint8)] = np.arange(4)
    return alphabet[np.frombuffer(sequence.encode("ascii"), dtype=np.uint8)]


def encode(genome):
    return np.stack((genome >> 1, genome & 1), axis=-1).reshape(-1)
```

```
def spectral_windows(bits, periods):
    phases = np.exp(-2j * np.pi * np.arange(bits.shape[1])[:, None] / periods)
    centered = bits - bits.mean(axis=1, keepdims=True)
    return np.square(np.abs(centered @ phases)).mean(axis=0)


def bit_spectrum(bits):
    windows = bits[: len(bits) // BIT_WINDOW * BIT_WINDOW].reshape(-1, BIT_WINDOW)
    first, second = np.split(windows, [len(windows) // 2])
    periods = np.arange(2, 129)
    rng = np.random.default_rng(SEED + 1)

    def ratios(group):
        shuffled = group.copy()
        for offset in range(2):
            shuffled[:, offset::2] = rng.permuted(group[:, offset::2], axis=1)
        observed = spectral_windows(group, periods)
        return observed / spectral_windows(shuffled, periods)

    discovery, test = ratios(first), ratios(second)
    top = np.argsort(-discovery, kind="stable")[:5]
    print(f"Blind spectrum, {len(first)} discovery and {len(second)} test windows")
    for index in top:
        print(f"  period {periods[index]:3d}: discovery {discovery[index]:.4f}x, "
              f"held out {test[index]:.4f}x")
    print(f"  period   2: held out {test[0]:.4f}x")


def figure_spectrum(genome, figure_data):
    usable = len(genome) // BASE_WINDOW * BASE_WINDOW
    windows = genome[:usable].reshape(-1, BASE_WINDOW)

    def mean_fft(group):
        bit_windows = encode(group).reshape(-1, 2 * BASE_WINDOW)
        return np.square(np.abs(np.fft.rfft(bit_windows, axis=1))).mean(axis=0)

    rng = np.random.default_rng(SEED)
    real_power = mean_fft(windows)
    null_power = np.mean(
        [mean_fft(rng.permuted(windows, axis=1)) for _ in range(BASE_SHUFFLES)],
        axis=0,
    )
    print(f"Base-shuffle spectrum, {len(windows)} windows of {BASE_WINDOW} bases")
    for period in (6, 3, 2):
        index = 2 * BASE_WINDOW // period
        print(f"  period {period}: {real_power[index] / null_power[index]:.5f}x")

    if figure_data is not None:
        figure_data.mkdir(parents=True, exist_ok=True)
        frequencies = np.fft.rfftfreq(2 * BASE_WINDOW)
        np.savetxt(
            figure_data / "spectrum.csv",
            np.column_stack((frequencies, real_power, null_power)),
            delimiter=",",
            header="cycles_per_bit,reference_power,shuffled_power",
            comments="",
        )

        def detrended_skew(bases):
            skew = np.cumsum((bases == 2).astype(np.int64) - (bases == 1))
            return skew - np.arange(1, len(bases) + 1) * (skew[-1] / len(bases))
```

```
        actual = detrended_skew(genome)
        skew_rng = np.random.default_rng(SEED + 2)
        shuffled = detrended_skew(skew_rng.permutation(genome))
        keep = np.r_[np.arange(0, len(genome) - 1, 1000), len(genome) - 1]
        np.savetxt(
            figure_data / "strand_skew.csv",
            np.column_stack((keep + 1, actual[keep], shuffled[keep])),
            delimiter=",",
            header="base,detrended_G_minus_C,shuffled_G_minus_C",
            comments="",
        )


def smoothed_log_loss(bits, length, train_end, test_start, test_end):
    context = np.zeros(len(bits) - length, dtype=np.int32)
    for offset in range(length):
        context = (context << 1) + bits[offset : len(bits) - length + offset]
    targets = bits[length:]
    train_slice = slice(0, train_end - length)
    test_slice = slice(test_start - length, test_end - length)
    joint_counts = np.bincount(
        2 * context[train_slice] + targets[train_slice],
        minlength=2 ** (length + 1),
    ).reshape(-1, 2)
    test_context = context[test_slice]
    test_targets = targets[test_slice]
    probability = (joint_counts[test_context, test_targets] + PSEUDOCOUNT) / (
        joint_counts[test_context].sum(axis=1) + 2 * PSEUDOCOUNT
    )
    return -np.log2(probability).mean()


def test_prediction(bits):
    middle, final_quarter = len(bits) // 2, 3 * len(bits) // 4
    scores = [
        smoothed_log_loss(bits, length, middle, middle, final_quarter)
        for length in HISTORY_LENGTHS
    ]
    chosen = int(np.argmin(scores))
    test_loss = smoothed_log_loss(
        bits, chosen, final_quarter, final_quarter, len(bits)
    )
    baseline = smoothed_log_loss(bits, 0, final_quarter, final_quarter, len(bits))
    saved = (baseline - test_loss) * (len(bits) - final_quarter)
    print(f"Prediction: k={chosen}; {test_loss:.8f} vs {baseline:.8f} bits/bit; "
          f"{saved:.1f} ideal bits saved before model cost")


def word_view(bits, width, offset):
    rows = (len(bits) - offset) // (2 * width)
    records = bits[offset : offset + rows * 2 * width].reshape(rows, 2, width)
    weights = 2 ** np.arange(width - 1, -1, -1)
    words = records @ weights
    return words[:, 0], words[:, 1]


def pair_tables(inputs, outputs, count):
    observed = np.bincount(count * inputs + outputs, minlength=count**2)
    observed = observed.reshape(count, count)
    expected = np.zeros_like(observed, dtype=np.float64)
    for start in range(0, len(inputs), PAIR_BLOCK):
        left = inputs[start : start + PAIR_BLOCK]
```

```
        right = outputs[start : start + PAIR_BLOCK]
        in_counts = np.bincount(left, minlength=count)
        out_counts = np.bincount(right, minlength=count)
        expected += np.outer(in_counts, out_counts) / len(left)
    return observed, expected


def test_word_maps(bits):
    for width in (2, 3, 4):
        count = 2**width
        row = np.arange(count)
        best = None
        for offset in range(2 * width):
            inputs, outputs = word_view(bits, width, offset)
            middle = len(inputs) // 2
            train, expected = pair_tables(inputs[:middle], outputs[:middle], count)
            excess = (train - expected) / middle
            _, mapping = linear_sum_assignment(excess, maximize=True)
            score = excess[row, mapping].sum()
            if best is None or score > best[0]:
                best = (
                    score, offset, mapping, train,
                    inputs[middle:], outputs[middle:],
                )

        _, offset, mapping, train, inputs, outputs = best
        observed, expected = pair_tables(inputs, outputs, count)
        accuracy = observed[row, mapping].sum() / len(inputs)
        null_accuracy = expected[row, mapping].sum() / len(inputs)
        identity = observed.trace() / len(inputs)
        print(f"Width {width}, offset {offset}, map {mapping.tolist()}: "
              f"{100 * accuracy:.5f}% vs {100 * null_accuracy:.5f}% "
              f"local independence; identity {100 * identity:.5f}%")

        if width == 2:
            marginal = (train.sum(axis=0) + PSEUDOCOUNT) / (
                train.sum() + count * PSEUDOCOUNT
            )
            full = (train + PSEUDOCOUNT) / (
                train.sum(axis=1, keepdims=True) + count * PSEUDOCOUNT
            )
            prediction = np.eye(count)[mapping]

            def mixed_loss(weight, counts):
                model = (1 - weight) * marginal + weight * prediction
                return -(counts * np.log2(model)).sum() / counts.sum()

            fit = minimize_scalar(
                mixed_loss, bounds=(0.0, 1.0 - 1e-12), args=(train,),
                method="bounded",
            )
            conditional_loss = -(observed * np.log2(full)).sum() / len(inputs)
            print("  held-out bits/output: "
                  f"marginal {mixed_loss(0, observed):.6f}, "
                  f"map mixture {mixed_loss(fit.x, observed):.6f}, "
                  f"conditional {conditional_loss:.6f}")
            if offset % 2 == 0:
                bases = bits[:len(bits) // 2].reshape(-1, 2) @ np.array([2, 1])
                adjacent = np.bincount(4 * bases[:-1] + bases[1:], minlength=16)
                adjacent = adjacent.reshape(4, 4)
                agreement = adjacent[row, mapping].sum() / adjacent.sum()
                print("  discovery dinucleotide agreement "
```

```
                    "across all adjacent bases: "
                    f"{100 * agreement:.5f}%")


def word64_seeds(bits, offset):
    length = (len(bits) - offset) // 8 * 8
    packed = np.packbits(bits[offset : offset + length])
    windows = np.lib.stride_tricks.sliding_window_view(packed, 8)
    return windows.view(">u8").reshape(-1)


def matching_prefix(first, second):
    step = 8192
    for start in range(0, len(first), step):
        end = start + step
        errors = np.flatnonzero(first[start:end] != second[start:end])
        if errors.size:
            return start + int(errors[0])
    return len(first)


def extended_repeat(bits, a, b, seed_length=64):
    left = min(a, b)
    right = len(bits) - max(a, b) - seed_length
    before = matching_prefix(bits[a - left : a][::-1], bits[b - left : b][::-1])
    after = matching_prefix(
        bits[a + seed_length : a + seed_length + right],
        bits[b + seed_length : b + seed_length + right],
    )
    return a - before, b - before, before + seed_length + after


def longest_seeded_repeat(bits):
    anchors = word64_seeds(bits, 0)[::8]
    order = np.argsort(anchors)
    sorted_anchors = anchors[order]
    longest = (0, 0, 0)
    checked = {}
    for offset in range(8):
        words = word64_seeds(bits, offset)
        starts = np.searchsorted(sorted_anchors, words, side="left")
        ends = np.searchsorted(sorted_anchors, words, side="right")
        for index in np.flatnonzero(ends > starts):
            seed = int(index) * 8 + offset
            for anchor in order[starts[index] : ends[index]]:
                anchor = int(anchor) * 64
                if anchor == seed:
                    continue
                a, b = sorted((anchor, seed))
                distance = b - a
                intervals = checked.setdefault(distance, [])
                if any(start <= a < end for start, end in intervals):
                    continue
                start_a, start_b, length = extended_repeat(bits, a, b)
                intervals.append((start_a, start_a + length - 63))
                if length > longest[2]:
                    longest = (start_a, start_b, length)
    return longest


def shuffled_chunks(bits, rng):
    shuffled = bits.copy()
```

```
    for start in range(0, len(bits), REPEAT_REGION):
        window = shuffled[start : start + REPEAT_REGION]
        length = len(window) // REPEAT_CHUNK * REPEAT_CHUNK
        chunks = window[:length].reshape(-1, REPEAT_CHUNK)
        chunks[:] = chunks[rng.permutation(len(chunks))]
    return shuffled


def test_repeats(bits):
    a, b, length = longest_seeded_repeat(bits)
    print(f"Longest seeded repeat: {length} bits at zero-based offsets "
          f"{a} and {b}")
    if a >= 256 and b + length + 256 <= len(bits):
        before = np.count_nonzero(bits[a - 256 : a] == bits[b - 256 : b])
        after = np.count_nonzero(bits[a + length : a + length + 256]
                                 == bits[b + length : b + length + 256])
        print(f"  flanking matches: {before}/256 before, {after}/256 after")
    rng = np.random.default_rng(SEED + 3)
    lengths = [longest_seeded_repeat(shuffled_chunks(bits, rng))[2]
               for _ in range(REPEAT_SHUFFLES)]
    print(f"  maximum seeded repeat in each chunk shuffle: {lengths}")


def main():
    parser = argparse.ArgumentParser(description=__doc__)
    parser.add_argument(
        "--fasta", type=Path, default=Path("data") / f"{ACCESSION}.fasta"
    )
    parser.add_argument(
        "--download", action="store_true", help="retrieve FASTA if missing"
    )
    parser.add_argument(
        "--figure-data", type=Path, help="directory for optional plot CSVs"
    )
    args = parser.parse_args()
    if not args.fasta.exists() and not args.download:
        parser.error(f"{args.fasta} is missing. Use --download or --fasta.")

    genome = read_genome(args.fasta, args.download)
    bits = encode(genome)
    print(f"{ACCESSION}: {len(genome):,} bases, {len(bits):,} bits")
    print(f"Normalized sequence SHA-256 verified: {SEQUENCE_SHA256}")
    figure_spectrum(genome, args.figure_data)
    bit_spectrum(bits)
    test_prediction(bits)
    test_word_maps(bits)
    test_repeats(bits)


if __name__ == "__main__":
    main()
```

hydropathy_test.py

```
#!/usr/bin/env python3
"""Enumerate the fixed-stop, equal-degeneracy hydropathy null in Chapter 12."""

from itertools import permutations, product
from math import log2

import numpy as np


# T/C/A/G order on each of the three positions, with the third varying fastest.
CODONS = list(product("TCAG", repeat=3))
ASSIGNMENTS = (
    "FFLLSSSSYY**CC*W"
    "LLLLPPPPHHQQRRRR"
    "IIIMTTTTNNKKSSRR"
    "VVVVAAAADDEEGGGG"
)
# Kyte–Doolittle hydropathy in integer tenths. Stops have a fixed placeholder.
HYDROPATHY = {
    "A": 18, "R": -45, "N": -35, "D": -35, "C": 25,
    "Q": -35, "E": -35, "G": -4, "H": -32, "I": 45,
    "L": 38, "K": -39, "M": 19, "F": 28, "P": -16,
    "S": -8, "T": -7, "W": -9, "Y": -13, "V": 42,
    "*": 0,
}


def permutation_sums(degeneracy, assignments, features):
    names = sorted(name for name in HYDROPATHY
                   if name != "*" and ASSIGNMENTS.count(name) == degeneracy)
    codon_counts = np.array(
        [features[assignments == name].sum(axis=0) for name in names]
    )
    value_orders = sorted(set(permutations(HYDROPATHY[name] for name in names)))
    return np.array(value_orders, dtype=np.int32) @ codon_counts


def main():
    assignments = np.array(list(ASSIGNMENTS))
    values = np.array([HYDROPATHY[name] for name in ASSIGNMENTS])
    features = np.array(
        [[int(codon[position] == base) for position in range(3) for base in "TCAG"]
         for codon in CODONS],
        dtype=np.int32
    )
    marginal_sums = values @ features
    score = int(np.square(marginal_sums).sum())
    mean_term = int(values.sum()) ** 2 / len(values)
    total_variance = int(np.square(values).sum()) - mean_term
    r_squared = (score / 16 - 3 * mean_term) / total_variance
    second_variance = np.square(marginal_sums[4:8]).sum() / 16 - mean_term
    second_position = second_variance / total_variance

    one, two, three, four, six = [
        permutation_sums(degeneracy, assignments, features)
        for degeneracy in (1, 2, 3, 4, 6)
    ]
    # Combine the smaller groups once, then stream the larger group in batches.
    smaller = (one[:, None, None, :] + four[None, :, None, :]
```

```
            + six[None, None, :, :]).reshape(-1, 12) + three
    count = at_least = tied = 0
    for start in range(0, len(two), 100):
        sums = two[start : start + 100, None, :] + smaller[None, :, :]
        scores = np.square(sums).sum(axis=2)
        at_least += np.count_nonzero(scores >= score)
        tied += np.count_nonzero(scores == score)
        count += scores.size

    fraction = at_least / count
    print(f"R² = {r_squared:.10f}, "
          f"second-position share = {second_position:.10f}")
    print(f"{at_least:,} / {count:,} scores at least as high ({tied} equal)")
    print(f"Tail fraction = {fraction:.11g}, "
          f"-log₂(tail fraction) = {-log2(fraction):.8f}")


if __name__ == "__main__":
    main()
```

plot_sequence.py

```
#!/usr/bin/env python3
"""Render the book's grayscale Figure 4 from the sequence-test CSVs."""

import argparse
from io import BytesIO
from pathlib import Path

import matplotlib
matplotlib.use("Agg")
import matplotlib.pyplot as plt
import numpy as np
from PIL import Image


def main():
    parser = argparse.ArgumentParser(description=__doc__)
    parser.add_argument("data", type=Path, help="directory of sequence-test CSVs")
    parser.add_argument("output", type=Path, help="output PNG path")
    args = parser.parse_args()

    frequency, observed, shuffled = np.loadtxt(
        args.data / "spectrum.csv", delimiter=",", skiprows=1, unpack=True
    )
    position, reference_skew, shuffled_skew = np.loadtxt(
        args.data / "strand_skew.csv", delimiter=",", skiprows=1, unpack=True
    )

    plt.rcParams.update({
        "font.family": "DejaVu Sans",
        "font.size": 10,
        "axes.labelsize": 10,
        "axes.linewidth": 0.85,
        "xtick.labelsize": 9.4,
        "ytick.labelsize": 9.4,
        "legend.fontsize": 9.2,
        "pdf.fonttype": 42,
    })
    fig, (spectrum, skew) = plt.subplots(
        2, 1, figsize=(6, 4.15), facecolor="white"
    )
    fig.subplots_adjust(
        left=0.145, right=0.97, bottom=0.115, top=0.97, hspace=0.66
    )
    # The unnormalized FFT power is divided by the 6,000 bits in each window.
    reference_power = observed / 6000
    control_power = shuffled / 6000
    spectrum.plot(frequency[1:], reference_power[1:], color="black", linewidth=1.0,
                  label="Reference")
    spectrum.plot(frequency[1:], control_power[1:], color="0.34", linewidth=1.0,
                  linestyle="--", label="Shuffled bases")
    for numerator, fraction in ((1, 6), (1, 3), (1, 2)):
        target = numerator / fraction
        index = int(np.argmin(np.abs(frequency - target)))
        y = reference_power[index]
        spectrum.text(target, y + 0.11, f"{numerator}/{fraction}", ha="center",
                      va="bottom", fontweight="semibold")
    spectrum.set(
        xlim=(0, 0.517), ylim=(0, 2.21),
        xticks=[0, 0.1, 0.2, 0.3, 0.4, 0.5],
```

```
        yticks=[0, 0.5, 1, 1.5, 2],
        xlabel="Frequency (cycles per bit)", ylabel="Mean power per bit",
    )
    spectrum.legend(frameon=False, loc="upper right", bbox_to_anchor=(0.99, 0.99),
                    borderpad=0.1, handlelength=2)

    skew.axhline(0, color="0.57", linewidth=0.8)
    skew.plot(position / 1e6, reference_skew / 1e3, color="black", linewidth=1.15,
              label="Reference")
    skew.plot(position / 1e6, shuffled_skew / 1e3, color="0.34", linewidth=1.0,
              linestyle="--", label="Shuffled bases")
    skew.set(xlim=(0, 4.68), ylim=(-14, 29), xticks=[0, 1, 2, 3, 4],
             yticks=[-10, 0, 10, 20], xlabel="Genome position (million bases)",
             ylabel="Detrended cumulative\nG − C (thousands)")
    skew.legend(frameon=False, loc="upper right", bbox_to_anchor=(0.99, 0.99),
                borderpad=0.1, handlelength=2)

    for letter, panel in zip("AB", (spectrum, skew)):
        panel.text(0.025, 0.96, letter, transform=panel.transAxes, va="top",
                   ha="left", fontsize=11, fontweight="bold")
        panel.spines[["top", "right"]].set_visible(False)
        panel.tick_params(direction="out", length=3, width=0.85)

    args.output.parent.mkdir(parents=True, exist_ok=True)
    with BytesIO() as raster:
        fig.savefig(
            raster, format="png", dpi=450, facecolor="white", transparent=False
        )
        raster.seek(0)
        with Image.open(raster) as image:
            image.convert("L").save(args.output, dpi=(450, 450), optimize=True)


if __name__ == "__main__":
    main()
```

References

Akeson, Rachel, et al. 2021. “Precision Millimeter Astrometry of the α Centauri AB System.” The Astronomical Journal 162, 14. https://doi.org/10.3847/1538-3881/abfaff

Allen, Natalie H., et al. 2026. “JWST TRAPPIST-1 e/b Program: Motivation and First Observations.” *The Astronomical Journal* 171, 105. https://doi.org/10.3847/1538-3881/ae28cb

Anglada-Escudé, Guillem, et al. 2016. “A Terrestrial Planet Candidate in a Temperate Orbit Around Proxima Centauri.” *Nature* 536, 437–440. https://doi.org/10.1038/nature19106

Basant, Ritvik, et al. 2025. “Four Sub-Earth Planets Orbiting Barnard’s Star from MAROON-X and ESPRESSO.” The Astrophysical Journal Letters. https://doi.org/10.3847/2041-8213/adb8d5

Beichman, Charles, et al. 2025. “Worlds Next Door: A Candidate Giant Planet Imaged in the Habitable Zone of α Centauri A. I. Observations, Orbital and Physical Properties, and Exozodi Upper Limits.” *The Astrophysical Journal Letters* 989, L22. https://doi.org/10.3847/2041-8213/adf53f

Beisson, Janine, and Tracy M. Sonneborn. 1965. “Cytoplasmic Inheritance of the Organization of the Cell Cortex in *Paramecium aurelia*.” *Proceedings of the National Academy of Sciences* 53, 275–282. https://doi.org/10.1073/pnas.53.2.275

Bennett, Bryson D., et al. 2009. “Absolute Metabolite Concentrations and Implied Enzyme Active Site Occupancy in *Escherichia coli*.” *Nature Chemical Biology* 5, 593–599. https://doi.org/10.1038/nchembio.186

Biotic. 2026. “SpudCell.” Research overview. https://biotic.org/spudcell/

Boles, Kent S., et al. 2017. “Digital-to-Biological Converter for On-Demand Production of Biologics.” *Nature Biotechnology* 35, 672–675. https://doi.org/10.1038/nbt.3859

Booth, Mark, et al. 2023. “The Clumpy Structure of ϵ Eridani’s Debris Disc Revisited by ALMA.” *Monthly Notices of the Royal Astronomical Society* 521, 6180–6194. https://doi.org/10.1093/mnras/stad938

Chen, Cheng, Joshua S. MacCready, Daniel C. Ducat, and Katherine W. Osteryoung. 2018. “The Molecular Machinery of Chloroplast Division.” Plant Physiology 176, 138–151. https://doi.org/10.1104/pp.17.01272

Chen, Stephen S., et al. 2012. “Measuring the Dynamics of E. coli Ribosome Biogenesis Using Pulse-Labeling and Quantitative Mass Spectrometry.” Molecular BioSystems 8, 3325–3334. https://doi.org/10.1039/c2mb25310k

Clarke, Arthur C. 1986. *The Songs of Distant Earth*. https://www.penguinrandomhouse.com/books/28183/songs-of-distant-earth-by-arthur-c-clarke/

Consultative Committee for Space Data Systems (CCSDS). 2024. Reference Model for an Open Archival Information System (OAIS). CCSDS 650.0-M-3. https://ccsds.org/Pubs/650x0m3.pdf

Crick, F. H. C., and L. E. Orgel. 1973. “Directed Panspermia.” *Icarus* 19, 341–346. https://doi.org/10.1016/0019-1035(73)90110-3

Falkowski, Paul G., Tom Fenchel, and Edward F. DeLong. 2008. “The Microbial Engines That Drive Earth’s Biogeochemical Cycles.” *Science* 320, 1034–1039. https://doi.org/10.1126/science.1153213

Gaskill, Melissa L. 2023. “NASA Achieves Water Recovery Milestone on International Space Station.” National Aeronautics and Space Administration, June 20. https://www.nasa.gov/missions/station/nasa-achieves-water-recovery-milestone-on-international-space-station/

Gaut, Nathaniel J., Christopher Deich, Brock Cash, Tanner Hoog, Aaron E. Engelhart, and Katarzyna P. Adamala. 2026. “A Chemically Defined Synthetic Cell Capable of Growth and Replication.” *bioRxiv*, July. Preprint. https://doi.org/10.64898/2026.07.01.735724

Gibson, Daniel G., et al. 2010. “Creation of a Bacterial Cell Controlled by a Chemically Synthesized Genome.” *Science* 329, 52–56. https://doi.org/10.1126/science.1190719

Gillon, Michaël, et al. 2017. “Seven Temperate Terrestrial Planets Around the Nearby Ultracool Dwarf Star TRAPPIST-1.” *Nature* 542, 456–460. https://doi.org/10.1038/nature21360

Glidden, Ana, et al. 2025. “JWST-TST DREAMS: Secondary Atmosphere Constraints for the Habitable Zone Planet TRAPPIST-1 e.” *The Astrophysical Journal Letters* 990, L53. https://doi.org/10.3847/2041-8213/adf62e

Gros, Claudius. 2016. “Developing Ecospheres on Transiently Habitable Planets: The Genesis Project.” *Astrophysics and Space Science* 361, 324. https://doi.org/10.1007/s10509-016-2911-0

Gros, Claudius. 2017. “Universal Scaling Relation for Magnetic Sails: Momentum Braking in the Limit of Dilute Interstellar Media.” *Journal of Physics Communications* 1, 045007. https://doi.org/10.1088/2399-6528/aa927e

Gros, Claudius. 2019. “Why Planetary and Exoplanetary Protection Differ: The Case of Long Duration Genesis Missions to Habitable but Sterile M-Dwarf Oxygen Planets.” *Acta Astronautica* 157, 263–267. https://doi.org/10.1016/j.actaastro.2019.01.005

Harold, Franklin M. 2005. “Molecules into Cells: Specifying Spatial Architecture.” *Microbiology and Molecular Biology Reviews* 69, 544–564. https://doi.org/10.1128/MMBR.69.4.544-564.2005

Hassler, Donald M., et al. 2014. “Mars’ Surface Radiation Environment Measured with the Mars Science Laboratory’s Curiosity Rover.” *Science* 343, 1244797. https://doi.org/10.1126/science.1244797

Hatton, Ian A., et al. 2023. “The Human Cell Count and Size Distribution.” *Proceedings of the National Academy of Sciences* 120, e2303077120. https://doi.org/10.1073/pnas.2303077120

Hein, Andreas M., and Stephen Baxter. 2019. “Artificial Intelligence for Interstellar Travel.” *Journal of the British Interplanetary Society* 72, 125–143. https://www.bis-space.com/membership/jbis/2019/JBIS-v72-no04-April-2019-d94kwe.pdf

Heller, René, and Michael Hippke. 2017. “Deceleration of High-Velocity Interstellar Photon Sails into Bound Orbits at α Centauri.” *The Astrophysical Journal Letters* 835, L32. https://doi.org/10.3847/2041-8213/835/2/L32

Hoang, Thiem, Alex Lazarian, Blakesley Burkhart, and Abraham Loeb. 2017. “The Interaction of Relativistic Spacecrafts with the Interstellar Medium.” *The Astrophysical Journal* 837, 5. https://doi.org/10.3847/1538-4357/aa5da6

Hogan, James P. 1982. *Voyage from Yesteryear*. https://www.baen.com/voyage-from-yesteryear.html

Horneck, Gerda, H. Bücker, and Günther Reitz. 1994. “Long-Term Survival of Bacterial Spores in Space.” *Advances in Space Research* 14(10), 41–45. https://doi.org/10.1016/0273-1177(94)90448-0

Jones, Harry W. 2016. “Design and Analysis of a Flexible, Reliable Deep Space Life Support System.” NASA Ames Research Center, Table 1. https://ntrs.nasa.gov/citations/20160005782

Jones, Harry W. 2025. “Common Cause Failures Dominate and Defeat Redundancy.” Presentation at the 71st Annual Reliability and Maintainability Symposium. NASA Technical Reports Server, 20240013667. https://ntrs.nasa.gov/citations/20240013667

Kervella, Pierre, Frédéric Thévenin, and Christophe Lovis. 2017. “Proxima’s Orbit Around α Centauri.” *Astronomy & Astrophysics* 598, L7. https://doi.org/10.1051/0004-6361/201629930

Kopp, Greg, and Judith L. Lean. 2011. “A New, Lower Value of Total Solar Irradiance: Evidence and Climate Significance.” *Geophysical Research Letters* 38, L01706. https://doi.org/10.1029/2010GL045777

Kusuma, Paul, P. Morgan Pattison, and Bruce Bugbee. 2020. “From Physics to Fixtures to Food: Current and Potential LED Efficacy.” *Horticulture Research* 7, article 56. https://doi.org/10.1038/s41438-020-0283-7

Kyte, Jack, and Russell F. Doolittle. 1982. “A Simple Method for Displaying the Hydropathic Character of a Protein.” *Journal of Molecular Biology* 157, 105–132. https://doi.org/10.1016/0022-2836(82)90515-0

Lewin, Harris A., et al. 2018. “Earth BioGenome Project: Sequencing Life for the Future of Life.” *Proceedings of the National Academy of Sciences* 115 (17), 4325–4333. https://doi.org/10.1073/pnas.1720115115

MacGregor, Meredith A., et al. 2021. “Discovery of an Extremely Short Duration Flare from Proxima Centauri Using Millimeter through Far-Ultraviolet Observations.” *The Astrophysical Journal Letters* 911, L25. https://doi.org/10.3847/2041-8213/abf14c

Marino, Sebastián, et al. 2017. “ALMA Observations of the Multiplanet System 61 Vir: What Lies Outside Super-Earth Systems?” *Monthly Notices of the Royal Astronomical Society* 469, 3518–3531. https://doi.org/10.1093/mnras/stx1102

Mautner, Michael Noah. 1997. “Directed Panspermia. 3. Strategies and Motivations for Seeding Star-Forming Clouds.” *Journal of the British Interplanetary Society* 50, 93–102. https://scholarscompass.vcu.edu/chem_pubs/12/

Mautner, Michael Noah. 2000. *Seeding the Universe with Life: Securing Our Cosmological Future*. Legacy Books.

Moeller, Ralf, et al. 2014. “Resistance of *Bacillus subtilis* Spore DNA to Lethal Ionizing Radiation Damage Relies Primarily on Spore Core Components and DNA Repair, with Minor Effects of Oxygen Radical Detoxification.” *Applied and Environmental Microbiology* 80, 104–109. https://doi.org/10.1128/AEM.03136-13

NASA Exoplanet Archive. n.d. System overviews for Barnard’s Star, Epsilon Eridani, 61 Virginis, and TRAPPIST-1. https://exoplanetarchive.ipac.caltech.edu/overview/Barnards%20star https://exoplanetarchive.ipac.caltech.edu/overview/eps%20Eri https://exoplanetarchive.ipac.caltech.edu/overview/61%20Vir https://exoplanetarchive.ipac.caltech.edu/overview/TRAPPIST-1

NASA Goddard Space Flight Center. n.d. “The Nearest Neighbor Star.” *Imagine the Universe!* https://imagine.gsfc.nasa.gov/features/cosmic/nearest_star_info.html

NASA Jet Propulsion Laboratory. 2023. “NASA’s Oxygen-Generating Experiment MOXIE Completes Mars Mission.” September 6. https://www.jpl.nasa.gov/news/nasas-oxygen-generating-experiment-moxie-completes-mars-mission/

NASA Jet Propulsion Laboratory. 2024. “NASA’s Optical Comms Demo Transmits Data Over 140 Million Miles.” April 25. https://www.jpl.nasa.gov/news/nasas-optical-comms-demo-transmits-data-over-140-million-miles/

NASA Jet Propulsion Laboratory. n.d. “Deep Space Optical Communications (DSOC).” https://www.jpl.nasa.gov/missions/deep-space-optical-communications-dsoc/

NASA Science. 2019. “Proxima b 3D Model.” https://science.nasa.gov/resource/proxima-b-3d-model/

NASA Science. n.d. “Ladder of Life Detection.” https://science.nasa.gov/astrobiology/researchers/life-detection-resources/ladder-of-life-detection/

NASA. n.d. “Dawn Ion Propulsion.” https://science.nasa.gov/mission/dawn/technology/ion-propulsion/

National Center for Biotechnology Information. *Escherichia coli* Strain K-12 Substrain MG1655, Complete Genome. GenBank U00096.3. https://www.ncbi.nlm.nih.gov/nucleotide/U00096.3

National Center for Biotechnology Information. “The Genetic Codes.” Translation table 11. https://www.ncbi.nlm.nih.gov/Taxonomy/Utils/wprintgc.cgi#SG11

Nurk, Sergey, et al. 2022. “The Complete Sequence of a Human Genome.” *Science* 376, 44–53. https://doi.org/10.1126/science.abj6987

Ribas, Ignasi, et al. 2017. “The Full Spectral Radiative Properties of Proxima Centauri.” *Astronomy & Astrophysics* 603, A58. https://doi.org/10.1051/0004-6361/201730763

Salmon, Sébastien J. A. J., et al. 2021. “Reinvestigating α Centauri AB in Light of Asteroseismic Forward and Inverse Methods.” Astronomy & Astrophysics 646, A7. https://doi.org/10.1051/0004-6361/201937174

Schmidt, Alexander, et al. 2016. “The Quantitative and Condition-Dependent *Escherichia coli* Proteome.” *Nature Biotechnology* 34, 104–110. https://doi.org/10.1038/nbt.3418

Severinghaus, Jeffrey P., Wallace S. Broecker, William F. Dempster, Taber MacCallum, and Martin Wahlen. 1994. “Oxygen Loss in Biosphere 2.” *Eos, Transactions American Geophysical Union* 75, 33–37. https://doi.org/10.1029/94EO00285

SIMBAD Astronomical Database. n.d. “Barnard’s Star.” Centre de Données astronomiques de Strasbourg. https://simbad.harvard.edu/simbad/sim-id?Ident=Barnard%27s+star

Sleator, Roy D., and Niall Smith. 2017. “Directed Panspermia: A 21st Century Perspective.” *Science Progress* 100, 187–193. https://doi.org/10.3184/003685017X14901006155062

Tilman, David, Peter B. Reich, and Johannes M. H. Knops. 2006. “Biodiversity and Ecosystem Stability in a Decade-Long Grassland Experiment.” *Nature* 441, 629–632. https://doi.org/10.1038/nature04742

Tiwari, Sharmila, et al. 1997. “Prediction of Probable Genes by Fourier Analysis of Genomic Sequences.” *Computer Applications in the Biosciences* 13, 263–270. https://doi.org/10.1093/bioinformatics/13.3.263

Tyser, Richard C. V., et al. 2021. “Single-Cell Transcriptomic Characterization of a Gastrulating Human Embryo.” *Nature* 600, 285–289. https://doi.org/10.1038/s41586-021-04158-y

United Nations. 1967. Treaty on Principles Governing the Activities of States in the Exploration and Use of Outer Space, Including the Moon and Other Celestial Bodies. Article IX. https://treaties.un.org/pages/showdetails.aspx?objid=0800000280128cbd

Vinge, Vernor. 1972. “Long Shot.” *Analog Science Fiction/Science Fact*, August. https://sfadb.com/db/Vernor_Vinge_Titles

Wang, ZiMian, et al. 2003. “Total Body Protein: A New Cellular Level Mass and Distribution Prediction Model.” *American Journal of Clinical Nutrition* 78, 979–984. https://doi.org/10.1093/ajcn/78.5.979

Wheeler, Raymond M. 2020. “NASA's Contributions to Vertical Farming.” NASA Technical Memorandum 2020-5008832. https://ntrs.nasa.gov/citations/20205008832

Xue, Zhigang, et al. 2013. “Genetic Programs in Human and Mouse Early Embryos Revealed by Single-Cell RNA Sequencing.” *Nature* 500, 593–597. https://doi.org/10.1038/nature12364

Yang, Mary M., William J. Coleman, and Douglas C. Youvan. 1990. “Genetic Coding Algorithms for Engineering Membrane Proteins.” In *Reaction Centers of Photosynthetic Bacteria*, edited by Maria-Elisabeth Michel-Beyerle, 209–218. Springer. https://doi.org/10.1007/978-3-642-61297-8_21

Zeitlin, C., et al. 2013. “Measurements of Energetic Particle Radiation in Transit to Mars on the Mars Science Laboratory.” *Science* 340, 1080–1084. https://doi.org/10.1126/science.1235989