

Start-Up or Sustained Supply?

Tritium Fuel-Cycle Margins and Customer Commitments

Norio Yamazaki

Independent Researcher, Shizuoka, Japan

3 October 2026 | Revision 11

Abstract

I evaluate when a receiver's initial tritium requirement becomes a continuing external supply obligation. A serial aggregate fuel-cycle model gives a closed maintenance threshold in the inner-loop non-return fraction. For the declared parameters, including loop decay and a maintained 1.944 kg reserve lowers that threshold from 7.99×10^{-4} to 7.34×10^{-4} . A 10^{-3} non-return fraction requires at least 0.757 kg/year of continuing receipts in the continuous-maintenance limit; fixed-cadence delivery adds storage and transit decay. A controlled comparison holds non-return fixed when varying reserves. I convert the full receiver trajectory into dated, batched requests and evaluate them against stochastic supplier histories, capacity constraints, customer priorities and qualification holds. Counterfactual pairs share the same operating appointments. The fixed 41-date benchmark supports no admissible start under its 1% total missed-commitment criterion; this is a result for the specified supply system, not a deployment forecast. Source replenishment and receiver needs are evaluated over matching horizons, while rejected tritium remains in the full-site inventory. Separate single-order experiments show the trade-off between incumbent protection and new-customer delay. The contribution is a reproducible interface from a declared fuel-cycle margin to customer-resolved supply commitments within established inventory theory. It provides pre-feasibility consistency tests, rather than validated operating limits or a new general self-sufficiency criterion.

Keywords: tritium; fuel self-sufficiency; available-to-promise; reserve inventory; continuing supply; customer commitments.

AI-use disclosure. ChatGPT (OpenAI) assisted with literature checking, drafting, mathematical review, preparation and review of computational code, and document production. Results were obtained by executing the supplied programs. Scientific interpretation and submission responsibility remain with the human author.

1 Introduction

A fuel supplier needs to know whether a new receiver requires a finite startup inventory or continuing external support. The distinction depends on usable recovery, circulating flow, retained inventory and operating history. It also changes the obligations that must be protected for existing customers. A contract assessed against one year of receiver demand may appear feasible while a continued-operation contract from the same receiver is not.

I study this interface using established fuel-cycle and inventory models. Tritium accountancy has long distinguished physical material from hidden or estimated inventory [7]. Abdou et al. analyze the coupled effects of fuel processing, reserves, startup inventory and self-sufficiency [8]. Meschini et al. provide both detailed ARC/STEP-class models and a transparent aggregate transient [9]. On the service side, available-to-promise (ATP) determines quantities and dates after prior commitments [1, 2]; deteriorating-inventory and rationing theory address decay, lead times, priority classes and backorders [4, 5, 16]. These foundations already establish the principal mathematical mechanisms.

The contribution here is an integrated computational application. I derive an explicit maintenance threshold for a declared receiver model, distinguish initial reserve formation from ongoing imports, and pass the full resulting receipt schedule to a stochastic supplier. Within the same service boundary I compare current-stock funding with a lower-forecast reservation and chronological dispatch. Customer-specific outcomes show who bears a shortfall; extended horizons show which obligations were omitted by a short demand record. The study does not introduce a new general ATP algorithm.

The analytical threshold is an algebraic specialization of the component balance. Its contribution is the explicit connection between a receiver regime, its dated import sequence and customer-specific consequences at the supplier. Numerical verification supports that connection within the stated model; distinguishing the application from established inventory theory is essential to its interpretation.

The principal results address three questions. First, how do the receiver reserve and non-return fraction determine whether import demand persists? Second, what changes when a supplier credits conservative future receipts instead of either ignoring them or spending freely from current stock? Third, does a dated offer remain supported when initial filling, continuing imports, parcel cadence and the remaining contract horizon are considered together?

2 Evidence and input conventions

2.1 Published equations and declared extensions

I use the serial aggregate topology in Meschini et al., equations (26)–(30) and Figure 4 [9]. It has three differential inventories and two initial-condition statements. The outer-loop return enters the inner loop before reaching storage. The reference illustration uses a burn rate $b = 9 \times 10^{-7}$ kg/s, burn-to-feed ratio $\beta = 0.01$, generation ratio $\Lambda = 1.08$, inner residence time $\tau_I = 4$ h, outer residence time $\tau_O = 24$ h, and 1.5 kg initially in storage. Its loss-free equation replay remains a separate verification case.

The reserve expression in that paper’s equation (25) is

$$R_{\text{formula}} = F\chi t_{\text{res}}, \quad F = b/\beta. \quad (1)$$

Here χ is the fraction of the feed to be covered and t_{res} the specified reserve time. Combining the aggregate example’s feed, 7.776 kg/day, with the illustrative choices $\chi = 0.25$ and $t_{\text{res}} = 1$ day gives 1.944 kg. The full ARC baseline in the same paper uses a different burn rate and burn-to-feed ratio. The 1.944 kg case is therefore a transparent mixed-parameter scenario, not a quoted universal ARC requirement.

I compare this formula-based scenario with a synthetic 0.2 kg storage floor. The latter is not the reserve calculated by Eq. (1). With initial storage 1.5 kg, the 1.944 kg requirement already has a 0.444 kg gap

before loop filling. I maintain the receiver reserve in normal planned operation; emergency operation that deliberately consumes it would require a different operating policy. The supplier's withdrawal heel is a separate service allocation.

Conditional startup quantities in Parisi et al. [10] illustrate the sensitivity of required initial mass to fuel assumptions. Historical global supply scenarios [15] do not supply the inventory or interruption distribution used here. The evidence register identifies inspected primary sections, bibliographic records and synthetic inputs. The documentary cutoff is 3 October 2026. Direct verification of Eq. (1) was through Meschini et al.; the Abdou institutional abstract supports the broader fuel-cycle context.

2.2 Mass, time and model scope

Masses are kilograms of tritium, not gas-mixture mass. Receiver time is in days; service time is in years of 365.25 days. The selected half-life is 4500 days with standard uncertainty 8 days [13]; numerical examples use its central decay constant, $\lambda_d = \ln 2/4500 \text{ day}^{-1}$ and $\lambda_y = 0.0562604462 \text{ yr}^{-1}$. The selected value is an inherited evaluated convention, rather than a claim to have re-audited the latest nuclear-data libraries.

The aggregate feed is 7.776 kg/day, whereas burn is 0.07776 kg/day, or 28.40184 kg per declared year. A non-return fraction acts on the processed return, which is much larger than burn. The physical interpretation of that fraction must specify the branch and the period over which recovery is unavailable. The numerical extension below treats it as irreversible transfer to on-site rejected material during the modeled horizon. The ratio $\beta = b/F$ is defined at this aggregate feed boundary; it is distinct from a single-pass plasma burn fraction or a gas-purity specification. Whyte et al. derive burn-efficiency constraints from the vacuum-vessel boundary, helium exhaust and plasma performance [21]. A sensitivity that changes β at fixed burn is therefore conditional, rather than evidence that the plasma can realize those settings independently.

3 Receiver maintenance margin and continuing demand

3.1 Conserved serial-loop model

Let O, I, S, W denote outer-loop, inner-loop, usable storage and on-site rejected inventories. With outer usable recovery η_O , inner non-return fraction ϵ , and external receipt rate x , I use

$$\begin{aligned}\dot{O} &= g - O/\tau_O - \lambda_d O, \\ \dot{I} &= F - b + \eta_O O/\tau_O - I/\tau_I - \lambda_d I, \\ \dot{S} &= (1 - \epsilon)I/\tau_I - F - \lambda_d S + x, \\ \dot{W} &= (1 - \eta_O)O/\tau_O + \epsilon I/\tau_I - \lambda_d W.\end{aligned}\tag{2}$$

Summation gives $\dot{O} + \dot{I} + \dot{S} + \dot{W} = g - b + x - \lambda_d(O + I + S + W)$. Transfer to W removes usable fuel, not atoms from the site. In the receipt calculations, $x = 0$ between events and a discrete quantity x_k changes S by x_k at its arrival. Each receipt is counted once.

For the published loss-free reference, the independently implemented serial model has minimum storage 0.142528 kg at 2.78508 days. Routing the outer return directly to storage instead gives 0.155391 kg at 2.60289 days. Supplement S1 gives both sets of equations and their closed solutions. The published illustrative figure is not a quantitative validation trace for the chosen inputs; the comparison establishes equation-level reproduction.

3.2 An explicit threshold

For continuous operation, stationary loop inventories and $\eta_O = 1$, the no-decay storage input is

$$v_S = (g - b) - \epsilon(F - b + g). \quad (3)$$

Setting it to zero gives

$$\epsilon_0 = \frac{(\Lambda - 1)\beta}{1 - \beta + \Lambda\beta}. \quad (4)$$

This elementary balance explains the change of regime between small non-return fractions. It is not a full-plant self-sufficiency criterion.

Including residence-time decay and a maintained storage quantity R , define the stationary inner return before non-return,

$$J_I = \frac{F - b + \eta_O g / (1 + \lambda_d \tau_O)}{1 + \lambda_d \tau_I}. \quad (5)$$

The usable-storage maintenance threshold and necessary stationary import rate are

$$\epsilon_c(R) = 1 - \frac{F + \lambda_d R}{J_I}, \quad (6)$$

$$x_\infty(R, \epsilon) = [F + \lambda_d R - (1 - \epsilon)J_I]_+. \quad (7)$$

A negative ϵ_c means that even zero inner non-return cannot maintain that reserve under the stated other conditions. For $\epsilon < \epsilon_c$, the stationary storage balance at R is positive; it may still require startup imports. Equality provides no positive operating margin. Above the threshold, continuous operation requires ongoing imports or a change in the operating/design assumptions. Finite inventory can postpone that need but cannot remove a sustained negative balance.

The comparisons at $\epsilon/\epsilon_c(R) = 0.5$ and 1.25 are *regime-normalized* experiments: changing R also changes their absolute ϵ . They cannot isolate the causal effect of reserve alone. For a controlled reserve comparison I hold $b, \beta, \Lambda, \eta_O, \tau_I, \tau_O$ and ϵ fixed. The change in annual maintenance margin is exactly $-\lambda_y \Delta R$. At $\epsilon = 7.5 \times 10^{-4}$ the 0.2 kg floor has a positive stationary margin, whereas 1.944 kg has a negative one. Section S7 reports this fixed-parameter comparison, including the startup receipts. It is a declared parameter example, not a measured loss threshold.

Table 1: Maintenance thresholds for the declared aggregate parameters. Non-return fractions are dimensionless.

Convention	R (kg)	ϵ_c
No radioactive decay	–	0.00079936
Decay; synthetic floor	0.200	0.00076809
Decay; formula scenario	1.944	0.00073357

At $\epsilon = 10^{-3}$, zero-decay continuous operation has a 0.570309 kg/year deficit. With loop decay and $R = 1.944$ kg, Eq. (7) gives 0.757296 kg/year; with the synthetic 0.2 kg floor it gives 0.659177 kg/year. The year convention is stated because the corresponding 365-day values differ slightly.

3.3 Batched receipts and the terminal obligation

For receipt opportunities a_k separated by a declared cadence, I first propagate Eq. (2) over $[a_k, a_{k+1}]$ without a new receipt. Because extra storage mass contributes $x_k e^{-\lambda_d(t-a_k)}$, the minimum receipt preserving $S \geq R$ throughout the interval is

$$x_k = \left[\max_{a_k \leq t \leq a_{k+1}} e^{\lambda_d(t-a_k)} \{R - S_k^0(t)\} \right]_+. \quad (8)$$




Figure 1: Stationary imports as the inner non-return fraction crosses the zero-decay reference threshold. All curves use the same aggregate feed and generation. Reserve and residence-time decay move the maintenance boundary.

I compute the extrema using analytic exponentials and scalar root finding within constant-operation segments and split at planned operating changes. The corresponding peak determines the storage capacity required by the plan. An insufficient receiving capacity rejects the plan; the code does not silently discard material or force a feasible receiver trajectory.

The main continuous-operation cases use $\epsilon/\epsilon_c(R) \in \{0.5, 1.25\}$. Sensitivities include 0, 1 and 2, fixed 10^{-3} , the two storage requirements, and 7-, 30- and 90-day receipt cadences. A separate annual 30-day maintenance calendar replaces a weekly workweek pattern. It is a declared calendar sensitivity, not an inferred commercial outage distribution. Processing continues during its plasma-off intervals.

For a receiver starting in service year C , receipts continue through $(H - C)365.25$ days, where H is the supplier horizon. I also report Eq. (7) at the terminal boundary. A finite schedule is not relabeled a lifetime fuel guarantee. Equation (7) is the continuous-maintenance lower bound, rather than the exact annual mass of a parcel contract. With stationary loop inventories, constant operation and $D = x_\infty > 0$, receipts separated by T days that return storage to R at each interval end have

$$q_T = \frac{D}{\lambda_d}(e^{\lambda_d T} - 1), \quad \bar{x}_T = D \frac{e^{\lambda_d T} - 1}{\lambda_d T}. \quad (9)$$

A fixed transit time L multiplies the dispatch mass by $e^{\lambda_d L}$. These stationary formulas exclude startup filling and calendar changes. The annual-maintenance sensitivity is reported with its full receipt history; its continuous-operation reference tail is not a calendar-specific forecast. The assessment adds a ten-operating-year comparison with $H = C + 10$ to distinguish a genuinely later start from the artificial benefit of fewer modeled operating years.

3.4 Usable-fuel feasibility and retained site inventory

The rejected stream is unavailable to the modeled fuel cycle but remains in W . With no removal route, its stationary value is

$$W_\infty = \frac{(1 - \eta_O)O_\infty/\tau_O + \epsilon J_I}{\lambda_d}. \quad (10)$$

This quantity follows the assumed retention model; adding a real processing or disposal route requires a corresponding state or outward flow. I report $O + I + S + W$ alongside usable storage in Section S7. A

passing check on $S \leq C_r$ addresses only the usable-storage compartment. It cannot certify a total-site inventory allowance. Physical removal of W would change that allowance, but its already-accounted non-return must not be deducted from usable fuel a second time.

4 Supplier execution and customer protection

4.1 Separate physical boundaries

The source offers a net stream after its own operating consumption and retained-inventory requirements. A supply pool B has capacity C_B . A separate optional upstream holding stock U has capacity C_U . Both decay. Between review events the pool accepts offered input until its cap, with excess moving to U when the retry option is enabled. At review opportunities, available upstream material refills the pool. Only overflow beyond both admitted capacities is recorded as externally unaccepted, with no later credit. The aggregate check includes B , U , decay, handovers and external refusals.

This explicit carry-over alternative avoids assuming either that all refused offers disappear or that unlimited upstream storage is available. The single-order reference retains the no-retry convention for a controlled comparison; the principal coupled experiment enables the finite upstream buffer. Quantity allowances are synthetic capacities, not statements of authorized possession. A real site-wide limit must additionally count processing, held and waste stocks.

Receiver and supplier boundaries are operational/accounting boundaries. Co-location alone neither proves that they are one legal facility nor removes the need for measurement and custody control. In the coupled experiment, prescribed receiver receipts are dispatched seven days earlier with mass $x_k e^{7\lambda_d}$. Missed dispatches imply missed original receipt commitments. Parcelized make-up is recorded subsequently. No transport claim is inferred from a mass conversion alone.

4.2 Three dispatch policies

At an event, a customer receives the smaller of its remaining nominal claim, available pool stock above its applicable reservation, and remaining parcel allowance, provided its qualification gate is open. Nominal overdue claims do not decay; physical stock does. The chronological policy serves the oldest due claim first, with incumbents first at equal dates.

The current-stock-only reservation funds every outstanding incumbent and the heel without crediting future input:

$$F_E(t) = m e^{\lambda_y(H-t)} + \sum_{j \in E} \ell_j(t) e^{\lambda_y \max(d_j-t, 0)}. \quad (11)$$

This is a zero-future-receipts stress comparator. It releases to a new customer only above F_E , after serving due incumbents. Its mass protection also requires adequate handover capacity and incumbent eligibility. A pre-existing funding gap under this stronger comparator is not created by the new request.

To examine the intermediate case, let $\underline{s}$ be a declared lower offered-supply curve. The lower-forecast reservation is

$$F_L(t) = \max_{u \in \mathcal{E}_t} \left\{ m e^{\lambda_y(u-t)} + \sum_{j \in E: d_j \leq u} \ell_j(t) e^{\lambda_y \max(d_j-t, 0)} - \int_t^u \underline{s}(v) e^{\lambda_y(v-t)} dv \right\} \vee m. \quad (12)$$

The event set $\mathcal{E}_t$ includes the current date, due dates, flow changes and horizon; it is sufficient under the piecewise-constant assumptions. Here each of the two rate factors is credited only at its marginal 2.5% lower quantile, and the whole specified interruption interval is removed. This gives a time-zero conservative supply comparator, not a guarantee under distribution shift, capacity rejection or unforeseen outages. It

uses no realized future scenario rates. It is distinct from reproducing the optimized rationing policy of Deshpande et al. [16].

4.3 Customer-specific outcomes and emergency use

For a customer's original due dates, let $L_j(q)$ indicate any late quantity with a proposed new schedule of total nominal scale q . I compare it with the paired no-new-order outcome:

$$\Delta p_j = \Pr[L_j(q) = 1] - \Pr[L_j(0) = 1], \quad p_j^{\text{newlate}} = \Pr[L_j(q) = 1, L_j(0) = 0]. \quad (13)$$

The two statistics differ if changed receipts or policy improve some histories. I retain original-date missing mass, subsequent completion delay and unresolved claims separately. Mean or percentile delay conditions on completion where stated; unresolved cases remain explicitly censored.

The emergency target includes the new accepted schedule, rather than only incumbent quantities. With a withdrawal heel m , a restart allowance and an explicitly defined annualized contractual-service quantity determine $e(t)$; normal target is $h(t) = m + e(t)$. The executable specification is in Supplement S3. Utilization is measured on actual stock,

$$\Psi = \max_t \min\{1, [h(t) - B(t)]_+/e(t)\}. \quad (14)$$

A target reduction may end an episode without replenishing material, so reported duration means time below the contemporaneous target. I move CVaR severity screening to supplementary analysis because it is not the main quantity-selection mechanism. Its 99% bootstrap upper diagnostic uses 1999 resamples and separate populations. General loss distributions and bootstrap approximations follow established methods [14, 6, 19].

For each paired contract comparison, the no-new-order history retains the same scheduled review and dispatch appointments as the history with the receiver. Only the new claims are deleted. This controls the opportunity to make up an incumbent order or refill the pool from upstream stock: the new contract does not obtain credit merely by creating extra review dates. Appointment dates may differ between candidate service schedules, and are declared with each candidate.

5 Computational protocol

The service-risk criterion differs from the earlier newly-late increment screen: the primary coupled decision tests the total event that any incumbent or receiver parcel is late. At each fixed candidate I form a one-sided Clopper–Pearson bound with tail budget $0.01/82$ [18]. The first passing date in ascending order would be selected. The grid, thresholds and inputs were fixed before the final independent screening and validation; preliminary computational runs informed the sample budget, as recorded in the protocol. No external preregistration is asserted. Figures and tables separate screening from independent evaluation.

The receiver equations generate required receipts for a planned operating calendar. The supplier experiment tests whether it can fulfill them. Successful fulfillment supports that prescribed operation; failed fulfillment does not justify continuing to call the receiver's nominal burn physically executed. The work is therefore a stochastic qualification of model-derived contracts, not a feedback-controlled simulation of a starving plasma. Adaptive derating, repair and reactor restart remain additional models.

Table 2: Principal service specification. These are synthetic contract and supply assumptions, except for the stated decay constant.

Item	Value or convention
Initial pool	Four lognormal lots, each mean 1.5 kg, coefficient of variation 5%, pairwise Pearson correlation 0.5.
Offered rate	Mean 0.6 kg/yr before year 8; 0.9 kg/yr after year 8. Two correlated lognormal factors: coefficient of variation 15%, Pearson correlation 0.6.
Interruption	Probability 0.12 of no offered supply during years 4.5–6; independent of initial stock and rate factors.
Incumbents	2 kg at year 4 (E1), 3.5 kg at year 8 (E2).
Physical service	Withdrawal heel 0.2 kg; pool cap 12 kg; optional upstream cap 12 kg; monthly reviews; up to eight 0.5 kg parcels at an event. Due events are also service opportunities.
Qualification	A receiver-level hold occurs with probability 0.005 at its first request and lasts 0.25 yr. It is common to that receiver's affected parcels.
Single-order experiment	60,000 paired paths; policies and quantities 0, 0.5, 1, 1.5 and 2 kg at year 2.
Coupled receiver	$R = 1.944$ kg; empty initial loop/storage states; complete initial filling charged to supplier; 30-day receipt cadence, seven-day transit; continuous-operation maintenance ratios 0.5 and 1.25.
Coupled horizon and grid	$H = 14$; candidate starts 2–12 yr at 0.25-year spacing; all receiver demand through H ; synthetic receiver-storage allowance 25 kg; terminal import rate reported.
Final screening / evaluation	12,000 / 20,000 separately generated paths. Total all-customer late-commitment criterion 1%; 99% simultaneous sampling-confidence budget over 82 fixed candidates.

6 Results

6.1 The receiver changes from startup demand to continuing service

Equations (4)–(7) explain the sharp increase in external demand around the maintenance boundary. Figure 2 follows 14 years rather than erasing demand after the first year. Table 3 retains $\epsilon = 10^{-3}$ as a fixed-fraction comparison. Its receipts are additional to the initially owned 1.5 kg; in the coupled experiment all initial stock is instead supplied explicitly.

Table 3: Fixed $\epsilon = 10^{-3}$, 30-day receipt cadence. Receipts are kg additional to initial 1.5 kg; the final column is the total storage mass required if everything is supplied at startup.

Years	Receipts, $R = 0.2$	Receipts, $R = 1.944$	One initial load, $R = 1.944$
1	0.7403	2.5826	4.1026
5	3.3831	5.6188	7.6963
10	6.6866	9.4140	13.4895
14	9.3294	12.4502	19.4522

With $R = 1.944$ kg and $\epsilon = 0.5\epsilon_c$, the first 30-day replenishment for an empty receiver is about 3.3114 kg; the initial 1.944 kg reserve alone is insufficient to fill the circulating loops. At $1.25\epsilon_c$, the first receipt is about 3.3666 kg and the stationary import obligation is 0.52127 kg/year. Receiving-capacity sensitivity is expressed as a required peak, rather than selecting one cap between the phased and all-at-start quantities. Supplement S5 gives the complete capacity frontier and the extra long-hold sensitivity.




Figure 2: Long-horizon receipts at 30-day opportunities, in addition to an initially owned 1.5 kg. All curves maintain 1.944 kg in planned storage and use their ratio to the reserve-adjusted threshold. Above-threshold imports continue after the transient.

6.2 Current stock, credited supply and whose promise is protected

At year 2, $F_E = 7.5364$ kg versus nominal available stock 6.4964 kg. In the 60,000-path base sample, 99.9783% are not fully funded under the zero-future-receipts comparator before any new order. The lognormal law has unbounded upper support, so a finite sample maximum cannot establish universal underfunding. Nor must every incumbent contract expire before the reservation can become funded: sufficiently favorable intervening receipts can fund it earlier.

The lower-forecast reservation is 5.0605 kg at the same date, between the withdrawal heel and the current-stock-only reservation. Table 4 and Figure 3 compare executed outcomes at equal requests. The lower-forecast policy reduces the incumbent late frequency modestly but shifts some risk to the new customer's due date. Its effect is a trade-off, not an optimality claim.

Table 4: Year-2 request of 1 kg, same 60,000 histories. Baseline incumbent late frequency is zero in this sample. Median new completion delay is conditional on completion.

Policy	Incumbent late (%)	New late (%)	New delay median (d)
Chronological	0.260	0.503	0.0
Lower-forecast reserve	0.208	5.388	0.0
Current stock only	0.000	100.000	1278.4

The current-stock-only policy remains useful as a structural comparator; its median delay is retained only in the detailed results. The delay primarily measures the initial zero-inflow funding gap and subsequent rebuilding, rather than an intrinsic property of the 1 kg order. Horizon sensitivities cover this policy as well as chronological and forecast-reserved dispatch.




Figure 3: Incumbent late-commitment frequency against request size. New-customer outcomes are separately reported in Table 4 and Supplement S4. A policy name has no physical effect when it produces the same dispatch history.

6.3 Stochastic qualification of complete receiver contracts

All prospective receiver demand through year 14 enters each trial. The principal model has a separate receiver boundary, seven-day decay-adjusted dispatch, common receiver-qualification holds and finite parcel opportunities. Figure 4 shows simultaneous screening bounds; Table 5 provides the independent evaluation at fixed comparison dates.

Table 5: Independent evaluation of complete receiver contracts: 20,000 common histories, chronological dispatch. Required kg includes initial filling and subsequent receipts, but excludes transit decay charged separately at dispatch.

ϵ/ϵ_c	Start	Required kg	Incumbent late (%)	Receiver late (%)
0.5	2.0	3.311	55.605	0.545
0.5	8.0	3.311	93.440	0.545
0.5	10.0	3.311	0.000	30.540
0.5	12.0	3.311	0.000	1.980
1.25	2.0	9.593	99.980	99.980
1.25	8.0	6.459	94.250	94.490
1.25	10.0	5.414	0.000	33.105
1.25	12.0	4.369	0.000	2.290

No candidate on the declared grid satisfies the total 1% screen for either principal receiver class. Failure to obtain an upper-bound certificate is, in general, different from evidence that risk exceeds the limit. I therefore also report simultaneous lower bounds, using a 1% sampling-error budget over both bounds and all 82 cases. Every tested case has a lower bound above 1% under the specified generating model; Section S7 supplies the complete classification. This rules out the candidate set for that model and criterion, not all




Figure 4: Upper screening bounds for the total event of any missed original commitment. Receiver operation extends from each candidate start through service year 14. Each curve uses 12,000 paths and the fixed familywise confidence budget. Later starts also model fewer operating years, addressed separately in Supplement S5.

possible schedules or designs. This is the result of the stated service scale, reserve, initial filling and supply distribution; I do not assign a replacement “first feasible year.” Existing-contract failure and receiver failure remain separate in the table. When a large initial fill is dispatched early, the new receiver can receive it on time while later incumbents carry the shortage. For a late start, the binding risk can instead be initial receiver filling after the legacy deliveries.

The fixed-loss, low-floor diagnostic isolates the effect of horizon truncation. Starting at year 2 and charging only one year of demand requires about 2.240 kg, including initial filling; carrying the same receiver through year 14 requires about 9.508 kg. The associated independent customer outcomes are in Supplement S4. This is not a fair comparison of two equal-duration services: it is a demonstration that the shorter contract omits material obligations. A full comparison must state the service duration and terminal continuing imports.

6.4 Parcel cadence, reserve scale and capacity

At year 10, I compare weekly, 30-day and 90-day receipt opportunities with the same seven-day transit. Their required peaks and stochastic delivery outcomes appear in Table 6. Cadence is varied within a fixed R and ϵ . Across the two reserve rows, however, $\epsilon = 1.25\epsilon_c(R)$ changes; that cross-row difference is a regime-normalized comparison. The fixed- ϵ reserve experiment in Section S7 supplies the separate controlled contrast. Increasing cadence spacing moves receipts forward to cover longer intervals. It can increase both initial demand and the necessary receiver capacity even when the stationary import rate is unchanged.

Supplier capacity sensitivity separately enables or disables finite upstream retry storage. Rejected material is kept in the upstream state only when that state is explicitly modeled and has capacity. Thus the comparison exposes, rather than hides, the assumption behind an offered-but-unavailable stream. Full tables identify population sizes, baseline failure and censored completion.

Table 6: Fixed start year 10, $\epsilon = 1.25\epsilon_c(R)$, operation through year 14; 8,000 independent histories. Each row includes all prescribed receiver receipts.

R kg	Cadence d	Required kg	Peak kg	Receiver late (%)
0.2	7.0	3.7640	1.5900	0.600
0.2	30.0	3.7679	1.6246	0.650
0.2	90.0	3.7779	1.7151	0.662
1.944	7.0	5.4099	3.3336	31.500
1.944	30.0	5.4136	3.3666	33.013
1.944	90.0	5.4231	3.4530	37.413

7 Interpretation and limits

The receiver threshold and the service policy answer different parts of one question. Below the maintenance threshold, startup filling and capacity can dominate the first commitment. Above it, a sustained fuel deficit becomes a repeated claim on the supplier. Even a stationary import rate smaller than mean source inflow can be incompatible with incumbent dates, correlated low receipts, parcel delays or initial filling. Mean resource sufficiency and service feasibility are therefore separate tests.

The analytical threshold is specific to the declared serial aggregate. First-wall retention, blanket trapping, state-dependent extraction, direct internal recycling, maintenance of processing equipment and plasma control are compressed or absent. A supplementary stationary sensitivity routes a fraction of newly bred material through a longer hold; its additional decay lowers the margin. That exercise flags the direction of a missing mechanism but does not estimate real trapping. The detailed fuel-cycle literature retains separate components [8, 9].

The reserve formula, 0.2 kg alternative, loss fractions, supply distributions, storage caps, risk criteria and operating calendars have explicit but heterogeneous provenance. The service inputs are synthetic. A low forecast credited under one distribution need not protect commitments under a different failure mechanism. A finite Monte Carlo confidence statement applies to its generating model, not to an operating facility. The two-node retry boundary and deterministic transit are declared simplifications. Actual transport permissions, assay acceptance, decay-produced helium handling and complete site allowances need facility-specific evidence [11, 12].

The coupled experiment imports receiver demand as a planned contract and evaluates stochastic fulfillment. It does not recalculate plasma operation after a late parcel. A feedback study would alter burn, recovery, restart requirements and subsequent demand together. Likewise, the current-stock and lower-forecast reservations are transparent comparison policies, not optimized customer-service contracts. A policy objective would require interruption costs, priority terms and procurement alternatives. This study establishes neither a preferred reactor architecture nor a commercial price.

7.1 Practical use and remaining evidence requirements

The implementable output of this study is a planning screen: a stated receiver model produces a receipt schedule, and an explicit supplier policy maps it to dated shortfalls. Its immediate uses are checking omitted continuing demand, comparing reserve definitions at fixed parameters, and identifying whose contract carries a shortage. Engineering deployment requires an independently validated receiver model, source arrival histories with uncertainty, actual capacity and package limits, and closed-loop responses to late fuel. A contractual missed date is not itself an observed plasma shutdown. Receipt schedules that maintain an assumed floor also do not prove qualified gas composition, safe waste handling or plant availability.

8 Conclusion

I find that the maintenance margin must be evaluated before a receiver is classified as a startup-only customer. Under the declared aggregate parameters, a non-return fraction near 8×10^{-4} marks the no-decay transition, while maintained reserve and processing decay lower that boundary. The 10^{-3} case requires continuing imports; truncating its demand after one year changes the obligation being tested.

A current-stock-only reservation is structurally underfunded in almost all year-2 reference histories because it disregards the future receipts on which existing contracts rely. A lower-forecast reservation supplies an intermediate, inspectable policy, with separate effects on incumbent and new-customer service. Complete receipt schedules, rather than a single startup mass, connect the receiver's margin to those outcomes. The stochastic and capacity results remain conditional on their declared model and show why contract duration, reserve semantics and initial filling must accompany a tritium ATP claim.

Data and code availability

The accompanying revision-11 archive contains source programs, input conventions, generated results, separate figures, manuscript sources and 64 executed verification checks. Independent ODE and matrix-exponential implementations check the receiver solution; supplier and upstream aggregate balances are reconciled separately. Supporting Information reports each check and the full numerical tables. The archive contains the new controlled-reserve and retained-stock results, and a code-checked execution specification. Principal random seeds come from a single configuration module. Supplement S7 distinguishes the reconstructed results from validation against a facility. A repository DOI has not been assigned.

References

- [1] Z. Zhao, M. O. Ball, and M. Kotake. Optimization-Based Available-To-Promise with Multi-Stage Resource Availability. *Annals of Operations Research* 135, 65–85 (2005). [10.1007/s10479-005-6235-7](https://doi.org/10.1007/s10479-005-6235-7).
- [2] Oracle. Available to Promise (ATP). *MRP and Supply Chain Planning Help*, especially Calculating ATP and Negative ATP. Undated documentation; accessed 3 October 2026. https://docs.oracle.com/cd/A60725_05/html/comnls/us/mrp/atp.htm.
- [3] P. Ghare and G. Schrader. A Model for Exponentially Decaying Inventories. *Journal of Industrial Engineering* 14, 238–243 (1963). As cited in Nahmias and Wang (1979), reference 6; consulted indirectly.
- [4] S. Nahmias and S. S. Wang. A Heuristic Lot Size Reorder Point Model for Decaying Inventories. *Management Science* 25(1), 90–97 (1979). [10.1287/mnsc.25.1.90](https://doi.org/10.1287/mnsc.25.1.90).
- [5] S. Nahmias. Perishable Inventory Theory: A Review. *Operations Research* 30(4), 680–708 (1982). [10.1287/opre.30.4.680](https://doi.org/10.1287/opre.30.4.680).
- [6] T. Hasuike. Risk Management of Two-period Inventory Models for Perishable or Deteriorating Products Based on Conditional Value-at-Risk. *Journal of Japan Industrial Management Association* 64(2E), 231–243 (2013). [10.11221/jima.64.231](https://doi.org/10.11221/jima.64.231).
- [7] R. Avenhaus and G. Spannagel. Tritium Accountancy and Hidden Inventories. *Fusion Technology* 28(3P2), 1003–1008 (1995). [10.13182/FST95-A30537](https://doi.org/10.13182/FST95-A30537).

- [8] M. Abdou, M. Riva, A. Ying, C. Day, A. Loarte, L. R. Baylor, P. Humrickhouse, T. F. Fuerst, and S. Cho. Physics and technology considerations for the deuterium–tritium fuel cycle and conditions for tritium fuel self sufficiency. *Nuclear Fusion* 61, 013001 (2021). 10.1088/1741-4326/abf35. Institutional abstract: <https://impact.ornl.gov/en/publications/physics-and-technology-considerations-for-the-deuterium-tritium-f/>.
- [9] S. Meschini, S. E. Ferry, R. Delaporte-Mathurin, and D. G. Whyte. Modeling and analysis of the tritium fuel cycle for ARC- and STEP-class D–T fusion power plants. *Nuclear Fusion* 63, 126005 (2023). Published 14 September 2023. 10.1088/1741-4326/acf3fc.
- [10] J. F. Parisi, A. Diallo, and J. A. Schwartz. Simultaneous enhancement of tritium burn efficiency and fusion power with low-tritium spin-polarized fuel. *Nuclear Fusion* 64, 126019 (2024). Published 4 October 2024. 10.1088/1741-4326/ad7da3.
- [11] Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Guide to the Expression of Uncertainty in Measurement*. JCGM 100:2008, Section 5.2. https://www.bipm.org/documents/20126/2071204/JCGM_100_2008_E.pdf.
- [12] Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—The Role of Measurement Uncertainty in Conformity Assessment*. JCGM 106:2012. https://www.bipm.org/documents/20126/2071204/JCGM_106_2012_E.pdf.
- [13] L. L. Lucas and M. P. Unterweger. Comprehensive Review and Critical Evaluation of the Half-Life of Tritium. *Journal of Research of the National Institute of Standards and Technology* 105(4), 541–549 (2000). 10.6028/jres.105.043. <https://pmc.ncbi.nlm.nih.gov/articles/PMC4877155/>.
- [14] R. T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance* 26(7), 1443–1471 (2002). 10.1016/S0378-4266(02)00271-6. Author manuscript: <https://sites.math.washington.edu/~rtr/papers/rtr187-CVaR2.pdf>.
- [15] M. Kovari, M. Coleman, I. Cristescu, and R. Smith. Tritium resources available for fusion reactors. *Nuclear Fusion* 58, 026010 (2018). 10.1088/1741-4326/aa9d25. Related author manuscript WPPMI-PR(17)17185: https://scipub.euro-fusion.org/wp-content/uploads/eurofusion/WPPMI-PR17_17185_submitted-1.pdf.
- [16] V. Deshpande, M. A. Cohen, and K. Donohue. A Threshold Inventory Rationing Policy for Service-Differentiated Demand Classes. *Management Science* 49(6), 683–703 (2003). 10.1287/mnsc.49.6.683.16022.
- [17] National Institute of Standards and Technology. Atomic Weights and Isotopic Compositions for Hydrogen. Hydrogen-3 relative atomic mass 3.0160492779(24); accessed 3 October 2026. https://physics.nist.gov/cgi-bin/Compositions/stand_alone.pl?ele=H.
- [18] C. J. Clopper and E. S. Pearson. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika* 26(4), 404–413 (1934). 10.1093/biomet/26.4.404.
- [19] B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* 7(1), 1–26 (1979). 10.1214/aos/1176344552.
- [20] E. B. Wilson. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* 22(158), 209–212 (1927). 10.1080/01621459.1927.10502953.
- [21] D. G. Whyte, R. Delaporte-Mathurin, S. E. Ferry, and S. Meschini. Tritium burn efficiency in deuterium–tritium magnetic fusion. *Nuclear Fusion* 63, 126019 (2023). Published 25 September 2023. 10.1088/1741-4326/acf3fb.

Supporting Information

Start-Up or Sustained Supply? Tritium Fuel-Cycle Margins and Customer Commitments

Norio Yamazaki | Revision 11 | 3 October 2026

S1 Receiver equations and reference replay

The published aggregate reference is Meschini et al., *Nuclear Fusion* 63, 126005 (2023), DOI 10.1088/1741-4326/acf3fc, equations (26)–(30). The outer-to-inner-to-storage connection is also shown in its Figure 4. Three equations describe differential inventories; the other two statements specify initial conditions. I inspected the publisher-provided text through its public ResearchGate page. The equation replay is independent of the original authors' full system implementation.

For constant operation, let $a_O = 1/\tau_O + \lambda_d$, $a_I = 1/\tau_I + \lambda_d$, and

$$O_\infty = g/a_O, \quad I_\infty = (F - b + \eta_O O_\infty / \tau_O) / a_I, \quad (S1)$$

$$c_O = \frac{\eta_O(O_0 - O_\infty)}{\tau_O(a_I - a_O)}, \quad c_I = I_0 - I_\infty - c_O. \quad (S2)$$

The solution for the two loop inventories is

$$O(t) = O_\infty + (O_0 - O_\infty)e^{-a_O t}, \quad (S3)$$

$$I(t) = I_\infty + c_O e^{-a_O t} + c_I e^{-a_I t}. \quad (S4)$$

Define $r_I = (1 - \epsilon)/\tau_I$ and $J_\lambda(t) = (1 - e^{-\lambda_d t})/\lambda_d$, with limit t at zero decay. Then

$$S(t) = S_0 e^{-\lambda_d t} + (r_I I_\infty - F) J_\lambda(t) + r_I c_O \frac{e^{-\lambda_d t} - e^{-a_O t}}{a_O - \lambda_d} + r_I c_I \frac{e^{-\lambda_d t} - e^{-a_I t}}{a_I - \lambda_d}. \quad (S5)$$

Total inventory obeys $T(t) = T_0 e^{-\lambda_d t} + (g - b) J_\lambda(t)$ between receipts. The rejected inventory is $W = T - O - I - S$. The supplied implementation assumes the unequal time constants in the stated parameter set; coincident-time-constant limits would require the appropriate limiting expression.

Table S1: Reference-replay parameters and scope.

Parameter	Value	Source/scope
Burn b	9×10^{-7} kg/s	Published aggregate illustration
Feed ratio β	0.01	Published aggregate illustration
Generation ratio Λ	1.08	Published aggregate illustration
Inner / outer times	4 h / 24 h	Published aggregate illustration
Initial (O, I, S)	(0, 0, 1.5) kg	Published aggregate illustration
Decay and non-return	Zero	Reference only, not extension
Extension R	0.2 or 1.944 kg	Synthetic floor / mixed-parameter reserve scenario
Extension λ_d	$\ln 2 / 4500$ day $^{-1}$	Selected evaluated half-life

The serial solution yields 0.142528 kg at 2.78508 days. A model with $\dot{I} = F - b - I/\tau_I$ and $\dot{S} = I/\tau_I + O/\tau_O - F$ instead yields 0.155391 kg at 2.60289 days. Its output is a distinct routing sensitivity, not a correction of the published topology. After the fast inner exponential is small, the serial minimum occurs near

$$t_{\min} \approx \tau_O \ln \left[\frac{\Lambda \tau_O}{(\Lambda - 1)(\tau_O - \tau_I)} \right] = 2.78501 \text{ days}. \quad (S6)$$

The source prose states 2.7 days separately. Its Figure 3 explicitly uses illustrative curves rather than this parameter set. The manuscript reports the textual value, equation result and figure scope separately, without attributing the difference to an unverified cause.




Figure S1: Independent equation solutions for the serial reference and the alternative direct return. The figure is generated here, not digitized from the original paper.

S2 Reserve, threshold and receipt construction

The formula $R = F\chi t_{\text{res}}$ is directly traceable to equation (25) of Meschini et al. The full-design baseline and the aggregate pedagogical example have different feeds. Combining the aggregate $F = 7.776$ kg/day with $\chi = 0.25$ and one day gives 1.944 kg. No claim is made that an actual ARC plant has this reserve. The synthetic 0.2 kg comparison is retained to expose the effect of changing that assumption rather than silently changing the benchmark.

Setting stationary loop derivatives to zero gives

$$O_{\infty} = \frac{g}{1/\tau_O + \lambda_d}, \quad \frac{I_{\infty}}{\tau_I} = J_I = \frac{F - b + \eta_O g / (1 + \lambda_d \tau_O)}{1 + \lambda_d \tau_I}. \quad (\text{S7})$$

At $S = R$ the margin is $(1 - \epsilon)J_I - F - \lambda_d R$. Its derivative with respect to ϵ is $-J_I$, and its derivative with respect to a separately varied R is $-\lambda_d$. The zero-decay formula in the main text follows by setting $\eta_O = 1$ and $g = \Lambda b$. Holding R fixed, a higher burn-to-feed ratio reduces the gross return stream to which non-return applies. If instead $R = F\chi t_{\text{res}}$ is varied with F , both effects must be recalculated. The file `threshold_surface.csv` varies $\beta \in \{0.005, 0.01, 0.02\}$ and $\Lambda \in \{1.04, 1.08, 1.12\}$ with that reserve convention.

For each fixed receipt interval, I maximize $(R - S^0(t))e^{\lambda_d(t-a)}$. Its derivative has the sign of $F + \lambda_d R - (1 - \epsilon)I(t)/\tau_I$. The inner trajectory has at most one extremum for the distinct two-exponential representation. Bracketing around that extremum identifies all roots needed for the lower-storage requirement. A scalar root solver and endpoint tests then determine the required receipt. Storage-capacity peaks are checked separately by the derivative of S and endpoint values. The independent ODE checks use the four original right-hand sides rather than differentiating the closed-form code.

Receipts are required quantities for a prescribed operating history. A supplier shortfall is recorded as failure of that contract. The receiver is not credited with unreceived fuel, and its planned operation is not

asserted to continue after such a failure. Delayed dispatch may satisfy a nominal mass claim later; it cannot retroactively restore uninterrupted operation at the original due date.

The long-residence sensitivity places a fraction f of newly bred material behind an additional first-order hold τ_L , without changing the inner path. At stationary state, replace the bred contribution by

$$\frac{\eta_{OG}}{1 + \lambda_d \tau_O} \left(1 - f + \frac{f}{1 + \lambda_d \tau_L} \right). \quad (S8)$$

This flags additional retained-inventory decay. It does not model measured wall trapping, chemical transport, or an arbitrary circulating-stream retention fraction. Values $f = 0.05$ and $\tau_L = 30, 180$ days are declared sensitivity inputs. Full fuel-cycle studies include distinct component hold-ups; the aggregate model cannot establish their inventory in a real receiver.

S3 Supplier implementation and distributional protocol

S3.1 Random variables

For mean μ , coefficient of variation c and equal Pearson correlation ρ , let $v = \log(1 + c^2)$ and $\rho_\ell = \log(1 + \rho c^2)/v$. Positive quantities are generated as

$$X_j = \mu \exp\{\sqrt{v}(\sqrt{\rho_\ell} Z_0 + \sqrt{1 - \rho_\ell} Z_j) - v/2\}, \quad (S9)$$

where the normals are independent. Initial lots use $(\mu, c, \rho) = (1.5, 0.05, 0.5)$; two rate factors use $(1, 0.15, 0.6)$. The domains and interruption draw are independent by assumption.

A left-tail stress retains first two moments while replacing the rate factors by

$$X_j = 1 + 0.15(\sqrt{\rho} V_0 + \sqrt{1 - \rho} V_j), \quad V_j = (0.1 - B_j)/\sqrt{0.1(0.9)}, \quad B_j \sim \text{Bernoulli}(0.1). \quad (S10)$$

The common variable represents a stylized common derating component, with individual components added. This is a mechanism sensitivity, not a fitted failure distribution. The combined stress raises correlations to 0.9, interruption probability to 0.24, and the end of the interruption to year 7.5. Base and stress baselines are always shown separately from the effect of a new order.

S3.2 Reservations and target accounting

Current-stock funding ignores future supply; it is stronger than an ordinary forecast-funded promise. The lower-forecast comparator takes the marginal 2.5% quantile of each base rate factor, 0.738240 of its mean, and removes all inflow in years 4.5–6. Two marginal 2.5% exceedance failures have total probability at most 5% by a union bound, without requiring independent rate factors. This is a statement about offered rates under the base law. Capacity refusal, more severe outages, qualification or miscalibration remain separate.

The forecast reservation is evaluated at current time, incumbent due dates, flow changes and horizon. For the piecewise-constant lower curve, the decay-discounted slack is monotone on each open interval. Event checks therefore cover the same linear lower-envelope constraints as a fine time grid. Changing a realized post-year-8 rate in the verification does not change the year-2 dispatch; the policy never reads a scenario's future rate realization.

The emergency target is $e(t) = r_{\text{restart}} + \tau_{\text{cover}}[\nu_E(t) + \nu_N(t)]$, with $r_{\text{restart}} = 0.2$ kg and $\tau_{\text{cover}} = 1$ year. The total target is $h = m + e$, with the separate service heel $m = 0.2$ kg. Incumbent service intensity is $\nu_E = 5.5/8$ kg/year until year 8 and zero thereafter. A new accepted schedule contributes its total nominal dispatch mass divided by the larger of its due-date span and one year, until its final due date. This is a declared contractual intensity, not a reactor feed rate. The allocation is present from acceptance, hence it

can affect the pre-first-delivery reserve metric. A comparison with incumbent-only targeting is supplied. Ending an allocation at the final due date changes the target, not the physical stock.

An optional upstream retry store retains refused pool receipts. Between reviews, the combined stock and the capped pool each obey the exact admitted-flow balance. Their difference is upstream stock. At a review it transfers into free pool space. Overflow beyond the two finite capacities is explicitly outside the system. Total-state conservation is checked with cumulative admitted/offered input, external refusal, handover and decay. There is no implicit credit for infinite upstream capacity.

S3.3 Sampling, selection and interpretation

Single-order reference paths: 60,000, seed 20261028. Supplementary bootstrap populations: 20,000 each, seeds 20261108 and 20261109; 1999 resamples. Horizon/capacity sensitivity: 8,000, seed 20269999. Final coupled screening: 12,000, seed 20261230; independent coupled evaluation: 20,000, seed 20261231. The 82 alternatives comprise two receiver classes and 41 dates $2 + 0.25k$, $k = 0, \dots, 40$. Candidates are examined in increasing date order; the first one whose all-customer failure upper bound is at most 1% and whose required receiver peak is at most 25 kg would pass.

A pilot using 4,000 and subsequently 12,000 paths was used for computational development and sample-budget assessment. Final screening used a new seed after the grid, model and sample size had been fixed; final independent evaluation used another new seed. This protocol is reproducible computational separation, not externally registered research. No date was added because a favorable result occurred there. The final grid supports no passing date. The independently reported fixed comparison dates are 2, 8, 10 and 12 years.

For k failed histories in n samples the one-sided upper binomial bound is the appropriate beta quantile (Clopper–Pearson). The confidence tail is $0.01/82$ in screening, covering possible date selection. Wilson intervals in the machine-readable evaluation tables quantify Monte Carlo sampling, not unknown real-plant model error. The total criterion is distinct from an incumbent incremental-risk allowance. Bootstrap upper CVaR is reported only as an approximate severity diagnostic.

The scalar Cantelli comparison retained in the archive uses the rounded prior path-minimum estimates $\mu_Z = 3.2330$ kg and $u_Z = 0.8996$ kg. At $\alpha = 0.005, 0.01, 0.05$ it is respectively negative in all cases. These are alternative diagnostics based on estimated, rounded moments; they neither certify the new coupled contracts nor establish a general weakness of moment methods.

S4 Customer risk and operational sensitivities

Table S2: Supplementary severity assessment. Each point is evaluated on 20,000 paths; the approximate upper bootstrap statistic uses 1999 whole-path resamples.

Population	q kg	CVaR	99% bootstrap upper	Resamples
base	1.0	0.0921	0.1130	1999
base	2.0	1.0000	1.0000	1999
left tail	1.0	0.4597	0.5112	1999
left tail	2.0	1.0000	1.0000	1999

For the left-tail population and a 1 kg chronological request, the upper bootstrap result exceeds the illustrative 0.5 severity level even though the point estimate is below it. At 2 kg the upper tail exhausts the emergency tranche and both values equal one. Thus a severity criterion can bind, but no universal optimum follows from selecting that threshold. An atom at the maximum also limits what resampling can learn about unseen failures.




Figure S2: New-customer on-time failure versus request size. The same 60,000 histories are used as in the main incumbent-risk figure. Current-stock-only reservation is structurally underfunded at year 2 on almost all reference paths.

Table S3: Horizon sensitivity for all three policies, 8,000 common paths. Delays condition on completion. Capacity and available parcel terms are unchanged.

Horizon	Policy	E2 late (%)	New late (%)	New median d
10	Chronological	0.300	0.550	0.0
10	Lower-forecast reserve	0.250	5.475	0.0
10	Current stock only	0.000	100.000	1217.5
12	Chronological	0.300	0.550	0.0
12	Lower-forecast reserve	0.250	5.475	0.0
12	Current stock only	0.000	100.000	1247.9
14	Chronological	0.300	0.550	0.0
14	Lower-forecast reserve	0.250	5.475	0.0
14	Current stock only	0.000	100.000	1278.4
16	Chronological	0.300	0.550	0.0
16	Lower-forecast reserve	0.250	5.475	0.0
16	Current stock only	0.000	100.000	1308.8

Table S4: Explicit finite upstream storage versus no-retry offered supply, 8,000 common paths. Refusals are quantities outside the combined declared boundary.

Pool cap	Upstream retry	E2 late (%)	External refusal kg	Upstream end kg
6	no_retry	0.425	0.7604	0.0000
6	retry	0.300	0.0000	0.3424
8	no_retry	0.300	0.0129	0.0000
8	retry	0.300	0.0000	0.0125
12	no_retry	0.300	0.0000	0.0000
12	retry	0.300	0.0000	0.0000

The current-stock reservation depends on its protection horizon through the future heel and obligations. This dependence is part of the comparator. A term that remains unchanged under chronological dispatch need not remain unchanged under a reservation policy. Empty terminal backlog is reported separately from being on time.

Table S5: One-year demand truncation versus the same receiver continued through supplier year 14, 20,000 independent paths. Fixed $\epsilon = 10^{-3}$; synthetic floor 0.2 kg; initial filling is charged in both cases.

Start	Demand years	Required kg	Incumbent late (%)	Receiver late (%)
2.0	1.0	2.240	13.705	0.545
2.0	12.0	9.508	99.845	99.840
8.75	1.0	2.240	0.000	3.235
8.75	5.25	5.048	0.000	3.650
10.0	1.0	2.240	0.000	0.600
10.0	4.0	4.222	0.000	0.740

This table does not compare equal service durations. It quantifies the demand excluded by a short contract record. The fixed stationary import rate after the last recorded receipt remains positive. Each row separately records both incumbent and receiver late outcomes.

S5 Equal-duration, capacity and reserve sensitivities

The equal-duration comparison charges the same model-derived receiver quantity for a fixed loss class at every start date. Calendar position can still change conflict with the incumbent deliveries. These runs were declared as a sensitivity and are not additional candidates used to select a favorable start from the primary grid.

The all-at-start curve is constant with respect to the candidate replenishment cadence because that plan has no subsequent receipt. All mass values include the initial receiver storage when calculating the capacity requirement. Additional-receipt totals in the long-term table are explicitly reported relative to initial 1.5 kg, whereas the stochastic coupled contracts begin with empty receiver inventories.

Table S6: Equal ten-operating-year receiver services, 8,000 paths. Supplier horizon changes with the receiver start. Future supply after year 14 extends the declared synthetic rate law; it is not a long-range forecast.

Start	Horizon	ϵ/ϵ_c	Required kg	Incumbent late (%)	New late (%)
2.0	12.0	0.5	3.311	54.975	0.512
2.0	12.0	1.25	8.548	99.987	99.987
8.0	18.0	0.5	3.311	93.088	0.512
8.0	18.0	1.25	8.548	93.987	94.175
12.0	22.0	0.5	3.311	0.000	1.700
12.0	22.0	1.25	8.548	0.000	1.988




Figure S3: Receiving capacity required by timed and all-at-start plans over 365 days, fixed $\epsilon = 10^{-3}$. Solid curves use fixed replenishment opportunities. Dashed lines are the total initial mass required by the zero-subsequent-receipt plan. Any candidate cap can be compared against these curves rather than being chosen between two outcomes.

The data include an annual 30-day planned-maintenance case over five years. It changes burn and generation, while processing remains available. The arrangement is a controllable scenario, not an empirical failure calendar. A randomly failed processing train would change recovery, required reserve and supply simultaneously and requires its own branch model.

S6 Evidence record and verification

S6.1 Source access and interpretation

Verification identifiers describe the tested implementation, including `alternative_routing_numbers`. The accompanying evidence register records the boundary between published inputs and calculations performed for this study.

The current revision directly consulted Meschini et al.’s publisher-provided text through ResearchGate for equations (25)–(30), Figure 3’s caption, Figure 4’s stated topology and the distinction between the aggregate example and the complete ARC parameter table. The reserve time and failed fraction are scenario choices in that paper, not regulatory quantities. The full model’s component-residence structure supports the stated limitation of the lumped extension.

Abdou et al.’s institutional abstract was used for the general dependence of fuel self-sufficiency on throughput, inventory and processing; I do not claim direct independent inspection of its reserve equation in this revision. Rationing and ATP publisher records support the established methodological context. Nuclear-data and other references retained from earlier verified editions have their access status identified in `evidence/source_register.csv`. The present check does not claim an exhaustive review of every newest fuel-cycle publication.

Shipper and receiver assays require a common reference time and a covariance-aware measurement difference. In-transit radioactive survival is represented in the numerical mass transfer. Daughter helium and chemical purity require separate assay/cleanup evidence; no numerical impurity threshold or license limit is inferred here. The simulated qualification hold is a receiver-level eligibility event. It does not model a specific physically quarantined batch, an observed transport failure, or a regulatory refusal frequency.

S6.2 Executed checks

The following 64 checks were executed. Their tolerances evaluate equations and software behavior, rather than the physical truth of the synthetic inputs. Some checks are algebraic identities; others compare independently assembled numerical systems. The machine-readable table includes any recorded residual. The public archive does not treat the count of checks as a count of scientific discoveries.

ID	Executed check	Scope
1	<code>half_life_day_year_conversion</code>	Same decay constant in day and year units.
2	<code>fuel_supply_flow_basis</code>	b/beta, not burned mass per day.
3	<code>reserve_quantity_and_initial_gap</code>	Declared mixed-parameter reserve illustration.
4	<code>zero_decay_threshold_identity</code>	Independent algebra against stationary balance.
5	<code>decay_reduces_tolerable_nonreturn</code>	Loop and maintained stock decay charged once.
6	<code>threshold_sign_0p5</code>	Storage-maintenance balance sign at a multiplier of the threshold.
7	<code>threshold_sign_1</code>	Storage-maintenance balance sign at a multiplier of the threshold.
8	<code>threshold_sign_1p25</code>	Storage-maintenance balance sign at a multiplier of the threshold.
9	<code>threshold_sign_2</code>	Storage-maintenance balance sign at a multiplier of the threshold.

ID	Executed check	Scope
10	receiver_closed_form_matches_independent_ode	Independent four-state right-hand side; 51 times.
11	receiver_closed_form_matches_matrix_exponential	Independently assembled affine matrix.
12	receiver_site_balance_identity	Rejects remain on-site, rather than disappear.
13	published_serial_routing_numbers	Equation-level replay, not full published system validation.
14	alternative_routing_numbers	Outer returns directly to storage in this distinct topology.
15	receiver_replenishment_preserves_floor	All constant-operation intervals checked for extrema.
16	receiver_replenishment_mass_balance	Generated, received, burned, rejected, remaining and decayed atoms reconcile.
17	analytic_receipt_bound_dominates_dense_grid	Independent dense scan agrees with analytical extremum calculation.
18	smaller_receipt_violates_floor	One milligram removed from the minimum required receipt.
19	contract_charges_initial_fill	Initial receiver stock is supplied rather than created.
20	contract_includes_full_horizon	New receiver demands extend to service horizon.
21	transit_decay_receipt_equivalence	Each dispatch supplies the prescribed mass after seven days.
22	perpetual_import_tail_is_disclosed	Positive stationary obligation survives horizon truncation.
23	maintenance_calendar_preserves_balance	Processing remains on during explicitly planned plasma maintenance.
24	upstream_and_pool_total_conservation	Retry state and rejected external flow included.
25	capacity_bounds_both_compartments	Finite pool and upstream buffer capacities.
26	executed_stock_nonnegative	Short deliveries are recorded rather than executed from negative stock.
27	source_retry_state_nonzero	Overflow storage is active, not an unused label.
28	parcel_shared_allowance	Dated obligations and delivered quantities are separate.
29	cash_fence_amount	Zero-future-receipt funding of heel and prior promises.
30	forecast_fence_is_intermediate	Declared lower receipts reduce the current-stock reservation.
31	forecast_event_extrema_match_dense_horizon	Independent 0.001-year horizon scan.
32	lognormal_unbounded_support	No finite sample maximum is a distributional upper bound.
33	new_customer_updates_emergency_target	Reserve metric includes accepted newcomer until its last due date.
34	dispatch_has_no_realized_future_rate_lookahead	Changing only post-year-8 realized rates leaves year-2 fill unchanged.
35	cvar_at_zero_and_full_use	Correct handling of point masses.
36	cvar_tail_mass	A full upper five-percent tranche yields CVaR one.
37	bootstrap_replicates_exceed_999	Upper bootstrap diagnostic, not a distribution-free guarantee.
38	binomial_upper_zero_has_positive_limit	No observed failures does not prove zero risk.
39	cantelli_risk_levels_explicit	Supplementary diagnostics distinguish alternative tail probabilities.
40	date_grid_has_41_dates_per_receiver_class	Fixed equal grid from year 2 to year 12.
41	date_rule_uses_same_total_failure_criterion	No hindsight choice of date-specific risk budget.
42	independent_validation_uses_distinct_seed	Screening and validation histories separately generated.
43	initial_fence_not_claimed_always_unfunded	Actual sample contains funded high-stock paths; no false universal claim.
44	ongoing_contracts_longer_than_first_year	Short versus full-horizon diagnostics are retained.
45	equal_duration_sensitivity_executed	Late starts also compared over a common operating duration.
46	long_hold_sensitivity_lowers_margin	Separate stationary held-bred-stream sensitivity.
47	principal_seeds_match_execution_config	Documented seeds are checked against the module used by the runner.
48	principal_sizes_match_execution_config	Declared sample sizes match the executable principal protocol.
49	fixed_loss_reserve_comparison_crosses_boundary	A declared fixed-loss counterexample separates reserve effect from changing the loss parameter.
50	reserve_margin_change_equals_decay_cost	Changing only maintained stock changes annual usable margin by lambda times the stock change.
51	periodic_receipt_closes_stationary_storage_cycle	The exact periodic receipt returns stationary-loop storage to R at each cadence.
52	batching_and_transit_raise_continuous_lower_bound	Receipt cadence and transport decay are additional to the continuous stock-maintenance lower bound.

ID	Executed check	Scope
53	retained_site_stock_reconciles_all _compartments	Full-site inventory includes rejected stock and both loops, rather than storage alone.
54	retained_site_balance_over_long_horizons	Repeated receipts, burn and decay retain the whole-site balance.
55	usable_capacity_does_not_bound_total_site _inventory	A usable-storage check cannot stand in for a site-wide possession allowance.
56	paired_counterfactual_has_identical_event _calendar	Deleting the new claim does not remove incumbent make-up or upstream-transfer opportunities.
57	paired_calendar_preserves_offered_source _history	The paired contrast changes claims while retaining exogenous offers and review dates.
58	paired_calendar_preserves_site_balance	Calendar matching does not create or destroy material.
59	two_sided_screen_intervals_bracket_observed _rates	Failure to certify and evidence above the criterion are separately identified.
60	declared_grid_above_model_risk_limit	This fixed synthetic grid is above the limit under simultaneous two-sided sampling bounds; other architectures are untested.
61	receiver_parameter_domain_is_enforced	Physical parameter ranges and the distinct-time-constant scope are enforced before computation.
62	reject_invalid_period	A nonpositive duration or receipt cadence is an explicit input error.
63	reject_invalid_duration_days	A nonpositive duration or receipt cadence is an explicit input error.
64	periodic_tail_not_misassigned_to_maintenance _calendar	Continuous-reference tail is labeled separately from an unsolved periodic maintenance tail.

S6.3 Reproduction boundary

The archive provides a complete executable specification for the present results: receiver constants, source laws, random seeds, parcel conventions, risk levels, date grid and independent validation. Rebuilding and compilation commands belong to its README rather than the scientific argument. Code generated for this study is distinct from the public authors' original full ARC/STEP software, which was not executed here. No confidential facility data are used. Published-equation agreement and software verification remain distinct from experimental validation.

S7 Controlled scope checks and interpretation

S7.1 Fixed loss versus normalized loss classes

Table S8 varies reserve at fixed ϵ . By contrast, the main cadence table normalizes ϵ to each reserve's threshold. Both experiments are useful but answer different questions. The numerical margins below are model consequences, not measured fuel-cycle efficiencies. Each trajectory starts with empty loops and empty storage and uses 30-day receipt opportunities over one 365.25-day year.

S7.2 A continuous lower bound and a periodic shipment contract

For stationary loops let $v = (1 - \epsilon)J_I - F$ and $D = \lambda_d R - v > 0$. Starting just after receipt with $S = R + q$, the equation $\dot{S} = v - \lambda_d S$ returns to R after T precisely when $q = (D/\lambda_d)(e^{\lambda_d T} - 1)$. Hence annual batched receipts exceed the continuous lower bound by the factor $(e^{\lambda_d T} - 1)/(\lambda_d T)$. The transit column includes a further seven-day survival correction. These are stationary tails, excluding startup. The continuous-operation expression is kept only as a reference for the annual-maintenance calendar; it is not substituted for that calendar's long-run periodic solution.

Table S8: Controlled reserve comparison at a fixed physical non-return fraction. Positive margin means net usable inflow at the maintained floor; initial filling is still required.

ϵ	R kg	Margin kg/yr	First receipt kg	Year-1 receipts kg
0.00050	0.200	+0.762012	1.5701	1.5701
0.00050	1.944	+0.663894	3.3152	3.3152
0.00075	0.200	+0.051417	1.5789	1.5789
0.00075	1.944	-0.046701	3.3278	3.3707
0.00100	0.200	-0.659177	1.6339	2.2403
0.00100	1.944	-0.757296	3.3859	4.0826

Table S9: Continuing import and dispatch rates for $\epsilon = 10^{-3}$ and $R = 1.944$ kg. All rate columns are kg/year.

Cadence d	Continuous bound	Batched receipts	7-d dispatch
7.0	0.757296	0.757704	0.758521
30.0	0.757296	0.759048	0.759867
90.0	0.757296	0.762569	0.763392

S7.3 Rejected material remains a physical inventory

The following trajectories retain all rejected material on site. Reported total inventory is $O + I + S + W$ at the stated operating age, whereas the capacity screen in the service benchmark addresses usable storage S . The differing columns show why those boundaries must not be interchanged. A real site-wide restriction must additionally include any unmodeled compartments and transportation custody. No legal inventory ceiling is inferred from this table.

Table S10: Retained tritium under the two declared non-return classes, initially empty receiver, 30-day receipts. W is rejected material with no recovery or removal route in this model.

ϵ/ϵ_c	Years	S kg	W kg	Total kg
0.5	1	2.9454	1.0133	5.3397
0.5	5	6.4779	4.5433	12.4022
0.5	10	9.9099	7.9730	19.2639
0.5	14	12.0389	10.1006	23.5204
1.25	1	1.9440	2.5332	5.8582
1.25	5	1.9440	11.3583	14.6833
1.25	10	1.9440	19.9325	23.2575
1.25	14	1.9440	25.2514	28.5763

S7.4 Pairing, sampling classification and parameter validation

The new no-order counterfactual retains the candidate's appointment dates. This matching removes extra incumbent make-up opportunities as a hidden difference between the two histories. The code checks identical event calendars, identical offered-flow integrals and separate conservation. It changes only attribution baselines, rather than actual dispatch in the histories with the candidate contract.

The additional grid classification uses two one-sided Clopper–Pearson bounds per case. The tail probability is $0.01/(2 \times 82)$. A case is supported if its upper bound is at most 1%, above the model criterion if its lower bound exceeds 1%, and otherwise unresolved. Every case in this fixed benchmark is above the criterion. These sampling statements condition on the synthetic generating law; they do not quantify facility-model error. The machine-readable file `screen_classification.csv` records the counts and both bounds.

The receiver code enforces finite parameters, positive residence times and cadence, $0 < \beta \leq 1$, and recovery/non-return fractions in $[0, 1]$. Coincident residence times are explicitly rejected by this particular closed-form implementation; the model then requires the matrix-exponential or limiting implementation. Input validation is an implementation guard, not evidence for a physically achievable parameter combination.

S7.5 Reproduction record

The complete numerical study was rebuilt for this revision. Principal seeds and sample sizes are imported from `source/execution_config.py` and checked against the documented protocol. The executable configuration and documented protocol identify the same executed sampling populations. The author-facing review record retains the correction history. Numerical reproduction remains separate from empirical validation.