# Energy service and peroxide handling in a proteome-constrained red-cell model: exact inventory certificates, a certified turnover plateau, and kinetic limits

16 September 2026

**Abstract**

Whether stored red cells can sustain an ATP-dependent service while consuming an oxidative challenge is usually approached through a metabolic flux envelope. We ask what such an envelope actually decides. Working with the proteome-constrained S7 instance of RBC-GEM under an explicit finite-horizon relaxation, we prove three exact statements about one concrete model. First, a complete source-matrix inventory certificate: for every cumulative extent vector obeying the declared bounds, balances and terminal floors, the weighted combination of Na/K-ATPase service $M$ and two-route peroxide turnover $Q$ obeys $M + 4Q \le w^{\top}(x_0 - \ell) + 25.439042\,T + 4(R_{\mathrm{CYSGLTH}} + R_{\mathrm{LHCYSTIN}} + R_{\mathrm{SSGTHRD}})$, with all 19 620 columns and their decimal coefficients retained; the concrete vector inequality is verified in Lean 4. Second, and in a deliberately restricted diagnostic medium, an exact rational primal witness together with two exact dual certificates shows that the optimal turnover is uniformly enclosed, $6.50048030 \le F(m) \le 6.50048050$ mmol gDW$^{-1}$, for *every* required service $0 \le m \le 1.0345238$, while no feasible point at all has service above $1.0345238944791832$. The certified interval therefore covers all but a window of width $9 \times 10^{-8}$ of the attainable service range, and across it the variation of the optimum is at most $1.93 \times 10^{-7}$, about $3 \times 10^{-6}$ percent of the turnover level. Third, the apparent tradeoff reported when four internal reverse sulfur directions are blocked is not reproduced by closing external sulfur supply: a stoichiometric identity shows that the relevant carriers cancel against an independent reductant, so supply accounting cannot bound repeated turnover. We then show what the envelope does not decide. A returned optimum imports 1000 peroxide units and leaves 993.4995 of them in the terminal pool; minimising terminal peroxide merely diverts 981.05 units through uncounted haemoglobin reactions, and only a subsequent parsimony step removes the gratuitous handling. Finite carrier pools supply the missing constraint, and we prove a sharp finite-time recycling bound, $V(T) \le \frac{ab}{a+b} C_{\mathrm{tot}} T + \frac{a}{a+b}(g(0) - \frac{bC_{\mathrm{tot}}}{a+b})(1 - e^{-(a+b)T})$, which is attained, implies the harmonic average-rate ceiling with an $O(1/T)$ correction, and starts quadratically from a fully oxidised pool. A complementary paired-signal recovery identity gives a necessary consistency test for inferring integrated effective reduction from treated and untreated assays. All results are conditional on the stated model instance; a source-supported theorem of stored-cell oxidative tolerance with preserved ATP-dependent function remains open, and we state precisely which measurements would close it.

# 1 Introduction

## 1.1 The question

A quantitative claim that a preparation of red cells tolerates an oxidative challenge while preserving an ATP-dependent function needs five things: a preparation, an initial inventory, a substrate supply, a rate law, and an endpoint that separates useful peroxide handling from injury. A metabolic model supplies some of these and not others. It can say which integrated combinations of work and peroxide consumption obey its own resource and allocation constraints. It does not by itself say whether finite carrier pools and supported kinetics can realise such a combination on the required timescale, nor whether the resulting trajectory preserves the function of interest.

Keeping those three questions apart is the organising principle of this paper. We take one published proteome-constrained reconstruction of the human red blood cell, add an explicit finite-horizon relaxation, and ask what follows exactly. The answer is a set of conditional exclusion statements that are sharp about the model and silent about the cell, together with an identification of the specific assumptions that would have to be measured before the model could speak about a challenge assay.

> Which service–turnover restrictions follow exactly from one published proteome-constrained red-cell model, and which additional assumptions are needed before its flux envelope can be read as oxidative resilience?

The source reconstruction is RBC-GEM with the enzyme information of its supplementary file S6 and the proteome-constrained model of its supplementary file S7 [1]. The analysis below is an additional finite-inventory use of those inputs. It is not a claim that the source publication established stored-cell challenge tolerance, and no defect in that work is alleged anywhere in this paper.

## 1.2 Main results

Section 2 fixes the relaxation and the observables: $M$ is the cumulative Na/K-ATPase turnover over a horizon $T$, and $Q = 2\xi_{\mathrm{CAT}} + \xi_{\mathrm{GTHP}}$ counts peroxide molecules consumed through catalase and the glutathione-peroxidase column. Our results are the following.

(i) *A complete inventory certificate* (Theorem 3.1). For the explicit S7 matrix, all 19 620 columns, the declared sink closure and every source decimal read as a rational number,

$$M + 4Q \;\le\; w^{\top}(x_0 - \ell) + 25.439042\,T + 4\big(R_{\mathrm{CYSGLTH}} + R_{\mathrm{LHCYSTIN}} + R_{\mathrm{SSGTHRD}}\big),$$

where $w \ge 0$ has 53 nonzero chemical entries and $R_j$ is the negative part of the net cumulative extent of reaction $j$. The exact time coefficient is a rational number below 25.439042, displayed in (5). The concrete finite vector inequality compiles in Lean 4.

(ii) *Recycling defeats supply accounting* (Section 4). Blocking four internal reverse sulfur directions produces a declining numerical service–turnover frontier of slope $-1/4$; that effect is *not* reproduced by closing all 56 identified sulfur-containing exchange imports. The reason is an exact stoichiometric identity: the chemical parts of HCYSTRDX + TRDRY − LHCYSTIN + GTHP − ESTRONEDHY sum to estradiol + $H_2O_2$ → estrone + 2 $H_2O$, with every sulfur, thioredoxin and nicotinamide carrier cancelling. A budget on carrier *uptake* cannot bound carrier *reuse*.

(iii) *A plateau certified over the whole service range* (Theorem 5.4). In a restricted diagnostic medium, one exact rational feasible point and two exact dual certificates give

$$6.50048030 \;\le\; F(m) \;\le\; 6.50048050 \text{ mmol gDW}^{-1} \qquad \text{for every } 0 \le m \le 1.0345238,$$

where $F(m)$ is the maximal turnover subject to service at least $m$, together with the exact ceiling $M \le 1.0345238944791832$ for every feasible point. The certified endpoint is 99.99999% of that ceiling, so the statement covers all but a window of width $9 \times 10^{-8}$ of the attainable range of $m$, rather than a sampled sub-interval, and the variation of $F$ across the certified interval is at most $1.93 \times 10^{-7}$. The general principle behind this is a one-witness lemma (Lemma 5.1): an enclosure valid on an entire interval follows from a single feasible point plus a single global upper certificate, with no grid. That lemma, the shape of the value function and the ceiling argument are verified in Lean 4.

(iv) *Turnover does not constrain peroxide burden* (Section 6). The returned optimum imports 1000 units of peroxide, counts 6.50048 through the two routes and leaves 993.49952 in the terminal pool. Minimising terminal peroxide at fixed service and turnover substitutes 981.05 units of uncounted haemoglobin-peroxide consumption for that accumulation; only a further parsimony step removes the gratuitous handling and brings import down to the counted consumption. A terminal value is not a burden ceiling, and none of these solutions carries a damage endpoint.

(v) *A sharp finite-time carrier bound* (Theorem 7.1). For a closed two-state carrier with total amount $C_{\mathrm{tot}}$, reduced amount $g$, oxidation $v \le ag$ and regeneration $u \le b(C_{\mathrm{tot}} - g)$,

$$V(T) = \int_0^T v \;\le\; \frac{ab}{a+b}\, C_{\mathrm{tot}} T + \frac{a}{a+b}\Big(g(0) - \frac{b\, C_{\mathrm{tot}}}{a+b}\Big)\big(1 - e^{-(a+b)T}\big),$$

and the bound is attained. It requires no terminal condition, it implies the harmonic average-rate ceiling $V(T)/T \le C_{\mathrm{tot}}/(1/a + 1/b)$ with an $O(1/T)$ correction, and for an initially fully oxidised pool it starts as $\frac{1}{2} ab\, C_{\mathrm{tot}} T^2$: regeneration must precede oxidation. This is the smallest additional structure that separates the stoichiometric feasibility of repeated recycling from its kinetic feasibility, and the S7 instance shows that the separation matters, since the certified witness recycles carriers whose modelled initial amount is zero. The inequality is verified in Lean 4.

(vi) *A recovery consistency test* (Section 8). For paired treated and untreated signals with a common initial value and a constant residual-input fraction $r$, the transformed signal $z = (C_s - rD)/((1-r)D_0)$ equals $e^{-\int_0^t k}$ exactly, for arbitrary time-varying oxidative input. Hence $z$ is positive, bounded by one and nonincreasing, which is a falsifiable requirement on paired observations, and an explicit discrepancy bound converts an observation interval into an interval for the integrated effective reduction.

### 1.3 What this paper does not establish

Every statement above is conditional on a declared model instance. None of them establishes an admissible human preparation, a tolerable external dose, a kinetic trajectory, or a clinical threshold. In particular the coefficient 4 in (i) is a certificate weight, not a universal ATP cost of removing a peroxide molecule; the potential $w^\top x$ is a mathematical resource account, not a Gibbs energy; and the flatness in (iii) is a property of one restricted diagnostic medium, not a demonstration that storage energy and oxidative resilience are unrelated in real cells. The original biological target, a source-supported theorem of stored-cell oxidative tolerance with preserved ATP-dependent function, remains open; Section 9 states the smallest extension we think could close it.

### 1.4 Evidence classes

Different results in this paper rest on different kinds of evidence, and we label them explicitly rather than letting the strongest label spread.

| Result | Evidence class | Boundary |
|---|---|---|
| Inventory certificate, Theorem 3.1 | Lean 4 kernel check of the concrete vector inequality, plus an exact source-data audit | XML parsing, reaction-name interpretation and biological admissibility are external premises |
| Uniform enclosure, Lemma 5.1 and Propositions 5.2–5.3 | Ordinary proofs, also verified in Lean 4 | Nothing model-specific |
| Plateau instance, Theorem 5.4 | Exact rational primal and dual audit in Python | Not claimed as a Lean-replayed source certificate |
| Carrier bounds, Theorem 7.1 and corollaries | Ordinary analytic proofs, verified symbolically; the main inequality is also verified in Lean 4 for everywhere-differentiable $g$ | Coefficients $a, b, C_{\mathrm{tot}}$ are not calibrated for any preparation |
| Recovery identity and error bounds, Section 8 | Lean 4 for the identity, the discrepancy bound and the logarithmic interval | No validated intervention model; no reanalysis of donor data |
| Numerical optima, sensitivities and cleanup | Floating-point linear programs | Feasibility of a relaxation, not kinetic realisability |

### 1.5 Relation to existing work

None of the ingredients is individually new, and we claim no new duality theory. Flux balance analysis and its envelope calculations are standard [2]. Non-steady intracellular accounting has direct prior art in unsteady-state FBA, which uses quantitative time-course metabolomics to constrain changing pools and relaxes unmeasured nodes when required [3]; our finite-inventory relaxation uses much less dynamic information and consequently establishes much less about trajectories. Dynamic FBA couples metabolic optimisation to changing extracellular conditions [4], and personalised kinetic red-cell models address a different, more dynamically specified inference problem [5]. Selecting among optimal flux distributions by a parsimony criterion is established practice [6], and thermodynamic or loopless constraints remove certain unsupported internal cycles [7]; neither addresses finite carrier abundance or a substrate-powered recycling route, which is the obstruction we isolate in Section 4. The paired-signal cancellation in Section 8 is an integrating-factor argument of a familiar kind.

What is specific here is the concrete instance and its traceability: a complete exact audit of a named published model, at stated premises, with the decisive inequality verified by a proof kernel; an enclosure that holds uniformly over an entire certified service interval rather than at sampled points; and a diagnosis of which modelling assumptions actually move the answer.

## 2 The model instance

### 2.1 Finite-horizon relaxation

Partition the expanded source matrix into chemical rows $N_{\mathrm{c}}$ and auxiliary allocation rows $N_{\mathrm{a}}$. Let $\xi \in \mathbb{R}^{19620}$ be the vector of net signed cumulative reaction extents over a horizon $T \geq 0$, let $x_0$ be the initial chemical amounts, $\ell$ the required terminal chemical floors, and $L_j, U_j$ the source lower and upper rate bounds. The relaxation is

$$x_0 + N_{\mathrm{c}}\xi \;\geq\; \ell, \qquad N_{\mathrm{a}}\xi = 0, \qquad L_j T \leq \xi_j \leq U_j T, \qquad T \geq 0. \tag{1}$$

The expanded S7 instance has 19 620 columns and 10 411 rows, of which 2 157 are chemical and 8 254 are auxiliary. Auxiliary rows encode enzyme allocation and proteome budgets; they are algebraic relations, not inventories of ordinary chemicals, and must never be described as

terminal amounts of metabolites. Following the model experiments of the source-model audit, imports through every sink reaction except the three ionic sinks for sodium, potassium and calcium are closed, while exports remain allowed and the exchange feed bounds retain their source values integrated over $T$.

Two features of (1) deserve emphasis because the conclusions depend on them. Terminal floors are necessary conditions only: they constrain neither intermediate availability nor any rate law. And signed extents are equivalent to splitting reversible directions, so using them discards no sign or bound information.

The principal diagnostic family takes $T = 1$ hour, zero terminal floors and zero initial chemical stocks. This is an algebraic diagnostic family and not a viable initial red-cell state; in particular it can admit balanced turnover without the carrier inventory needed to start it, which is exactly the gap quantified in Section 7.

### 2.2 Observables

Write

$$M = \xi_{\mathrm{NAKT}}, \qquad Q = 2\xi_{\mathrm{CAT}} + \xi_{\mathrm{GTHP}}, \qquad R_j = \max\{-\xi_j, 0\}. \tag{2}$$

The Na/K-ATPase column consumes one ATP per turnover; catalase consumes two peroxide molecules per extent and the glutathione-peroxidase column consumes one, and both source directions are nonnegative. Thus $Q$ is a peroxide equivalent through two specified routes. It is not total peroxide disappearance, not total antioxidant capacity, and not a damage-free challenge dose. The reverse quantity $R_j$ is the negative part of the *net* cumulative extent; if a reaction changes direction inside the interval its gross reverse turnover may exceed $R_j$, so any justified bound on gross reverse turnover also bounds $R_j$, if less sharply.

### 2.3 Units

| Quantity | Unit or convention |
|---|---|
| Chemical rate; horizon | $\mathrm{mmol\,gDW^{-1}h^{-1}}$; hours |
| Chemical amount, net cumulative extent | $\mathrm{mmol\,gDW^{-1}}$ |
| $M$, $Q$ in the one-hour experiments | cumulative amounts in $\mathrm{mmol\,gDW^{-1}}$ |
| Source protein abundance | $\mathrm{nmol\,gDW^{-1}}$; auxiliary columns keep their encoded scales |
| Carrier coefficients $a, b$ | $\mathrm{h^{-1}}$ |
| Concentration | amount divided by a measured solution volume per gDW |

Because $T = 1$ h in the numerical experiments, $M$ and $Q$ are numerically equal to the corresponding average rates; they remain amounts. No compatible water volume is supplied for the peroxide diagnostics below, so a reported terminal peroxide value of 993.5 is 993.5 $\mathrm{mmol\,gDW^{-1}}$ and *not* 993.5 mM. Turning it into a concentration requires a justified compartment volume together with explicit conversions among packed-cell volume, cell water, haemoglobin mass and dry mass.

### 2.4 Two declared media

Results below use two constraint sets. The *original feed envelope* is (1) with the sink closure just described. The *restricted diagnostic medium* additionally closes the 56 identified sulfur-containing exchange imports and the 435 carbon-containing exchange imports, with the named D-glucose, carbon-dioxide and bicarbonate exchanges retained. Membership is decided by the source chemical formulas together with the actual singleton exchange stoichiometry; the two closure sets overlap, so 56 and 435 must not be added. Exact reaction identifiers matter: only the named D-glucose exchange `R_EX_glc__D_e` is retained, and the source's separately encoded glucose anomer exchanges must not be silently merged. Other permitted inorganic exchanges

remain open and some formulas are unknown, which limits what the phrase “restricted medium” can mean. It is a declared diagnostic, not a reconstruction of any experimental buffer.

## 3 A complete finite-inventory certificate

**Theorem 3.1** (Source-model inventory bound)**.** *Let $\xi$ satisfy (1) for the explicit S7 matrix and the original feed envelope of Section 2.4, with $T \ge 0$. Then, with the certificate vectors of Appendix A,*

$$M + 4Q \;\le\; w^{\top}(x_0 - \ell) + 25.439042\,T + 4\big(R_{\mathrm{CYSGLTH}} + R_{\mathrm{LHCYSTIN}} + R_{\mathrm{SSGTHRD}}\big), \tag{3}$$

*where $w \ge 0$ has 53 nonzero chemical entries, among them 2 for ATP, 1 for ADP, 2 for reduced glutathione and 1 for haemoglobin-bound 2,3-DPG.*

*Proof.* Let $c^{\top}\xi = M + 4Q$, so that the only nonzero objective coefficients are $c_{\mathrm{NAKT}} = 1$, $c_{\mathrm{CAT}} = 8$ and $c_{\mathrm{GTHP}} = 4$. The certificate supplies nonnegative chemical weights $w$, signed auxiliary weights $y$ with 290 nonzero entries, and nonnegative supply coefficients $\alpha$ equal to 4 on the three displayed sulfur columns and zero elsewhere. For each column put

$$d_j = c_j + \alpha_j + w^{\top}N_{\mathrm{c},j} - y^{\top}N_{\mathrm{a},j}, \qquad \kappa = \sum_j \max\{d_j L_j,\; d_j U_j\}. \tag{4}$$

Three elementary facts finish the argument. For a signed extent inside its interval and $T \ge 0$ we have $d_j\xi_j \le T\max\{d_jL_j, d_jU_j\}$, by cases on the sign of $d_j$. For a nonnegative supply coefficient, $-\alpha_j\xi_j \le \alpha_j\max\{-\xi_j, 0\} = \alpha_j R_j$. And summing the exact column identities gives

$$c^{\top}\xi = -w^{\top}N_{\mathrm{c}}\xi + y^{\top}N_{\mathrm{a}}\xi + \sum_j (d_j - \alpha_j)\xi_j \;\le\; w^{\top}(x_0 - \ell) + \kappa T + \sum_j \alpha_j R_j,$$

where nonnegativity of $w$ together with the chemical floors of (1) bounds the first term, and the auxiliary balances annihilate the second. Evaluating (4) over all 19 620 source columns with every decimal read as an exact rational number gives

$$\kappa = \frac{1283461673726143619667945455169487587}{5045243750000000000000000000000000} = 25.439041943734505\ldots \;<\; 25.439042, \tag{5}$$

and rounding the coefficient upward proves (3). □

Two aspects of this proof are worth isolating. No optimisation algorithm is assumed correct anywhere: verifying a supplied upper certificate needs neither strong duality nor optimality of any solver output, only the three inequalities above. And the three priced sulfur columns keep their residual $-4$, so their reverse directions remain permitted and appear explicitly in the bound rather than being forbidden. A fourth reaction restricted in an early diagnostic, `HTSULGTHST`, has residual zero in this certificate; restoring its direction while retaining the other three restrictions left both tested endpoint optima unchanged, exactly as the coefficient analysis predicts.

**Corollary 3.2.** *Suppose independently justified budgets give $w^{\top}(x_0 - \ell) \le I$ and $R_{\mathrm{CYSGLTH}} + R_{\mathrm{LHCYSTIN}} + R_{\mathrm{SSGTHRD}} \le S$. Then a service requirement $M \ge m$ forces*

$$Q \;\le\; \frac{I + 25.439042\,T + 4S - m}{4}. \tag{6}$$

*If the right-hand side is negative, the service requirement and the nonnegative-turnover assumption are jointly infeasible under those premises.*

| Declared constraints | Required service $m$ | Optimised $Q$ |
|---|---|---|
| Original feed envelope | 0 to 3.64222 | 46.93125 at all tested points |
| Sulfur exchange imports closed | 0 and 3.64222 | 35.07139 at both endpoints |
| Four internal reverse directions blocked | 0; 3.64222 | 6.35976; 5.44921 |
| Restricted carbon and sulfur feeds | 0; 1.02418 | 6.50048; 6.50048 |
| Same restricted feeds, old demand | 3.64222 | infeasible without any challenge |

Table 1: Numerical optima at a one-hour horizon under four declared constraint sets, in $\mathrm{mmol\,gDW^{-1}}$. These are floating-point optima of the stated relaxation. The internal-direction restriction is a model intervention and is distinct from the external-supply restrictions. Each medium has its own service scale: the restricted medium's service ceiling is 1.0345239, so the old demand 3.64222 is infeasible there for reasons that have nothing to do with a peroxide requirement.

The corollary is a useful exclusion only when $I$ and $S$ have independent support. That the right-hand side decreases in $m$ does not by itself exhibit a biological tradeoff: the inequality may be loose, and the internal reverse budget may itself depend on the protocol. Nor does exact arithmetic upgrade estimated enzyme capacities, an assumed medium or uncertain initial inventories into exact biological facts; it controls only the arithmetic of the stated model. Finally, being below the bound constructs nothing. Theorem 3.1 excludes; it does not produce a trajectory, control a peroxide burden, or show that ATP is preserved throughout the interval.

## 4 Recycling defeats supply accounting

The first discriminating experiment asked whether finite inventory alone exposes a service–turnover tradeoff in the original envelope. It does not. Over the sampled original-direction family the optimal $Q$ is constant at 46.93124973 while the required service ranges over $[0, 3.6422153]$. Blocking four internal reverse sulfur directions (`CYSGLTH`, `LHCYSTIN`, `HTSULGTHST`, `SSGTHRD`) produces instead a frontier of slope approximately $-1/4$, descending from 6.3597605 to 5.4492067 across the same service range. That restriction was a diagnostic intervention in the model. It is not an experimental finding that those directions vanish in stored cells, and the induced tradeoff should not be reported as a property of the reconstruction.

The natural repair is to price sulfur where it enters the cell rather than where it turns over internally. It fails, for a reason worth recording. Closing all 56 sulfur-containing exchange imports lowers the optimum only from 46.93125 to 35.07139, and leaves it unchanged as the service requirement rises (Table 1). A structured replacement certificate that assigns weights to cysteine, homocysteine and hydrosulfide pools does cancel the three targeted residuals exactly, but it changes 34 columns and degrades the rate coefficient from 25.44 to approximately 10 492.57, of which 4000 comes from the extracellular $H_2S$ dissociation column alone, because the candidate weights hydrosulfide and not $H_2S$. Chemically interconverting protonation states must be weighted consistently unless their imbalance is deliberately budgeted.

The deeper obstruction is structural. In a sulfur-closed witness, glutathione regeneration runs through reverse `LHCYSTIN` coupled to `HCYSTRDX` and `TRDRy`, and reverse `ESTRONEDHy` supplies 23.40156 of the 28.17063 NADPH units that `TRDRy` consumes. Summing the chemical parts of

$$\begin{aligned} &\textsc{hcystrdx} + \textsc{trdry} - \textsc{lhcystin} + \textsc{gthp} - \textsc{estronedhy} \\ &\text{gives} \qquad \text{estradiol} + \mathrm{H_2O_2} \longrightarrow \text{estrone} + 2\,\mathrm{H_2O}. \end{aligned} \tag{7}$$

Homocysteine, glutathione, thioredoxin and nicotinamide all cancel; eight auxiliary enzyme terms do not cancel and are retained separately. This is a chemical stoichiometric identity, not a self-contained feasible cycle of the expanded model and not an observed human pathway.

Its content is that a carrier can be reused indefinitely while the net conversion is paid for by a different reductant, so a budget on sulfur *uptake* cannot bound sulfur-carrier *turnover*. Consistently with this reading, removing steroid imports does not restore a tradeoff: the witness simply switches supply routes, with `G6PDH2` and `GND` each contributing 0.53024 to support `TRDRy` at 1.06047 while `GTHOx` contributes 2.19940.

Two cautions. This is not automatically a thermodynamically impossible internal loop, since a carrier cycle with a net external substrate conversion is not an unpowered closed flux loop, and removing cycles indiscriminately would delete legitimate catalysis along with possible artefacts [7]. And the failure of this particular replacement certificate does not show that no better potential exists; it shows that the missing constraint is about carrier abundance and rates, which is what Section 7 supplies.

# 5 A plateau certified over the whole service range

## 5.1 One witness and one certificate suffice

Equal floating-point optima at two service levels suggest a plateau but prove nothing between them. The following elementary lemma removes the temptation to certify a grid: an enclosure valid on an entire interval follows from one feasible point and one global upper bound.

Let $P \subseteq \mathbb{R}^n$ be a feasible set and $M, Q : \mathbb{R}^n \to \mathbb{R}$ two observables. For $m \in \mathbb{R}$ write $P_m = \{\xi \in P : M(\xi) \geq m\}$ and let

$$F(m) = \max\{Q(\xi) : \xi \in P_m\} \tag{8}$$

whenever the maximum exists.

**Lemma 5.1** (Uniform enclosure from one witness)**.** *Suppose $Q(\xi) \leq U$ for every $\xi \in P$ with $M(\xi) \geq 0$, and suppose there is $\xi_h \in P$ with $M(\xi_h) \geq m_h \geq 0$ and $Q(\xi_h) \geq L$. Then for every $0 \leq m \leq m_h$ the set $P_m$ is nonempty and*

$$L \;\leq\; F(m) \;\leq\; U.$$

*Proof.* Fix $0 \leq m \leq m_h$. Since $M(\xi_h) \geq m_h \geq m$, the witness lies in $P_m$, so $P_m \neq \emptyset$ and $F(m) \geq Q(\xi_h) \geq L$. If $F(m)$ is attained at $\xi \in P_m$ then $M(\xi) \geq m \geq 0$, so the upper certificate applies and $F(m) = Q(\xi) \leq U$. □

Lemma 5.1 is verified in Lean 4 as `PlateauEnclosure.uniform_enclosure`, with the feasibility half separated as `feasible_nonempty`. It is deliberately stated for any attained maximum, so no compactness hypothesis is hidden in the notation; in the application $P$ is a nonempty compact polytope, so the maxima exist. Two further elementary facts describe the shape of $F$ and the domain on which the question is meaningful.

**Proposition 5.2.** *$F$ is nonincreasing on its domain. If moreover $P$ is convex and $M, Q$ are linear, then $F$ is concave: for $m = am_1 + bm_2$ with $a, b \geq 0$ and $a + b = 1$,*

$$F(m) \;\geq\; aF(m_1) + bF(m_2).$$

*Proof.* If $m_1 \leq m_2$ then $P_{m_2} \subseteq P_{m_1}$, so $F(m_2) \leq F(m_1)$. For concavity let $\xi_i$ attain $F(m_i)$. Then $a\xi_1 + b\xi_2 \in P$ by convexity, its service is $aM(\xi_1) + bM(\xi_2) \geq am_1 + bm_2 = m$, and its turnover is $aF(m_1) + bF(m_2)$; hence it is a feasible competitor at level $m$. □

**Proposition 5.3.** *If $M(\xi) \leq U_M$ for every $\xi \in P$, then $P_m = \emptyset$ for every $m > U_M$.*

Proposition 5.2 is verified in Lean 4 as `PlateauEnclosure.value_antitone` and `value_concave`, and Proposition 5.3 as `infeasible_above_ceiling`.

Proposition 5.3 is what converts an enclosure on $[0, m_h]$ into a statement about the *whole* question: outside $[0, U_M]$ there is nothing to enclose. The gap between the certified endpoint $m_h$ and the certified ceiling $U_M$ is then the only part of the service axis left undecided, and in the instance below it has width below $10^{-7}$.

## 5.2 The S7 instance

**Theorem 5.4** (Certified plateau and ceiling)**.** *Let $P$ be the restricted-medium polytope of Section 2.4 with $T = 1$ h, zero initial chemical stocks and zero terminal floors, retaining all source bounds, directions and auxiliary balances. Then*

$$6.50048030 \;\le\; F(m) \;\le\; 6.50048050 \text{ mmol gDW}^{-1} \qquad \text{for every } 0 \le m \le 1.0345238, \tag{9}$$

*and $M(\xi) \le 1.0345238944791832$ for every $\xi \in P$, so $P_m = \emptyset$ for $m > 1.0345238944791832$.*

*Proof.* Three exact objects are supplied and checked, and Lemma 5.1 with Proposition 5.3 assembles them.

*A rational witness.* An explicit rational vector $\xi_h$ with 1580 nonzero entries is verified, in exact arithmetic and column by column, against every source lower and upper bound, every chemical terminal floor and every auxiliary equality. Its observables are

$$M(\xi_h) = \frac{4421041890341877}{4273504273504270} = 1.03452380234\ldots,$$

$$Q(\xi_h) = \frac{940018750568624598119532794660 3}{144607583826053727575627416 4820} = 6.500480304680\ldots,$$

so $M(\xi_h) \ge 1.0345238$ and $Q(\xi_h) \ge 6.50048030$.

*A turnover certificate.* Nonnegative chemical weights with 194 nonzero entries and signed auxiliary weights with 128 nonzero entries give, by the argument of Theorem 3.1 with $c = Q$, $\alpha = 0$, $x_0 = \ell = 0$,

$$Q(\xi) \;\le\; U = \frac{81256006217701144379925098514861}{12500000000000000000000000000000} = 6.500480497416092\ldots$$

for every $\xi \in P$, hence in particular whenever $M(\xi) \ge 0$. Exactly eight interval-cost terms are nonzero: four dominant enzyme-allocation terms associated with `GTHOx`, `GND`, `CAT` and `G6PDH2`, and four small residual terms at the corresponding metabolic columns. The small terms are retained; describing the sum as containing only the four dominant terms would be incorrect.

*A service certificate.* The same construction with $c = M$ and a second pair of weight vectors, with 200 nonzero chemical and 251 nonzero auxiliary entries, gives

$$M(\xi) \;\le\; U_M = \frac{25863097361979580707739836 82319}{25000000000000000000000000000000} = 1.0345238944791832$$

for every $\xi \in P$, with exactly two nonzero interval-cost terms, at `R_PYK` and at the pyruvate-kinase enzyme-allocation column `R_ENZDL_enzyme_PYK_total_pc`.

Lemma 5.1 with $m_h = 1.0345238$, $L = 6.50048030$ and $U$ as above gives (9) after rounding $L$ down and $U$ up; Proposition 5.3 gives the emptiness claim. □

**Corollary 5.5.** *For $0 \le m_1 \le m_2 \le 1.0345238$,*

$$0 \;\le\; F(m_1) - F(m_2) \;\le\; U - L \;<\; 1.93 \times 10^{-7} \text{ mmol gDW}^{-1},$$

*which is at most $3.0 \times 10^{-8}$ relative to the turnover level, that is about $3 \times 10^{-6}$ percent.*

*Remark* 5.6. The certified endpoint is 1.0345238 and the certified ceiling is 1.0345238944791832, so (9) covers a fraction 0.99999990... of the exactly certified service range. The residual window of width $9.2 \times 10^{-8}$ is not certified; nothing in this paper asserts what $F$ does there beyond the global upper bound $U$, which holds everywhere.

*Remark* 5.7. Corollary 5.5 is an exclusion of a substantial objective tradeoff *within this model*, not an assertion of exact constancy and not a statement of biological measurement precision. It also does not licence the converse constructions: a requirement $Q > U$ is excluded in $P$, and a requirement $Q \geq q$ with $q \leq L$ together with $m \leq 1.0345238$ is met by the same witness, but neither statement establishes a tolerable external dose or constructs a trajectory realising a prescribed exact value of $Q$.

### 5.3 How the certificate is found, and what decides it

The generator starts from floating-point solver information to identify a useful face and a candidate dual. Sparse rational elimination reconstructs the active equalities, a small numerical program in the remaining nullspace locates an interior candidate, and rational free parameters then determine a complete rational vector. The final all-column rational audit decides feasibility; the success or failure of the discovery program decides nothing. This separation matters in practice as well as in principle: a first attempt that pushed the witness to the extreme corner of the face, requesting $Q \geq 6.50048049$, failed the exact audit by about $4 \times 10^{-15}$ on 287 columns, and the certified statement is the one that survived the audit, not the one the solver reported.

Verifying a saved sparse primal and dual certificate costs one traversal of the relevant matrix entries and bounds, so its arithmetic-operation count is proportional to the sparse matrix size, with rational bit costs depending on numerator and denominator sizes. The proof avoids certifying a dense service grid. It establishes no new complexity result for finding certificates or for solving linear programs in general.

## 6 Turnover does not constrain peroxide burden

The returned restricted-medium optimum has $Q \approx 6.5004805$, peroxide influx 1000 and terminal intracellular peroxide 993.4995195, all as source-normalised cumulative amounts. This follows directly from the peroxide row: almost all imported peroxide simply remains, which the constraint set permits because a terminal lower bound carries no protective upper bound.

A sequential cleanup shows that the accumulation is removable degeneracy rather than a necessary feature of the optimum. Holding service and counted turnover within $10^{-8}$ of their selected values and minimising terminal peroxide drives the terminal pool to zero, but redirects approximately 981.05 units through two uncounted haemoglobin reactions, `HBFE4ORPX` and `METHBPX`, at 490.52428 each. Minimising total absolute flux while retaining that result removes the gratuitous routes and brings the import down to 6.50048, matching the counted consumption exactly (Table 2 and Figure 2A).

Three readings of this should be resisted. The solver did not hallucinate a reaction or violate a supplied constraint: it optimised the problem it was given, and the concern is that the problem leaves crucial outcomes unconstrained. Haemoglobin-peroxide reactions are not intrinsically meaningless; their large use in this particular optimum is gratuitous with respect to the chosen objectives and may be consequential for damage, which is a different and more interesting statement. And absolute-flux minimisation is a selection rule whose result depends on the norm, the weights, reaction splitting and the scaling of the expanded model, especially when chemical and auxiliary columns are mixed; parsimonious flux selection has substantial prior art [6] and no claim of a uniquely justified physiological objective is made here.

| Saved stage | $M$ | $Q$ | influx | terminal | main uncounted consumption |
|---|---|---|---|---|---|
| returned optimum | 0 | 6.5004805 | 1000 | 993.4995195 | none in the peroxide trace |
| minimise terminal peroxide | 1.0241786 | 6.5004805 | 987.5490388 | 0 | 490.5242791 each through `HBFE4ORPX` and `METHBPX` |
| then minimise absolute flux | 1.0241786 | 6.5004805 | 6.5004805 | 0 | the listed gratuitous routes disappear |

Table 2: Peroxide accounting across three saved solutions, in $\mathrm{mmol\,gDW^{-1}}$. The rows are different primal vectors at different service values, obtained by sequential objectives with a $10^{-8}$ tolerance; they are not one vector viewed three ways. Service and turnover values are as recorded in the saved result file.

What a challenge-tolerance claim needs instead is an explicit endpoint. For a compatible dynamic model with intracellular peroxide amount $H(t)$, a complete balance has the form

$$H(t) = H(0) + I_H(t) + P_H(t) - Q_{\mathrm{counted}}(t) - Q_{\mathrm{other}}(t) - E_H(t), \tag{10}$$

with import $I_H$, internal production $P_H$, export $E_H$ and all remaining consuming fates $Q_{\mathrm{other}}$, every term on a consistent amount basis and with the correct stoichiometric multiplicity. Some consuming fates represent damage or irreversible loss of protective capacity. A zero value of $H(T)$ bounds neither $\sup_{t\le T} H(t)$, nor cumulative exposure, nor irreversible injury. A tolerance claim therefore needs a supported burden ceiling, a damage integral or a measured retained function, and the present model experiments supply a numerical threshold for none of these.

Transport data alone do not close this gap either. Human red-cell peroxide permeability has been measured at 37°C [8], but the reported experiments include fresh cells and millimolar exposure, whereas the stored-cell recovery assay considered in Section 8 uses micromolar peroxide. Temperature agreement is not transferability: membrane area, dry-mass normalisation, concentration range and preparation must also match before a permeability coefficient can be used to constrain an amount flux.

# 7 Carrier inventory and finite-time recycling

## 7.1 A consequential rate assumption

In the restricted-medium witness the two-route turnover decomposes as approximately 49.9% catalase, 33.8% glutathione recycling supported by the NADH-dependent `GTHOx` column, and 16.3% recycling supported through the pentose-phosphate and thioredoxin route. These are properties of one model witness, not measured pathway fractions or uniquely necessary allocations. Their importance is that scaling a single capacity moves the answer: at the same absolute service, reducing only the `GTHOx` upper flux bound to one half or to zero lowers the optimum from 6.500480 to 5.400780 and 4.301080 (Figure 2B). Three samples do not establish a globally affine sensitivity law, but they do identify a consequential assumption rather than just another omitted reaction.

The assumption deserves care. The source files map the NADH-dependent `GTHOx` and the NADPH-dependent `GTHOy` to the same glutathione reductase homodimer, `cplx_HOMODIMER_GSR`, and give both forward directions the same estimated effective catalytic rate, $2.19818 \times 10^5\,\mathrm{h^{-1}}$, with the source reaction forms

$$\mathrm{GSSG} + \mathrm{H^+} + \mathrm{NADH} \rightarrow 2\,\mathrm{GSH} + \mathrm{NAD^+}, \qquad \mathrm{GSSG} + \mathrm{H^+} + \mathrm{NADPH} \rightarrow 2\,\mathrm{GSH} + \mathrm{NADP^+}.$$

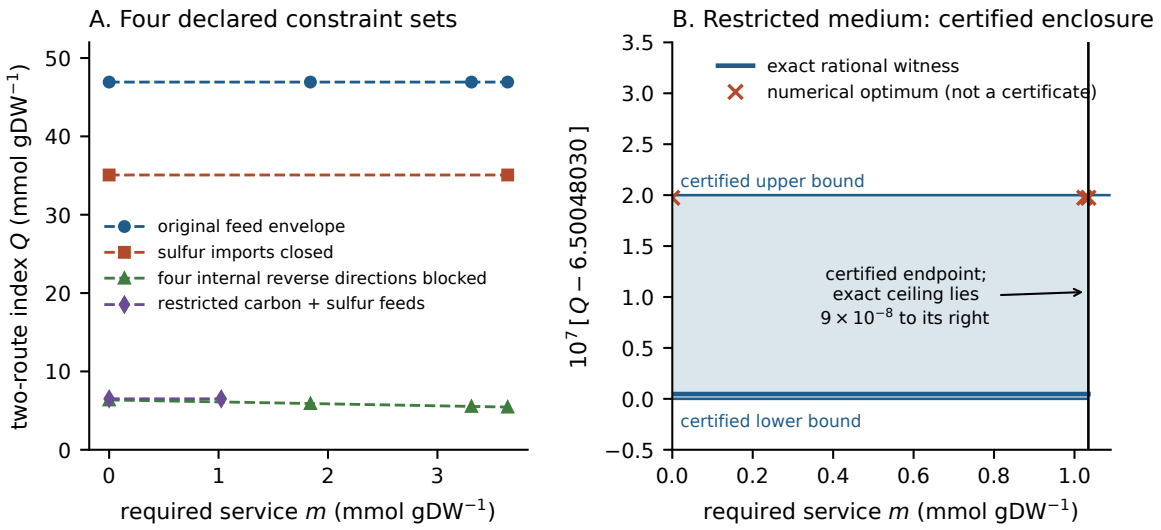


Figure 1: (A) Service–turnover comparisons at a one-hour horizon under the four declared constraint sets of Table 1. Dashed segments join computed values and are not additional certified points; the internal-direction restriction is distinct from the external-supply restrictions, and each medium has its own absolute service scale. (B) The certified enclosure of Theorem 5.4, on an offset ordinate $10^7[Q - 6.50048030]$, so that the certified band is $[0, 2]$. The shaded band is drawn only over the certified service interval: the upper bound also holds outside it, the lower bound is not claimed there. The lower horizontal line is the exact rational witness value and the crosses are floating-point optima, which are not certificates. The vertical line marks the certified endpoint; the exact service ceiling lies $9 \times 10^{-8}$ further right and is not resolvable at this scale. The band is an arithmetic enclosure for one model, not a measured uncertainty interval.

The expanded allocation equations couple the use of that shared complex, so the two individual flux maxima cannot simply be added, and the common estimated effective rate is not a pair of fitted concentration-dependent kinetic laws. Human erythrocyte glutathione reductase has been reported to use both NADPH and NADH [9], and its substrate-dependent behaviour has been studied [10]; deleting the NADH-dependent activity merely because NADPH is the more familiar donor would therefore not be justified. The zero-capacity experiment is a sensitivity test. What is actually needed is supported activity at the nucleotide and GSSG concentrations and the pH of a chosen preparation.

### 7.2 A sharp finite-time bound

Carrier stock constrains turnover even when it is perfectly conserved, and the constraint is invisible to a terminal balance. Consider a closed two-state carrier with total amount $C_{\mathrm{tot}}$, reduced amount $g(t)$, oxidation turnover $v(t)$ and regeneration $u(t)$, and positive first-order coefficients $a, b$:

$$g' = u - v, \qquad 0 \le v \le a\,g, \qquad 0 \le u \le b\,(C_{\mathrm{tot}} - g), \qquad 0 \le g \le C_{\mathrm{tot}}. \tag{11}$$

Take $g$ absolutely continuous and the rates integrable, so the inequalities may hold almost everywhere, and write $V(T) = \int_0^T v$. A simple closed glutathione module uses $g = \mathrm{GSH}/2$ and $C_{\mathrm{tot}} = \mathrm{GSH}/2 + \mathrm{GSSG}$, which counts oxidised-disulfide equivalents.

**Theorem 7.1** (Sharp transient recycling bound)**.** *Under* (11)*, for every* $T \ge 0$*,*

$$V(T) \;\le\; \frac{ab}{a+b}\,C_{\mathrm{tot}}\,T + \frac{a}{a+b}\Big(g(0) - \frac{b\,C_{\mathrm{tot}}}{a+b}\Big)\big(1 - e^{-(a+b)T}\big), \tag{12}$$

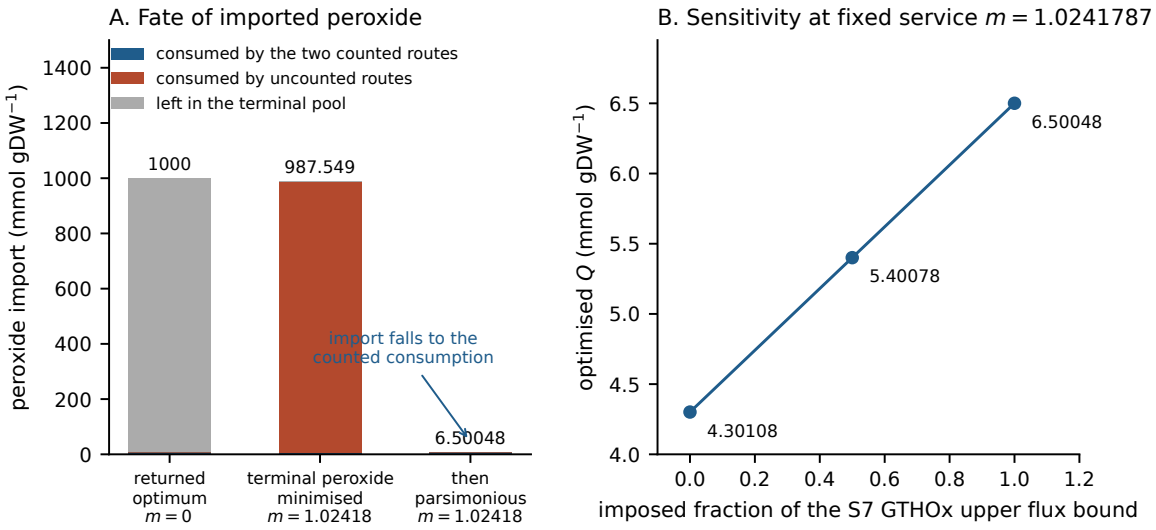


Figure 2: (A) Fate of imported peroxide in the three saved solutions of Table 2. Terminal-only minimisation substitutes uncounted consumption for accumulation; the subsequent parsimony step removes it and import falls to the counted consumption. The three bars are different primal vectors at different service values. (B) Optimised turnover when only the `GTHOx` upper flux bound is scaled, at fixed absolute service 1.0241787. The fractions are imposed model perturbations, not measured activity changes, and the line joining three points is a visual guide.

*and equality holds for the saturated controls* $v = ag$, $u = b(C_{\mathrm{tot}} - g)$.

Theorem 7.1 is verified in Lean 4 as `PlateauEnclosure.carrier_transient_bound`, in the scope recorded in Appendix A; the sharpness statement and the corollaries below are conventional.

*Proof.* Put $k = a + b$ and

$$\lambda(t) = \frac{a}{k}\bigl(1 - e^{-k(T-t)}\bigr),$$

so that $\lambda(T) = 0$, $0 \le \lambda(t) \le a/k < 1$ and $\lambda' = k\lambda - a$. Then, almost everywhere on $[0, T]$,

$$v + (\lambda g)' = (1-\lambda)v + \lambda u + \lambda' g \le (1-\lambda)ag + \lambda b(C_{\mathrm{tot}} - g) + (k\lambda - a)g = \lambda\, b\, C_{\mathrm{tot}},$$

where the inequality uses $v \le ag$ with $1 - \lambda \ge 0$, $u \le b(C_{\mathrm{tot}} - g)$ with $\lambda \ge 0$, and then cancels the $g$ terms because $(1-\lambda)a - \lambda b + k\lambda - a = 0$. Integrating over $[0, T]$ and using $\lambda(T) = 0$ gives

$$V(T) \le \lambda(0)g(0) + b\, C_{\mathrm{tot}} \int_0^T \lambda = \frac{a}{k}\bigl(1 - e^{-kT}\bigr)g(0) + \frac{ab\, C_{\mathrm{tot}}}{k}\Bigl(T - \frac{1 - e^{-kT}}{k}\Bigr),$$

which is (12) after collecting the two $(1 - e^{-kT})$ terms. For sharpness, the saturated controls give $g' = b(C_{\mathrm{tot}} - g) - ag$, hence $g(t) = \frac{bC_{\mathrm{tot}}}{k} + (g(0) - \frac{bC_{\mathrm{tot}}}{k})e^{-kt}$, which stays in $[0, C_{\mathrm{tot}}]$ whenever $g(0)$ does, and $V(T) = a\int_0^T g$ equals the right-hand side of (12). □

**Corollary 7.2.** *Under* (11)*:*

*(a) If* $g(T) = g(0)$ *and* $T > 0$ *then* $V(T)/T \le C_{\mathrm{tot}}/(1/a + 1/b)$*, and in general the right-hand side of* (12) *divided by* $T$ *exceeds that harmonic ceiling by at most* $O(1/T)$*.*

| Initial reduced fraction $\theta$ | Maximum turnover per unit total carrier |
|---|---|
| 0 | 0.45551 |
| 1/3, the saturated stationary fraction | 2/3 |
| 1 | 1.08898 |

Table 3: Bound (12) for an illustrative module with $a = 2\,\mathrm{h}^{-1}$, $b = 1\,\mathrm{h}^{-1}$ and $T = 1$ h, reported per unit total carrier so that no physiological pool is invented. More than one turnover per carrier is possible because the carrier recycles; equal total pools do not give equal short-time turnover when the initial oxidation states differ. These numbers illustrate the theorem and are not fitted red-cell predictions.

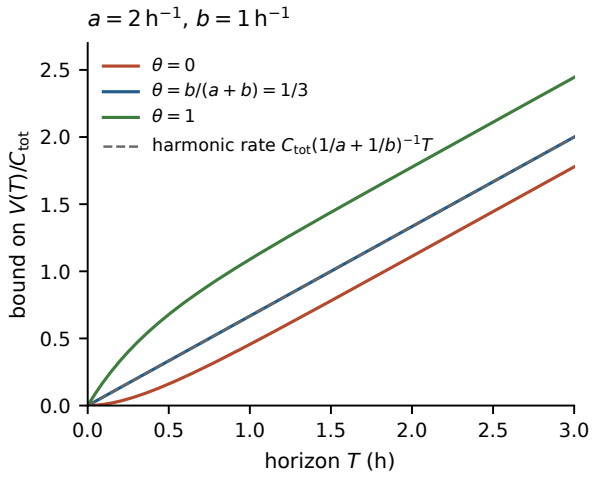


Figure 3: The bound (12) per unit total carrier for three initial reduced fractions, with the harmonic average-rate ceiling for comparison. The curves are upper envelopes for an abstract two-state module with the stated coefficients, not measured red-cell trajectories.

*(b) If* $g(0) = 0$ *then* $V(T) \le \frac{1}{2}ab\,C_{\mathrm{tot}}T^2 + O(T^3)$ *as* $T \to 0$*: the upper envelope starts quadratically, because regeneration must precede oxidation. An initially reduced pool instead has leading allowance* $a\,g(0)T$*.*

*(c) If* $\theta = g(0)/C_{\mathrm{tot}}$ *is known then* (12) *inverts to a necessary lower bound on carrier amount,*

$$C_{\mathrm{tot}} \;\ge\; \frac{V(T)}{\dfrac{ab}{a+b}T + \dfrac{a}{a+b}\Big(\theta - \dfrac{b}{a+b}\Big)\big(1 - e^{-(a+b)T}\big)}, \tag{13}$$

*whenever the denominator is positive.*

*(d)* $C_{\mathrm{tot}} = 0$ *forces* $g \equiv u \equiv v \equiv 0$*, so no positive turnover is possible for any finite* $a, b$*.*

Part (a) is the version of the bound that appears when a terminal balance is imposed; the point of (12) is that no terminal condition is needed, since a cell may legitimately draw down a reduced pool during a short experiment. Part (d) is the sharpest contrast with the stoichiometric relaxation: a balanced stoichiometric model can assign equal positive rates to opposing carrier reactions while the carrier amount is zero at every sampled time, and adding more balance time points alone does not remove that artefact. Pool-dependent rate constraints are the missing condition, which is precisely why the certified witness of Theorem 5.4, with zero modelled initial carriers, is not evidence of kinetic feasibility.

### 7.3 What this costs the model witness

For the returned `GTHP` extent of approximately 3.2598729 $\mathrm{mmol\,gDW^{-1}}$ over one hour, a balanced closed-carrier explanation requires $C_{\mathrm{tot}}/(1/a + 1/b) \geq 3.2598729\ \mathrm{mmol\,gDW^{-1}h^{-1}}$ by Corollary 7.2(a); if the pool is not restored, (12) or (13) applies instead. This is a necessary diagnostic, not an assertion that stored cells have zero carrier. It also suggests a concrete measurement set: the total relevant carrier amount, its initial oxidation state, and supported bounds on regeneration and oxidation at the relevant substrate concentrations. Mixed disulfides, synthesis, export and additional carrier states must either be included or justify the closed-module approximation, and the measurements must constrain instantaneous rate bounds rather than report a fitted average with no envelope interpretation. No source-supported values of $a$, $b$ or $C_{\mathrm{tot}}$ are asserted here for any stored preparation.

## 8 A complementary recovery consistency test

The capacity analysis identifies rate information that a stoichiometric envelope does not provide. Recovery assays are a possible source of such information, but their interpretation depends on how a treatment changes oxidative input and effective reduction. We therefore examine a paired-signal model and derive a necessary consistency test. This section specifies observational requirements; it does not calibrate the S7 instance, and the two analyses do not share a validated biological preparation.

### 8.1 Identity

Let $D$ be an untreated signal and $C_s$ a treated signal, with a common positive initial value $D_0$, a single constant residual-input fraction $0 \leq r < 1$, and

$$D' = q - kD, \qquad C_s' = rq - kC_s, \qquad D(0) = C_s(0) = D_0, \tag{14}$$

where the oxidative input $q$ and the effective reduction $k$ may both vary with time. Put $K(t) = \int_0^t k$ and $z(t) = (C_s(t) - rD(t))/((1-r)D_0)$.

**Proposition 8.1.** *Under (14), $z(t) = e^{-K(t)}$ for all $t$. If in addition $k \geq 0$ on $[0, t_2]$ then*

$$0 < z(t) \leq 1 \quad \textit{and} \quad z(t_2) \leq z(t_1) \qquad (0 \leq t_1 \leq t_2). \tag{15}$$

*Proof.* Set $Y = C_s - rD$. Subtracting the two equations cancels the oxidative input and gives $Y' = -kY$, so $(e^K Y)' = 0$ and $e^K Y$ is constant. The common initial value gives $Y(0) = (1-r)D_0$, whence the identity; the two consequences are immediate from monotonicity of $K$ and positivity of the exponential. □

The identity conditionally identifies *integrated effective reduction*. It does not identify an intrinsic enzyme constant, and it makes no assumption of constant oxidative input.

### 8.2 Discrepancy, observation error and the test

Suppose instead that the treated signal obeys $C_s' = rq + \eta - (k+\delta)C_s$, and set $e = \eta - \delta C_s$. Variation of constants gives $z(t) - e^{-K(t)} = \frac{1}{(1-r)D_0}\int_0^t e^{-(K(t)-K(s))}e(s)\,\mathrm{d}s$, and when $k \geq 0$ the kernel lies in $(0, 1]$, so

$$\left|z(t) - e^{-K(t)}\right| \;\leq\; \Gamma(t) := \frac{\int_0^t |\eta(s) - \delta(s)C_s(s)|\,\mathrm{d}s}{(1-r)D_0}. \tag{16}$$

This allows both imperfect proportional suppression and a treatment effect on effective reduction; a high treatment occupancy does not by itself bound either term for every source of measured

oxidation. If the transformed observation error is at most $\varepsilon$, put $\zeta = \Gamma + \varepsilon$, $L_z = \hat{z} - \zeta$ and $U_z = \min\{1, \hat{z} + \zeta\}$. Provided $0 < L_z \le U_z$,

$$-\log U_z \;\le\; K(t) \;\le\; -\log L_z. \tag{17}$$

An empty interval means the stated model and error assumptions are incompatible. A nonpositive lower endpoint supplies no finite upper bound on $K$, which is a failed identification margin and not evidence of an infinite biological rate. The inference is sensitive near zero and as $r$ approaches one, because the signal difference is divided by $1 - r$.

Two practical points govern how the test is applied. First, preserve the raw paired values and their common initial peak $P$: a normalised untreated signal $D/P$ and a ratio $C_s/D$ share a denominator, their product being $C_s/P$, so treating them as independent observations invents information and discards correlated normalisation uncertainty. For fixed $r$ and raw absolute errors bounded by $\epsilon_C, \epsilon_D, \epsilon_P$ with $\epsilon_P < P$, a conservative transformed-error bound is

$$\epsilon_z \;\le\; \frac{\epsilon_C + r\epsilon_D}{(1-r)(P-\epsilon_P)} + \frac{Z_{\max}\epsilon_P}{P-\epsilon_P}, \tag{18}$$

provided the true transformed magnitude is at most $Z_{\max}$; this is a deterministic bound, not a confidence interval. Second, use one common $r$ at every time. For a fixed raw pair with $0 \le C_s \le D$,

$$\frac{\partial}{\partial r}\frac{C_s - rD}{(1-r)P} = \frac{C_s - D}{(1-r)^2 P} \le 0,$$

so a specified interval for $r$ gives an exact endpoint range for that pair; choosing a separate favourable value at each observation weakens the model being tested.

The test itself is (15): a later transformed interval lying strictly above an earlier one rejects the model at the stated error budget. For a synthetic illustration take a common normalised initial peak of 1, observations $(D, C_s) = (0.8, 0.5)$ followed by $(0.7, 0.6)$, raw absolute errors at most 0.01, and one common $r \in [0, 0.05]$. The smallest possible later-minus-earlier numerator is $(C_{s,2} - C_{s,1}) + r(D_1 - D_2) \ge 0.08 + 0.08r > 0$, while the denominator is positive and common, so every allowed realisation violates (15). This is an illustrative incompatibility certificate. It is not a finding about any published observation, and if intervention discrepancy is permitted then its propagated interval must be included before a rejection is claimed. Passing the test is only a necessary consistency check: it does not validate the intervention mechanism, the complete kinetics or any future prediction.

### 8.3 Applicability

The stored-cell peroxide recovery experiments that motivate this model used washed stored human cells with micromolar peroxide at 37°C and introduced carbon monoxide after a shared challenge period [11]; a later comparison used a different treatment timing and displayed normalisation [12]. The two must not be merged by assuming a common initial condition, and neither the original matched unnormalised donor signals nor calibrated intervention errors were available to us. The consistency test above is therefore a measurement consequence of the identity, not a reanalysis claiming to resolve those experiments.

## 9 Discussion

### 9.1 What the results mean

The practical finding is not that ATP service and oxidative resilience are unrelated. It is that the tested resource relaxation does not impose the substantial tradeoff one might expect,

that the tradeoff which does appear is an artefact of imposed internal reverse restrictions and is not reproduced by supply accounting, and that the assumptions which actually move the answer are about rates, carrier pools and the fate of the oxidant. That narrows what should be measured next.

The results are useful in different ways to different readers. A metabolic modeller can check a proposed integrated output against an exact certificate and inspect the routes that relax it, provided the same source constraints and units are used. A bioinformatician can separate solver-discovered candidates from checkable rational witnesses and upper bounds. An experimental red-cell researcher gets a concrete list of quantities to distinguish: total carrier amount, initial oxidation state and regeneration rate are three different measurements, and matched raw recovery signals with their own initial normalisation are worth more than an additional treatment-to-control ratio. A bioengineer designing a challenge assay should specify absolute service, external exposure and a damage-related endpoint before comparing preparations, since relative service recalibration and total peroxide disappearance answer different questions.

### 9.2 The next bounded target

The recommended next question is smaller than fitting a new red-cell model and more informative than another unrestricted parameter grid:

> For one documented stored-cell preparation, do measured carrier inventories and supported nucleotide-dependent regeneration bounds exclude the high-turnover restricted-medium witness, while preserving a feasible baseline service condition?

The glutathione and glutathione-reductase module is the place to start, because the sensitivity experiment of Section 7 shows that its assumptions materially change $Q$. Five steps make the question decidable. Choose one preparation and fix storage age, medium, temperature, normalisation, treatment timing and assay compartment, rather than assembling a fictitious donor from incompatible studies. Acquire the quantities that enter the bound: GSH/GSSG and other relevant carrier pools, NADH and NADPH availability, shared reductase activity at supported substrates, and the absolute service requirement, retaining uncertainty as ranges with a stated basis. Insert the smallest supported constraint, such as a justified pool-dependent rate envelope or the integrated bound (12), rather than an arbitrary capacity deletion. Retest the baseline before the challenge, since if service alone becomes infeasible the diagnosis lies in the preparation or the capacity assumptions and not in oxidative competition. And test a held-out time or perturbation, because predicting a withheld recovery segment is more informative than refitting the same two ratios. If the new bound does not move the witness, its slack and the remaining active capacities identify the next consequential assumption; if it excludes every baseline state, the units, pool closure and kinetic envelope are the place to look. Both are informative outcomes.

Adding exposure and a protective endpoint requires compatible data: extracellular peroxide, transport, intracellular burden and a supported functional or injury endpoint, with the matching conditions discussed after (10). A dynamic implementation can track substrate depletion [4] and a calibrated kinetic ensemble can express parameter uncertainty [5], but time discretisation alone, unconstrained ensembles and flux minimisation do not supply the missing biological premise. Balances alone cannot prevent turnover of a zero carrier pool.

### 9.3 Scope statement

The source-model inventory theorem is compiled at its stated conditional scope, and the restricted-medium service–turnover enclosure is exactly rationally certified over the whole

certified service range. The numerical diagnostics and the carrier and recovery arguments explain the limits of interpreting these results physiologically. A source-supported theorem of stored-red-cell oxidative tolerance with preserved ATP-dependent function remains open.

# A Reproducibility and verification scope

## Source data

The primary data are the supplementary files S6 and S7 accompanying Haiman et al. [1]. The S7 XML has SHA-256 `b2703b5a730cf12890034f01df92506de5facd9ecea5751a795e982b88557f64`. All decimal coefficients are interpreted as exact rational numbers throughout; no rounded float is ever a certificate input. A release should preserve the original source identifiers, the row classification, the medium masks and the exact XML version, since a list of reaction counts does not determine the restriction.

## Formal verification and its boundary

The declaration `StoredRedCells.S7Certificate.source_joint_inventory_bound` in `proofs/StoredRedCells/S7Source.lean` states Theorem 3.1 as a concrete inequality between finite vector sums over the source family. The recorded strict receipt reports successful compilation under Lean 4 [13] with the repository's pinned Mathlib [14], warnings treated as errors, no `sorry` and no additional axioms; the proof-file SHA-256 is `9419514870b60a803f4904eac1f31acfad9383b84e4e2ff11d10c48aead96bd1`. The supporting modules separate concerns: `FiniteCertificate.inventory_with_reverse_supplies` is the general weighted-balance inequality with signed extents and reverse charges, `S7Bridge` converts between the integer representation and physical scales, `S7Application.inventory_from_chunks` applies the general result to a family assembled from checked blocks, `S7Data.Part0` through `Part19` hold the concrete data in modules of at most 1024 columns each, and `S7Certificate.source_cost_exact` evaluates the complete residual-cost sum. The recorded integer scales are $W = 403619500$, $S = 10^{16}$ and $B = 5 \times 10^{21}$: chemical weights divide by $W$, stoichiometry by $S$, rate bounds by $B$ and objective and supply fields by $WS$. These are representation choices, not physiological hypotheses.

What the kernel checks is the finite vector arithmetic and the implication. What remains outside it is XML parsing, source-version selection, the interpretation of reaction names as source coordinates, and biological admissibility. The reaction names in (3) identify source coordinates through the external audit and manifest; a separate compiled named-coordinate simplification is not claimed. The recovery identity, the discrepancy bound and the logarithmic interval of Section 8 are compiled separately as `PairedRecovery.normalized_integrated_identity`, `paired_robustness` and `observation_rate_interval`, with global derivative hypotheses on the real line; an interval-only minimisation of those hypotheses is not claimed. Theorem 5.4 is established by exact rational arithmetic in Python, and is *not* claimed as a Lean-replayed source certificate.

The abstract results of Sections 5 and 7 are compiled separately in `proofs/StoredRedCells/PlateauEnclosure.lean` as `feasible_nonempty`, `uniform_enclosure`, `value_antitone`, `value_concave`, `infeasible_above_ceiling` and `carrier_transient_bound`, in one strict compilation under Lean 4.30.0 with warnings treated as errors, no `sorry`, and the standard axioms `Classical.choice`, `Quot.sound` and `propext` only; the proof-file SHA-256 is `c01bafb0d3e3200aba595782a95ca691ee7766c2e749aaf144175db084909a21`. Three scope points should be recorded exactly. The compiled carrier statement assumes that $g$ is differentiable everywhere with $g' = u - v$ and that $u, v$ are continuous, which is stronger than the almost-everywhere hypotheses used in Theorem 7.1; the conventional proof above is the one

that covers the weaker scope. Conversely the compiled statement is more general in one respect: it never uses $0 \le g \le C_{\mathrm{tot}}$, since only the two caps enter the pointwise estimate. Finally, the nonnegativity hypotheses $0 \le v$ and $0 \le u$ are inert for the upper bound, as is $0 \le m_h$ in the enclosure lemma; they are retained in the compiled statements as part of the declared model scope, and the compiled files say so. The sharpness claim in Theorem 7.1, the corollaries and everything in Sections 4–6 are ordinary proofs or numerical records, not compiled statements.

### Exact values for the plateau certificate

The rational witness, the turnover weights and the service weights are deposited in the audit record together with every nonzero entry. The decisive exact values are

$$M(\xi_h) = \frac{4421041890341877}{4273504273504270}, \qquad Q(\xi_h) = \frac{940018750568624598119532794 6603}{144607583826053727575627416 4820},$$

$$U = \frac{8125600621770114437992509851 4861}{125000000000000000000000000000000}, \qquad U_M = \frac{2586309736197958070773983682319}{250000000000000000000000000000000},$$

with decimal summaries 1.03452380234, 6.500480304680002, 6.500480497416092 and 1.0345238944791832, and $U - L < 1.93 \times 10^{-7}$ for the displayed rounded $L = 6.50048030$. The audit reports exact primal feasibility with no inconsistent reconstructed equation, no bound violation and no chemical or auxiliary balance violation, from 2009 selected free variables and 1651 pivots; those labels describe reconstruction stages and are not a claimed nullity. The witness has 1580 nonzero entries, the turnover certificate has 194 chemical and 128 auxiliary weights with 8 nonzero interval-cost terms, and the service certificate has 200 chemical and 251 auxiliary weights with 2 nonzero terms. An earlier, narrower audit of the same medium certified the same band only for $m \le 1.02417$; the present certificate supersedes it and both records are retained.

A verification-only entry point should consume the frozen model, the explicit row classification and the medium mask, align rows and columns by recorded source identifiers while rejecting duplicates or omissions, interpret every decimal and every certificate fraction exactly, fill omitted sparse entries with exact zero, check every bound, floor and equality, recompute $M$ and $Q$ from their named coordinates, check that the dual chemical weights are nonnegative, recompute every dual residual and interval maximum while retaining all nonzero terms, and emit a PASS or FAIL record tied to the input hashes. Checking that array lengths agree is not a proof of semantic alignment; identifier-based verification and a frozen restriction manifest are what address that risk.

### Numerical computations

All floating-point linear programs use single-threaded HiGHS [15] through SciPy [16] on the complete sparse model. The cleanup fixes service and turnover within $10^{-8}$, minimises terminal peroxide, then minimises total absolute flux while preserving that result. The plateau witness search maximises service subject to a turnover floor, and the sensitivity experiment changes only the `GTHOx` upper flux bound. No parameter grid, donor-cohort fit, concentration conversion or measured kinetic validation is implied anywhere. The figures are generated directly from the saved records, and the symbolic identities behind Theorem 7.1 and Corollary 7.2 are checked with a computer algebra system.

## References

[1] Zachary B. Haiman, Alicia Key, Angelo D'Alessandro, and Bernhard O. Palsson. RBC-GEM: A genome-scale metabolic model for systems biology of the human red blood cell.

*PLOS Computational Biology*, 21(3):e1012109, 2025. doi: 10.1371/journal.pcbi.1012109. Supplementary files S6 (enzyme information) and S7 (proteome-constrained model).

[2] Jeffrey D. Orth, Ines Thiele, and Bernhard Ø. Palsson. What is flux balance analysis? *Nature Biotechnology*, 28(3):245–248, 2010. doi: 10.1038/nbt.1614.

[3] Aarash Bordbar, James T. Yurkovich, Giuseppe Paglia, Ottar Rolfsson, Ólafur E. Sigurjónsson, and Bernhard O. Palsson. Elucidating dynamic metabolic physiology through network integration of quantitative time-course metabolomics. *Scientific Reports*, 7:46249, 2017. doi: 10.1038/srep46249.

[4] Radhakrishnan Mahadevan, Jeremy S. Edwards, and Francis J. Doyle. Dynamic flux balance analysis of diauxic growth in Escherichia coli. *Biophysical Journal*, 83(3):1331–1340, 2002. doi: 10.1016/S0006-3495(02)73903-9.

[5] Aarash Bordbar, Douglas McCloskey, Daniel C. Zielinski, Nikolaus Sonnenschein, Neema Jamshidi, and Bernhard O. Palsson. Personalized whole-cell kinetic models of metabolism for discovery in genomics and pharmacodynamics. *Cell Systems*, 1(4):283–292, 2015. doi: 10.1016/j.cels.2015.10.003.

[6] Nathan E. Lewis, Kim K. Hixson, Tom M. Conrad, Joshua A. Lerman, Pep Charusanti, Ashoka D. Polpitiya, Joshua N. Adkins, Gunnar Schramm, Samuel O. Purvine, Daniel Lopez-Ferrer, Karl K. Weitz, Roland Eils, Rainer König, Richard D. Smith, and Bernhard O. Palsson. Omic data from evolved E. coli are consistent with computed optimal growth from genome-scale models. *Molecular Systems Biology*, 6:390, 2010. doi: 10.1038/msb.2010.47.

[7] Jan Schellenberger, Nathan E. Lewis, and Bernhard O. Palsson. Elimination of thermodynamically infeasible loops in steady-state metabolic models. *Biophysical Journal*, 100(3): 544–553, 2011. doi: 10.1016/j.bpj.2010.12.3707.

[8] Florencia Orrico, Ana C. Lopez, Daniela Saliwonczyk, Cecilia Acosta, Ismael Rodriguez-Grecco, Isabelle Mouro-Chanteloup, Mariano A. Ostuni, Ana Denicola, Leonor Thomson, and Matias N. Möller. The permeability of human red blood cell membranes to hydrogen peroxide is independent of aquaporins. *Journal of Biological Chemistry*, 298(1):101503, 2022. doi: 10.1016/j.jbc.2021.101503.

[9] Koji Nakashima, Shiro Miwa, and Katsuyo Yamauchi. Human erythrocyte glutathione reductase. I. Purification and properties. *Biochimica et Biophysica Acta*, 445(2):309–323, 1976. doi: 10.1016/0005-2744(76)90085-1.

[10] Kenny K. Wong and John S. Blanchard. Human erythrocyte glutathione reductase: pH dependence of kinetic parameters. *Biochemistry*, 28(8):3586–3590, 1989. doi: 10.1021/bi00434a065.

[11] Victoria M. Harper, Joo Yeun Oh, Ryan Stapley, Marisa B. Marques, Landon Wilson, Stephen Barnes, Chiao-Wang Sun, Tim Townes, and Rakesh P. Patel. Peroxiredoxin-2 recycling is inhibited during erythrocyte storage. *Antioxidants & Redox Signaling*, 22(4): 294–307, 2015. doi: 10.1089/ars.2014.5950.

[12] Joo-Yeun Oh, Chae Yun Bae, Malgorzata Kasztan, David M. Pollock, Robert T. Russell, Jeffrey Lebensburger, and Rakesh P. Patel. Peroxiredoxin-2 recycling is slower in denser and pediatric sickle cell red cells. *FASEB Journal*, 36(4):e22267, 2022. doi: 10.1096/fj.202200052R.

[13] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.

[14] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381, 2020. doi: 10.1145/3372885.3373824.

[15] Qi Huangfu and J. A. Julian Hall. Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142, 2018. doi: 10.1007/s12532-017-0130-5.

[16] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.