The Saturated-Maintenance Projection and Tau-Maintenance-Suppression Theorem in VERSF

▲ Programme Milestone — Flavour-Magnitude and Mass-Hierarchy Series Gate SMPτ-1 / Saturated-Maintenance Projection, Tau-Branch Maintenance Suppression, Independent Demand Ordering, Non-Insertion Closure, and Mass-Interpretation Firewall

Keith Taylor VERSF Theoretical Physics Programme Flavour-Magnitude and Mass-Hierarchy Series Tier A conditional gate-theorem paper, SMPτ-1 July 2026

General Reader Summary

The electron has two heavier siblings: the muon and the tau. Together they make up the charged-lepton family. Simple counting says each of the three should account for one third of the family. This paper shows why one third need not remain the physical maintenance share once an independently derived, nonuniform maintenance demand is introduced.

The picture behind the result is a budget.

In the VERSF maintenance model studied here, sustaining a branch is represented as requiring finite substrate support — think of it as upkeep. Different branches can have different upkeep costs, and the model gives the substrate only a limited upkeep budget to spend. Whether VERSF itself produces exactly this model is one of the openly declared dependencies of the paper, not something the paper assumes settled.

So the question this paper asks is:

If the total upkeep budget is limited, and the tau is the most expensive branch to keep supported, how much support does the tau actually receive?

The paper does not answer by simply writing down a small number for the tau. That would be assuming the conclusion. Instead, it sets the situation up as an honest optimisation problem — the kind a careful accountant would recognise — and lets the mathematics decide.

The optimisation says: given the budget, support as much of the family as possible, paying each branch's upkeep cost as you go. The solution turns out to have a very rigid shape, and the paper proves that **every** solution has this shape, not just one convenient solution:

The cheap branches get filled first. The budget runs out before the expensive branch is fully covered.

It is the same logic as filling a shopping basket with a fixed amount of money: you buy the cheap essentials completely before you touch the expensive item, and if the money runs out, the expensive item is the one left partly or wholly on the shelf.

Applied to the charged leptons, with the tau assumed to be the most expensive branch, three regimes follow:

1. **Small budget.** The electron and muon are supported first. The tau gets nothing at all.
2. **Middling budget.** The electron and muon are fully supported, and the tau gets partial support — some, but not all.
3. **Full budget.** Everything is supported, and only then does the tau's share climb back to the naive one third.

So whenever the budget genuinely falls short of covering everything, the tau's share of the total support is strictly **less than one third**. The naive counting answer is not wrong — it is simply the answer that belongs to a world with an unlimited budget.

Once the electron and muon are fully supported and the tau is partially supported, the paper also proves that the shortfall is not a technicality that could be made invisibly small. In that regime the gap below one third is pinned between two limits set by how far the budget falls short: the bigger the shortfall in the budget, the bigger the tau's suppression must be, within a calculable window.

Three honest limitations are stated just as prominently as the result:

First, the key physical input is assumed, not yet derived. The theorem needs the tau to be the most expensive branch to maintain. That ordering must come from elsewhere in VERSF; this paper proves what follows *once* it is supplied. It would be circular to justify the ordering by pointing at the tau's observed heaviness — the paper explicitly forbids that move.

Second, nothing disappears. A tau receiving little or no maintenance support in this problem is not a claim that tau particles don't exist, can't be produced, or vanish from the theory. The result concerns one specific, carefully defined quantity — sustained maintenance support — and nothing else.

Third, this is not yet a mass or lifetime calculation. The tau's measured mass, its lifetime, and its coupling to the Higgs field all require additional bridges that this paper deliberately does not build. It says so, lists exactly which bridges are missing, and provides a table of forty-three ways the argument would fail if misused.

The central result, in plain words:

If keeping the tau supported costs more than keeping the electron and muon supported, and the total upkeep budget cannot cover everything, then the tau is necessarily the branch that gets squeezed — its share of maintained support falls below the naive one third, by an amount tied to how far the budget falls short.

Scope and Closure Box

Claim	SMPτ-1 status
A finite charged-lepton maintenance carrier exists	Inherited premise
Electron, muon, and tau projectors are consistently defined	Inherited/conditional
A positive maintenance-load map exists	Conditional premise
The maintenance-demand operator is positive	Closed
Maintenance capacity is finite and independently specified	Conditional premise
Admissible maintenance allocations are positive contractions	Defined
Neutral maintained support is the gate objective	Conditional physical premise
The saturated-maintenance problem has an optimum	Closed
Every optimum saturates the capacity constraint for 0 < C < C_full	Closed
An optimum can be chosen diagonal in the demand basis	Closed
Every optimum has threshold-projector form (full characterisation)	Closed
Lower-demand modes are filled before higher-demand modes	Closed
A dual certificate confirms the threshold structure by complementary slackness	Closed
Strict tau-demand separation forces tau-last ordering	Closed
Tau support vanishes below the complement-saturation capacity	Closed
Partial tau support begins only after the complement is full	Closed
Tau maintenance remains below its census fraction until full maintenance	Closed
Equality with the census occurs only at full maintenance	Closed
The census deficit obeys two-sided capacity control	Closed
The tau contrast is exactly the expectation of Q_τ	Closed
The tau contrast is strictly negative in the saturated regime	Closed
Soft monotone saturation also suppresses tau participation	Closed
Commuting positive benefit-to-demand ordering replaces equal-benefit assumptions	Closed conditionally
The maintained-support value function is concave and piecewise linear	Closed
Basis covariance holds	Closed
Branch-preserving perturbative stability holds	Closed
The suppression sign is independent of tau observations	Required for prediction
The numerical demand spectrum D_M	Open unless inherited

Claim	SMPτ-1 status
The numerical capacity C	Open unless independently derived
The numerical tau suppression magnitude	Open
A tau-mass relation	Not derived
A tau-lifetime or decay-width relation	Not derived
Electroweak and observable matching	Open
Empirical confirmation	Not established

Abstract

A raw charged-lepton census assigns one of three flavour directions to the tau branch and therefore gives the exact unweighted fraction ⅓. That rank ratio does not determine a physical maintenance, persistence, interaction, spectral, mass, or decay fraction.

This paper establishes the Saturated-Maintenance Projection and Tau-Maintenance-Suppression Theorem in VERSF.

The finite charged-lepton maintenance carrier is written

$$\mathcal{H}_\ell = \mathcal{H}_0 \oplus \mathcal{H}_\tau, \quad \dim \mathcal{H}_0 = h, \quad \dim \mathcal{H}_\tau = g,$$

with the physical three-branch case given by $h = 2$, $g = 1$. An independently specified maintenance-load map L_M and positive response metric G_M define the positive maintenance-demand operator

$$D_M = L_M^\dagger G_M L_M \geq 0.$$

A maintenance allocation is represented by a positive contraction R, with $0 \leq R \leq I$. For a finite capacity C, saturated maintenance is defined variationally by

$$R_C \in \arg\max \{ \operatorname{Tr} R : 0 \leq R \leq I, \operatorname{Tr}(D_M R) \leq C \}.$$

The feasible set is a compact convex spectrahedron and the objective is linear, so an optimum exists; the problem is a finite-dimensional linear programme over operator constraints. The paper proves a complete saturated-projection theorem: for $0 < C < C_{\text{full}}$ every optimum saturates the capacity constraint, every optimum may be replaced without loss by one commuting with D_M, and — strictly stronger — **every** optimum is exactly of the form

$$R_C = P_{\{<\lambda_C\}} + X_C, \quad 0 \leq X_C \leq P_{\{=\lambda_C\}},$$

where $P_{\{<\lambda_C\}}$ is the spectral projector below the saturation threshold and X_C is a partial occupancy of the threshold eigenspace with $\operatorname{Tr} X_C$ fixed by C. Once optimality fixes the low-

and high-demand diagonal occupancies at one and zero respectively, the operator inequalities $0 \leq R \leq I$ force every associated off-block component to vanish. A dual-certificate theorem independently confirms the threshold structure through complementary slackness.

Assume strict tau-demand separation,

$$\sup \sigma(D_M|\mathcal{H}_0) < \inf \sigma(D_M|\mathcal{H}_\tau).$$

Let

$$C_0 = \mathrm{Tr}(D_M \Pi_0), \; C_{full} = \mathrm{Tr}\, D_M.$$

The exact tau-last theorem then gives:

$$0 < C \leq C_0 \Longrightarrow \Pi_\tau R_C = 0,$$

while

$$C_0 < C < C_{full}$$

forces full maintenance of the complement and only partial maintenance of the tau sector.

Writing

$$q_\tau(C) = \mathrm{Tr}(\Pi_\tau R_C),$$

the saturated tau fraction is

$$f_\tau(C) = q_\tau(C) / (h + q_\tau(C))$$

in the partial-tau regime. Since $0 < q_\tau(C) < g$,

$$f_\tau(C) < g/(g+h) = f_\tau^{census}.$$

For the physical $1 \oplus 1 \oplus 1$ charged-lepton census,

$$f_\tau(C) < \tfrac{1}{3}$$

for every positive capacity below complete maintenance. In the partial-tau regime the census deficit obeys the quantitative two-sided control

$$h\,(C_{full} - C) / ((g+h)^2\, d_\tau^+) \leq g/(g+h) - f_\tau(C) \leq (C_{full} - C) / ((g+h)\, d_\tau^-),$$

with $d_\tau^- = \inf \sigma(D_M|\mathcal{H}_\tau)$ and $d_\tau^+ = \sup \sigma(D_M|\mathcal{H}_\tau)$, so the suppression sign is accompanied by a capacity-controlled magnitude window bounded away from zero and from census-level tuning alike.

With

$Q_\tau = \Pi_\tau - (g/(g+h))\, I,\ \rho_C = R_C / \mathrm{Tr}\, R_C,$

the departure from census is exactly

$f_\tau(C) - g/(g+h) = \mathrm{Tr}(Q_\tau\, \rho_C) < 0.$

The paper also proves a soft-saturation extension for positive non-increasing spectral profiles, a generalised benefit-to-demand theorem, concavity and piecewise linearity of the maintained-support value function, capacity monotonicity, perturbative sign stability, and a constructive no-insertion theorem showing that the numerical magnitude remains underdetermined unless the demand spectrum and capacity are independently fixed.

SMPτ-1 therefore closes the mathematical passage

finite maintenance capacity → spectral saturation projection → tau-last maintenance ordering → negative tau participation contrast.

It does not by itself derive the tau mass, lifetime, decay width, or Yukawa coupling. Those remain distinct dynamical and observable bridges.

Contents

Appendix A — Minimal algebraic skeleton Appendix B — Worked models Appendix C — General $g \oplus h$ theorem Appendix D — Entropy-regularised soft saturation Appendix E — Open bridge-debt ledger Appendix F — References

---

1. Aim, Inheritance, and Novelty

1.1 Inherited participation discipline

A prior VERSF gate established that a sector census is not automatically a physical participation fraction.

For a declared positive participation operator W_X, the share assigned to a sector projector Π_s is

$f_s^X = \mathrm{Tr}(\Pi_s W_X) / \mathrm{Tr}\, W_X$.

After normalisation,

$\rho_X = W_X / \mathrm{Tr}\, W_X$,

the departure from the raw census is the expectation of the centred sector projector.

SMPτ-1 inherits this discipline and supplies a specific physical mechanism for producing a negative tau-sector contrast.

1.2 The new problem

The new question is not merely:

What fraction of the charged-lepton carrier is labelled tau?

That answer is one branch out of three.

The new question is:

Under a finite maintenance capacity, which charged-lepton directions can be sustainably maintained, in what order are they admitted, and what fraction of the maintained response belongs to the tau branch?

1.3 Why saturation matters

If maintenance capacity were unlimited, every admissible direction could be fully maintained:

R = I.

The participation fraction would then reduce to the census fraction.

A nontrivial hierarchy requires at least one of the following:

1. unequal maintenance demand;
2. unequal closure benefit;
3. finite capacity;
4. a nontrivial response profile;
5. differential scale evolution.

SMPτ-1 studies the first three in their cleanest form.

1.4 Core novelty

The paper's programme-specific contribution is the conditional construction of a nonuniform physical-participation candidate from an explicitly stated finite-capacity optimisation principle.

It proves:

1. finite capacity produces a threshold spectral projector;
2. the projector is not inserted by hand — every optimum, not merely some convenient optimum, has threshold form, with the operator inequalities eliminating all cross-block freedom once optimality fixes the extreme occupancies, and with a convex-duality certificate confirming the structure independently;
3. lower-demand modes are maintained before higher-demand modes;
4. a strictly highest-demand tau branch is maintained last;
5. every genuinely saturated optimum has a tau fraction below census;
6. the census deficit obeys two-sided capacity control, so the sign result is accompanied by a magnitude window whenever the capacity gap and tau spectral edges are fixed;
7. the suppression sign is the negative expectation of one centred contrast operator;
8. the sign can be robust even when the numerical magnitude remains open.

1.5 Standard mathematics and programme-specific content

The finite-dimensional optimisation resembles the classical bathtub or fractional-knapsack principle, viewed here as a linear programme over a spectrahedron. Spectral calculus, trace inequalities, positive contractions, pinching maps, convexity, and perturbation bounds are standard mathematical tools.

The VERSF-specific content is their use to convert an independently derived maintenance-demand ordering into a non-inserted tau-sector suppression theorem, with declared premises, declared failure conditions, and firewalls separating the maintenance result from mass, lifetime, and observable claims.

1.6 Core aim

The core aim is:

To derive the saturated-maintenance projector from a finite-capacity variational principle, characterise the entire optimum set, prove that a strictly highest-demand tau branch is maintained last, and establish that the resulting normalised tau participation is strictly below its raw census value — with the deficit confined to a two-sided capacity-controlled window — until complete maintenance becomes possible.

---

2. Typed Charged-Lepton Maintenance Carrier

2.1 Carrier decomposition

Let the finite charged-lepton maintenance carrier be

$\mathscr{H}_\ell = \mathscr{H}_e \oplus \mathscr{H}_\mu \oplus \mathscr{H}_\tau$.

For the minimal physical census,

$\dim \mathscr{H}_e = \dim \mathscr{H}_\mu = \dim \mathscr{H}_\tau = 1$.

It is useful to combine the non-tau sectors:

$\mathscr{H}_0 = \mathscr{H}_e \oplus \mathscr{H}_\mu$.

Then

$\mathscr{H}_\ell = \mathscr{H}_0 \oplus \mathscr{H}_\tau$, $\dim \mathscr{H}_0 = h$, $\dim \mathscr{H}_\tau = g$.

In the minimal case,

$h = 2$, $g = 1$.

2.2 Projectors

Let

Π_e, Π_μ, Π_τ

be mutually orthogonal branch projectors satisfying

$\Pi_e + \Pi_\mu + \Pi_\tau = I_\ell$.

Define

$\Pi_0 = I_\ell - \Pi_\tau$.

The projectors must be fixed by inherited VERSF branch structure or an independently derived spectral identification. They may not be introduced solely by attaching observed particle names to an arbitrary basis.

2.3 Meaning of the carrier

The maintenance carrier is not automatically:

- the complete electroweak Hilbert space;
- the full one-particle Fock space;
- a mass-eigenstate space;
- a decay-channel space;
- a probability simplex;
- a claim that all chiral or spin degrees of freedom have already been counted.

It is the finite internal carrier on which the declared maintenance problem is posed.

2.4 Raw tau census

The tau census fraction is

$f_\tau^{census} = \mathrm{rank}\ \Pi_\tau / \dim \mathcal{H}_\ell = g/(g+h)$.

For $g = 1$, $h = 2$,

$f_\tau^{census} = \frac{1}{3}$.

This is exact as a rank statement.

2.5 Premise P1 — Carrier consistency

All directions included in $\mathcal{H}_\ell$ must use one common counting convention.

Falsifier: electron, muon, and tau sectors include different hidden spin, chirality, particle–antiparticle, or generation multiplicities.

2.6 Premise P2 — Branch-projector consistency

The tau projector must satisfy

$\Pi_\tau^2 = \Pi_\tau$, $\Pi_\tau^\dagger = \Pi_\tau$, $\mathrm{Tr}\, \Pi_\tau = g$,

and must be fixed by inherited or independently derived VERSF structure.

Falsifier: the tau branch is defined only by a preferred coordinate label, or is identified retrospectively from the observed mass ordering.

2.7 Premise P3 — Positive carrier metric

The maintenance carrier must admit the positive metric required for the demand and allocation operators below.

Falsifier: null or negative-norm directions are included without prior physical reduction.

---

3. Positive Maintenance Demand

3.1 Maintenance-load map

Let

$L_M : \mathscr{H}_\ell \to \mathcal{K}_M$

be the maintenance-load map into a response space $\mathcal{K}_M$.

Let

$G_M \geq 0$

be the positive metric that measures maintenance burden in that response space.

Define

$D_M = L_M^\dagger G_M L_M$.

3.2 Premise P4 — Independent demand construction

The pair (L_M, G_M) must be fixed by prior VERSF structure — closure-maintenance burden, ordered χ-increment demand, localisation stiffness, substrate restoration work, persistence cost, or an independently fixed spectral displacement — and not by the intended tau output.

Falsifier: D_M, or any of its spectral data, is fitted from the desired tau fraction, the observed tau mass, or the observed tau lifetime.

3.3 Theorem 1 — Positive Maintenance-Demand Theorem

Statement.

$$D_M \geq 0.$$

Proof. For every $|\psi\rangle \in \mathscr{H}_\ell$,

$$\langle\psi| D_M |\psi\rangle = \langle L_M \psi| G_M |L_M \psi\rangle \geq 0.$$

Therefore D_M is positive. ■

3.4 Interpretation

For a unit-normalised carrier direction $|\psi\rangle$,

$$\langle\psi| D_M |\psi\rangle$$

is its maintenance demand under the declared response metric.

It is not automatically:

- rest mass;
- energy;
- decay width;
- Yukawa coupling;
- inverse lifetime;
- interaction probability.

Those identifications require separate theorems.

3.5 Strict positivity

For the simplest saturated-capacity theorem, assume

$$D_M > 0.$$

This avoids zero-cost directions that could be maintained without consuming capacity.

The theorem extends to a positive semidefinite operator by fully including the zero-demand kernel before the positive-demand optimisation begins: on $\ker D_M$ the allocation is set to the identity at zero cost, and every statement below applies verbatim to the positive-demand complement.

3.6 Premise P5 — Branch covariance

Assume that the maintenance demand respects the inherited branch decomposition:

$$[D_M, \Pi_a] = 0, \quad a \in \{e, \mu, \tau\}.$$

Then

$$D_M = D_e \oplus D_\mu \oplus D_\tau.$$

In the one-dimensional branch case,

$$D_M = d_e \Pi_e + d_\mu \Pi_\mu + d_\tau \Pi_\tau, \quad d_a > 0.$$

Falsifier: the derived demand operator mixes the branch sectors and no block-diagonalisation theorem or physical spectral branch projector is supplied (Section 12.4).

3.7 Demand ordering

The decisive physical premise is the independently derived ordering

$$d_e \leq d_\mu < d_\tau$$

in the rank-one case.

More generally, define

$$d_0^+ = \sup \sigma(D_M|\mathcal{H}_0), \quad d_\tau^- = \inf \sigma(D_M|\mathcal{H}_\tau), \quad d_\tau^+ = \sup \sigma(D_M|\mathcal{H}_\tau).$$

The strict tau-demand gap is

$$\Delta_\tau = d_\tau^- - d_0^+.$$

SMPτ-1 requires

$$\Delta_\tau > 0.$$

This is formalised as Premise P8 in Section 6.

3.8 What counts as an independent derivation

The tau-demand gap is independently derived only if it follows from prior VERSF structure such as:

- ordered χ-increment demand;
- closure-maintenance burden;
- localisation stiffness;
- substrate restoration work;
- persistence cost;
- independently fixed spectral displacement.

It is not independently derived if d_τ is chosen because the observed tau is heavy or short-lived.

---

4. Finite-Capacity Maintenance Principle

4.1 Maintenance allocation

A maintenance allocation is a Hermitian operator R satisfying

$$0 \le R \le I_\ell.$$

The eigenvalue of R along a maintained mode lies between zero and one.

- $R = 0$: no maintained support;
- $R = I$: full maintained support;
- intermediate eigenvalues: partial or ensemble-averaged maintenance.

The operator R is not asserted to be a quantum density matrix.

4.2 Maintenance cost

The total maintenance cost is

$$\mathcal{C}(R) = \mathrm{Tr}(D_M R).$$

Because $D_M \ge 0$ and $R \ge 0$,

$$\mathcal{C}(R) \ge 0.$$

4.3 Maintenance capacity

Let

$C \geq 0$

be the available closure-maintenance capacity.

An allocation is admissible when

$\mathrm{Tr}(D_M R) \leq C$.

The admissible set is

$\mathcal{K}_C = \{ R : 0 \leq R \leq I, \mathrm{Tr}(D_M R) \leq C \}$.

The capacity C must be fixed independently of the intended tau fraction.

4.4 Neutral maintenance objective

At this gate, each fully maintained carrier direction contributes one unit of maintained support.

The objective is therefore

$\mathscr{M}(R) = \mathrm{Tr}\, R$.

The saturated-maintenance problem is

$R_C \in \arg\max \{ \mathrm{Tr}\, R : R \in \mathcal{K}_C \}$.

Structurally this is a linear programme: the objective is linear and $\mathcal{K}_C$ is a compact convex spectrahedron, the intersection of the operator interval [0, I] with one linear half-space.

4.5 Why this is a saturation principle

If

$C \geq C_{full}, C_{full} = \mathrm{Tr}\, D_M$,

then R = I is admissible and every direction can be fully maintained.

If

$C < C_{full}$,

full maintenance is impossible. The constraint is active, and the system must select which modes receive support.

4.6 Premise P6 — Finite capacity

A nontrivial saturated regime requires

$0 < C < C_{full}$,

with C fixed independently of the intended tau result.

Falsifier: C is infinite, exceeds C_full, or is selected to reproduce a target fraction.

4.7 Premise P7 — Neutral branch value

The objective Tr R assumes that one unit of maintained support has equal gate-level value in every admissible direction.

Falsifier: the dynamics supplies a nontrivial benefit operator that is ignored, or the objective is changed after inspecting the desired ordering.

Section 10 replaces this premise with a general benefit operator and shows that the result survives whenever the tau benefit-to-demand ratio is strictly lower.

4.8 Static character of the principle

The saturated-maintenance principle is a constrained optimisation, not yet a dynamics. This gate derives no time-evolution equation for R, no relaxation mechanism toward the optimum, no stability statement for the optimum under a dynamical flow, no substrate reason why the VERSF response maximises Tr R, and no conservation law producing C. Appendix D supplies a regularised variational model, not a VERSF evolution law. Throughout the paper the result is therefore described as finite-capacity maintenance optimisation; a later gate may derive a relaxation equation whose stable stationary points are exactly the saturated optimisers (debt D17).

---

5. Existence, Reduction, and the Saturated-Maintenance Projection Theorem

5.1 Demand eigenbasis

Let

$D_M = \Sigma_i d_i |i\rangle\langle i|, 0 < d_1 \leq d_2 \leq \cdots \leq d_n$,

where

$n = g + h$.

5.2 Lemma 1 — Existence

Statement. For every $C \geq 0$ the saturated-maintenance problem attains its maximum.

Proof. In finite dimension the set $\mathcal{K}_C$ is closed and bounded, hence compact, and it is convex as the intersection of the convex operator interval $[0, I]$ with a closed half-space. The objective $R \mapsto \operatorname{Tr} R$ is continuous and linear. A continuous function on a nonempty compact set attains its maximum ($0 \in \mathcal{K}_C$, so the set is nonempty). ■

5.3 Lemma 2 — Constraint saturation

Statement. If $0 < C < C_{full}$, then every optimum satisfies

$$\operatorname{Tr}(D_M R_C) = C.$$

Proof. If $R = I$ then $\operatorname{Tr}(D_M R) = C_{full} > C$, so $R = I$ is inadmissible; hence every admissible R satisfies $R \neq I$. Suppose now that $\operatorname{Tr}(D_M R) < C$ for an admissible R. For $\varepsilon \in (0, 1]$ define

$$R_\varepsilon = (1 - \varepsilon) R + \varepsilon I = R + \varepsilon (I - R).$$

Then $0 \leq R_\varepsilon \leq I$, and

$$\operatorname{Tr} R_\varepsilon = \operatorname{Tr} R + \varepsilon (n - \operatorname{Tr} R) > \operatorname{Tr} R,$$

$$\operatorname{Tr}(D_M R_\varepsilon) = \operatorname{Tr}(D_M R) + \varepsilon (C_{full} - \operatorname{Tr}(D_M R)).$$

Choosing ε small enough keeps $\operatorname{Tr}(D_M R_\varepsilon) \leq C$, so R was not optimal (the trace increase is strict because $R \neq I$ forces $\operatorname{Tr} R < n$). Hence every optimum has cost exactly C. ■

5.4 Dephasing map

Define the pinching (dephasing) map in the demand basis:

$$\mathcal{D}_D(R) = \Sigma_i |i\rangle\langle i| R |i\rangle\langle i|.$$

If $0 \leq R \leq I$, then $0 \leq \mathcal{D}_D(R) \leq I$, because pinching is a unital completely positive trace-preserving map. Moreover

$$\operatorname{Tr} \mathcal{D}_D(R) = \operatorname{Tr} R$$

and, since $\mathcal{D}_D$ is self-adjoint for the trace pairing and fixes D_M,

$\mathrm{Tr}(D\_M\ \mathcal{D}\_D(R)) = \mathrm{Tr}(\mathcal{D}\_D(D\_M)\ R) = \mathrm{Tr}(D\_M\ R)$.

5.5 Lemma 3 — Demand-Basis Reduction

Statement. Every admissible allocation can be replaced by one commuting with D_M without changing either its total maintained support or its maintenance cost. In particular, the optimal value is attained on diagonal allocations.

Proof. Apply the pinching map of Section 5.4. It preserves positivity, the upper bound by I, the trace, and the demand-weighted trace. ■

The optimisation therefore reduces to scalar occupancies

$0 \leq r_i \leq 1$,

maximising $\Sigma_i\ r_i$ subject to $\Sigma_i\ d_i\ r_i \leq C$.

5.6 Theorem 2 — Saturated-Maintenance Projection Theorem

Statement. For $0 < C < C\_full$, every diagonal saturated-maintenance optimum fills lower-demand modes before higher-demand modes: there exists a threshold $\lambda\_C \in \sigma(D\_M)$ such that

$r_i = 1$ whenever $d_i < \lambda\_C$, $r_i = 0$ whenever $d_i > \lambda\_C$,

with partial occupancy possible only at $d_i = \lambda\_C$. Equivalently, among allocations commuting with D_M,

$R\_C = P\_\{<\lambda\_C\} + X\_C$, $P\_\{<\lambda\_C\} = \mathbf{1}\_\{[0, \lambda\_C)\}(D\_M)$, $0 \leq X\_C \leq P\_\{=\lambda\_C\}$,

and the threshold occupancy is fixed by capacity saturation:

$\mathrm{Tr}\ X\_C = (\ C - \mathrm{Tr}(D\_M\ P\_\{<\lambda\_C\})\ ) / \lambda\_C$.

Proof. By Lemma 3 take R diagonal. Suppose an optimum has $r_i < 1$ and $r_j > 0$ for two modes with $d_i < d_j$. Perform the cost-preserving transfer

$r_i \mapsto r_i + \varepsilon$, $r_j \mapsto r_j - \varepsilon\ d_i/d_j$,

for $\varepsilon > 0$ small enough that both occupancies stay in [0, 1]. The maintenance cost changes by

$\varepsilon\ d_i - \varepsilon\ (d_i/d_j)\ d_j = 0$,

while the total maintained support changes by

$\varepsilon (1 - d_i/d_j) > 0$,

since $d_i < d_j$. This contradicts optimality. Hence no mode with strictly larger demand carries occupancy while a mode with strictly smaller demand remains unsaturated. In particular partial occupancy cannot occur at two distinct demand values, so all occupancies equal one below a single threshold demand λ_C, zero above it, and lie in [0, 1] on the threshold eigenspace. Lemma 2 fixes the threshold occupancy by $\mathrm{Tr}(D_M R_C) = C$. ■

5.7 Theorem 3 — Optimum-Set Characterisation Theorem

Statement. For $0 < C < C_{full}$, **every** optimum R of the saturated-maintenance problem — not only the diagonal representatives — has the exact form

$$R = P_{<\lambda_C} + X,\ 0 \le X \le P_{=\lambda_C},\ \mathrm{Tr}\, X = (C - \mathrm{Tr}(D_M P_{<\lambda_C})) / \lambda_C,$$

with

$$P_{<\lambda_C} R = R P_{<\lambda_C} = P_{<\lambda_C},\ P_{>\lambda_C} R = R P_{>\lambda_C} = 0.$$

All components of R connecting distinct spectral blocks of D_M vanish identically. The optimum is unique if and only if either the threshold eigenspace is one-dimensional or the threshold occupancy is 0 or its maximal value; otherwise the residual freedom is exactly the choice of the positive contraction X on the threshold eigenspace with fixed trace. (Uniqueness here is of the operator R, not of the presentation: when the threshold occupancy is zero, the same operator admits the descriptions $P_{<\lambda} + 0$ and $P_{\le\lambda'}$ for the adjacent demand value λ', and the pair (λ_C, X) is a bookkeeping convention rather than an invariant of the optimum.)

Proof. Let R be any optimum. By Lemma 3, $\mathcal{D}_D(R)$ is feasible with the same objective value, hence a diagonal optimum, so by Theorem 2 its diagonal entries satisfy $\langle i|R|i\rangle = 1$ for $d_i < \lambda_C$ and $\langle i|R|i\rangle = 0$ for $d_i > \lambda_C$.

Take a mode with $\langle i|R|i\rangle = 1$. Then $(I - R) \ge 0$ has vanishing (i, i) entry; for a positive operator, a vanishing diagonal entry forces the entire corresponding row and column to vanish, because $\langle i|(I - R)|i\rangle = 0$ implies $(I - R)^{1/2}|i\rangle = 0$ and hence $(I - R)|i\rangle = 0$. Therefore $R|i\rangle = |i\rangle$.

Take a mode with $\langle j|R|j\rangle = 0$. The same argument applied to $R \ge 0$ gives $R|j\rangle = 0$.

Hence the low-demand spectral subspace consists of eigenvectors of R with eigenvalue one, the high-demand subspace of eigenvectors with eigenvalue zero, and every off-block matrix element of R vanishes. What remains is the compression $X = P_{=\lambda_C} R P_{=\lambda_C}$, a positive contraction on the threshold eigenspace, and Lemma 2 fixes Tr X. Conversely, every operator of the stated form is feasible with cost C and maximal trace, hence optimal. The uniqueness classification reads off directly. ■

5.8 Canonical threshold representative

If the threshold eigenspace has rank $m > 1$ and the threshold occupancy is strictly between its extremes, X is not unique. A symmetry average within the degenerate eigenspace gives the canonical representative

$X_C = \alpha_C P_{\{=\lambda_C\}}, 0 \le \alpha_C \le 1.$

The trace $\mathrm{Tr}\, R_C$ is independent of the choice of X within the optimum set. The sector aggregates $\mathrm{Tr}(\Pi_\tau R_C)$ and $f_\tau(C)$ are independent of X **provided no demand eigenvalue is shared between $\mathscr{H}_0$ and $\mathscr{H}_\tau$**: Premise P5 alone permits a threshold eigenspace straddling both sectors, in which case X can weight the two parts differently and $\mathrm{Tr}(\Pi_\tau X)$ is not fixed. A minimal counterexample under P5 alone: on a two-dimensional carrier with $D_M = I$ and $C = 1$, every positive contraction X with $\mathrm{Tr}\, X = 1$ is optimal, yet $\mathrm{Tr}(\Pi_\tau X)$ ranges over the whole interval $[0, 1]$ even though $[D_M, \Pi_\tau] = 0$. Under Premise P8 the sectors have disjoint spectra, every demand eigenspace — in particular the threshold eigenspace — lies entirely inside one sector, Π_τ then commutes with each spectral block in which the freedom lives, and all aggregate quantities used from Theorem 4 onward are X-independent. Every invocation of this fact below occurs under P8.

5.9 Meaning of "projection"

The saturated solution is exactly a spectral projector except possibly at the capacity threshold, where a partial occupancy may occur.

This is not an imposed projector. It is the complete solution set of the finite-capacity variational problem: optimality pins the extreme occupancies, and positivity together with the contraction bound then eliminates every cross-block component.

5.10 Theorem 3b — Dual Threshold Certificate

The threshold structure admits an independent confirmation by convex duality.

Statement. For $\nu \ge 0$ define the Lagrangian

$\mathscr{L}(R, \nu) = \mathrm{Tr}\, R - \nu\, (\mathrm{Tr}(D_M R) - C) = \nu C + \mathrm{Tr}((I - \nu D_M) R)$

and the dual function

$g(\nu) = \sup_{\{0 \le R \le I\}} \mathscr{L}(R, \nu) = \nu C + \mathrm{Tr}\, (I - \nu D_M)_+,$

where $(\cdot)_+$ denotes the positive part. Then for $0 < C < C_{\mathrm{full}}$:

1. **Weak duality.** $\mathrm{Tr}\, R \le g(\nu)$ for every feasible R and every $\nu \ge 0$.
2. **Complementary slackness.** With $\nu_C = 1/\lambda_C$, any operator R satisfying $R = I$ on the spectral subspace $\{ I - \nu_C D_M > 0 \}$, $R = 0$ on $\{ I - \nu_C D_M < 0 \}$, partial occupancy

only on $\{ I - \nu_C D_M = 0 \}$, and $\mathrm{Tr}(D_M R) = C$ is primal feasible and satisfies $\mathrm{Tr}\, R = g(\nu_C)$.

3. **Strong duality and certification.** Consequently the threshold solutions of Theorem 3 are optimal, g attains its infimum over $\nu \geq 0$ at ν_C, and the pair (R_C, ν_C) is a self-contained optimality certificate.

Proof. For weak duality, take any feasible R and any $\nu \geq 0$. Since $\mathrm{Tr}(D_M R) \leq C$ and $\nu \geq 0$,

$\mathrm{Tr}\, R = \mathrm{Tr}((I - \nu D_M) R) + \nu\, \mathrm{Tr}(D_M R) \leq \mathrm{Tr}((I - \nu D_M) R) + \nu C$,

and for a Hermitian operator A the supremum of $\mathrm{Tr}(AR)$ over $0 \leq R \leq I$ is $\mathrm{Tr}\, A_+$, attained by the spectral projector onto the strictly positive part (with arbitrary occupancy on ker A). Hence $\mathrm{Tr}\, R \leq \nu C + \mathrm{Tr}(I - \nu D_M)_+ = g(\nu)$.

For complementary slackness, set $\nu_C = 1/\lambda_C$ and let R be as described. On modes with $d_i < \lambda_C$ the coefficient $1 - \nu_C d_i$ is strictly positive and the occupancy is one, so the block contributes its full positive part; on the threshold eigenspace the coefficient vanishes and the occupancy is irrelevant to $\mathrm{Tr}((I - \nu_C D_M)R)$; on modes with $d_i > \lambda_C$ the coefficient is strictly negative and the occupancy is zero. Therefore

$\mathrm{Tr}((I - \nu_C D_M) R) = \mathrm{Tr}\, (I - \nu_C D_M)_+$,

and with $\mathrm{Tr}(D_M R) = C$ exactly,

$\mathrm{Tr}\, R = \nu_C C + \mathrm{Tr}\, (I - \nu_C D_M)_+ = g(\nu_C)$.

Since every feasible allocation is bounded by $g(\nu_C)$ and R attains it, R is optimal and ν_C minimises g. By Theorem 3 the operators so described are exactly the saturated optima, so the primal–dual pair certifies the entire threshold structure without re-deriving it. ■

The dual multiplier $\nu_C = 1/\lambda_C$ is the marginal maintained support per unit capacity, matching the slope law of the value function proved in Theorem 11: where V is differentiable — equivalently, away from cumulative-cost breakpoints — $V'(C) = \nu_C = 1/\lambda_C$. A degenerate threshold eigenspace still gives this definite slope throughout the interior of its filling interval; the derivative is nonunique only at the breakpoints themselves.

---

6. Tau-Demand Separation and the Tau-Last Theorem

6.1 Premise P8 — Strict tau-demand separation

Assume

$d_0^+ < d_\tau^-$,

that is,

$\Delta_\tau = d_\tau^- - d_0^+ > 0$.

Every electron–muon maintenance mode is therefore strictly cheaper than every tau maintenance mode.

Falsifier: the tau spectrum overlaps or falls below the complement spectrum, or the gap is smaller than uncontrolled corrections (Section 13).

6.2 Complement-maintenance capacity

Define

$C_0 = \mathrm{Tr}(D_M \Pi_0)$.

This is the capacity required to maintain the full non-tau complement:

$R = \Pi_0$.

Define

$C_{full} = \mathrm{Tr}\, D_M = C_0 + \mathrm{Tr}(D_M \Pi_\tau)$.

6.3 Theorem 4 — Tau-Last Maintenance Theorem

Statement. Under Premises P1–P8:

1. for $0 < C \leq C_0$, every optimum has no tau support: $\Pi_\tau R_C = 0$;
2. for $C_0 < C < C_{full}$, every optimum fully maintains the complement, $\Pi_0 R_C = \Pi_0$, and only partially maintains the tau sector, $0 < \mathrm{Tr}(\Pi_\tau R_C) < g$;
3. for $C \geq C_{full}$, full maintenance $R_C = I$ is admissible and optimal.

Proof. By Theorem 3 every optimum is determined by a threshold $\lambda_C \in \sigma(D_M)$ and a threshold block.

For $0 < C \leq C_0$: the cumulative cost of fully occupying all complement modes is C_0, so with $C \leq C_0$ the threshold lies inside the complement spectrum or at its upper edge d_0^+. By strict separation every tau eigenvalue exceeds $d_0^+ \geq \lambda_C$, so the tau sector lies in the region $P_{>\lambda_C}$ where Theorem 3 forces R_C to vanish, except in the boundary case $\lambda_C = d_0^+$ where the threshold eigenspace is still contained in $\mathscr{H}_0$ by Premise P5 and strict separation. Hence $\Pi_\tau R_C = 0$.

For $C_0 < C < C\_full$: the capacity strictly exceeds the cost of full complement maintenance, so the threshold lies in the tau spectrum, $\lambda\_C \in \sigma(D\_M|\mathscr{H}\_\tau)$. Theorem 3 then forces every complement mode into the fully occupied block, $\Pi_0 R\_C = \Pi_0$, and forces the tau occupancy to be positive but incomplete because $C < C\_full$. Hence $0 < \mathrm{Tr}(\Pi\_\tau R\_C) < g$.

For $C \geq C\_full$: $R = I$ is admissible and attains the maximal possible trace $g + h$. ■

6.4 Last-maintained, first-released principle

As capacity rises from zero, tau support appears only after all lower-demand complement directions are maintained.

Equivalently, as capacity falls from full maintenance, tau support is the first branch support to be released.

This ordering is a theorem of the optimisation, not an interpretive metaphor.

---

7. Exact Tau-Maintenance-Suppression Theorem and Quantitative Deficit Bound

7.1 Saturated tau fraction

Whenever $\mathrm{Tr}\, R\_C > 0$, define

$f\_\tau(C) = \mathrm{Tr}(\Pi\_\tau R\_C) / \mathrm{Tr}\, R\_C,\ q\_\tau(C) = \mathrm{Tr}(\Pi\_\tau R\_C)$.

By Section 5.8, $f\_\tau(C)$ is well defined on the whole optimum set.

7.2 Low-capacity regime

For $0 < C \leq C_0$, Theorem 4 gives

$q\_\tau(C) = 0,\ f\_\tau(C) = 0$.

7.3 Partial-tau regime

For $C_0 < C < C\_full$, the complement is fully occupied, so

$\mathrm{Tr}\, R\_C = h + q\_\tau(C),\ 0 < q\_\tau(C) < g$.

Hence

$f_\tau(C) = q_\tau(C) / (h + q_\tau(C))$.

7.4 Theorem 5 — Exact Tau-Maintenance-Suppression Theorem

Statement. Under Premises P1–P8, every genuinely saturated optimum satisfies

$f_\tau(C) < f_\tau^{census} = g/(g+h)$

for all $0 < C < C_{full}$. In particular, equality with census cannot occur at any capacity below full maintenance; the full-capacity limit is treated in Section 11.4.

Proof. For $0 < C \leq C_0$,

$f_\tau(C) = 0 < g/(g+h)$.

For $C_0 < C < C_{full}$,

$f_\tau(C) = q_\tau / (h + q_\tau),\ 0 < q_\tau < g$.

The function $x \mapsto x/(h+x)$ is strictly increasing for $x \geq 0$, so

$q_\tau/(h+q_\tau) < g/(h+g)$.

Thus the tau fraction is strictly below census throughout the saturated regime. ■

7.5 Theorem 6 — Quantitative Deficit Bound

Statement. Write $\kappa = C_{full} - C$. In the partial-tau regime $C_0 < C < C_{full}$, the census deficit obeys the two-sided capacity control

$h \kappa / ((g+h)^2 d_\tau^+) \leq g/(g+h) - f_\tau(C) \leq \min\{ \kappa / ((g+h) d_\tau^-), g/(g+h) \}$,

where the second entry of the minimum is the trivial positivity bound $f_\tau(C) \geq 0$ and sharpens the capacity bound early in the partial regime, where κ is large. In the low-capacity regime $0 < C \leq C_0$ the deficit is maximal, $g/(g+h) - f_\tau(C) = g/(g+h)$. Consequently the census deficit is bounded away from zero on every capacity interval bounded away from C_{full}, and vanishes proportionally to κ as $C \uparrow C_{full}$ for fixed positive tau spectral edges.

Proof. In the partial-tau regime the complement is full, so the unspent full-maintenance cost is carried entirely by the unoccupied tau support:

$C_{full} - C = \mathrm{Tr}(D_M (I - R_C)) = \mathrm{Tr}(D_M (\Pi_\tau - \Pi_\tau R_C \Pi_\tau))$, using Premise P5 and $0 \le \Pi_\tau - \Pi_\tau R_C \Pi_\tau \le \Pi_\tau$. Bounding the tau spectrum from both sides,

$$d_\tau^- (g - q_\tau) \le C_{full} - C \le d_\tau^+ (g - q_\tau),$$

so

$$(C_{full} - C) / d_\tau^+ \le g - q_\tau \le (C_{full} - C) / d_\tau^-.$$

The deficit is, by Appendix C,

$$g/(g+h) - f_\tau(C) = h (g - q_\tau) / ((g+h)(h + q_\tau)).$$

For the lower bound, use $h + q_\tau \le h + g$:

$$g/(g+h) - f_\tau(C) \ge h (g - q_\tau) / (g+h)^2 \ge h (C_{full} - C) / ((g+h)^2 d_\tau^+).$$

For the upper bound, use $h + q_\tau \ge h$:

$$g/(g+h) - f_\tau(C) \le (g - q_\tau) / (g+h) \le (C_{full} - C) / ((g+h) d_\tau^-). \blacksquare$$

The theorem is a coarse but explicit two-sided capacity window: whenever the capacity deficit κ and the tau spectral edges d_τ^-, d_τ^+ are independently fixed, the tau fraction can neither be tuned arbitrarily close to census nor pushed arbitrarily far below the level the capacity shortfall dictates. The window fixes the proportionality class of the vanishing deficit, not one exact coefficient; near full maintenance the released support is drawn from the maximal-demand tau eigenspace, and the exact nonlinear conversion between capacity deficit and census deficit is deliberately reserved for a subsequent gate.

7.6 Physical three-branch form

For $g = 1$, $h = 2$, Theorem 5 becomes

$$0 < C < C_{full} \Longrightarrow f_\tau(C) < \tfrac{1}{3}.$$

7.7 Rank-one explicit formulae

Let

$$D_M = d_e \Pi_e + d_\mu \Pi_\mu + d_\tau \Pi_\tau, \quad d_e \le d_\mu < d_\tau.$$

Then

$$C_0 = d_e + d_\mu, \quad C_{full} = d_e + d_\mu + d_\tau.$$

For $C_0 < C < C_{full}$, the tau occupancy is

$$\alpha_\tau(C) = (C - C_0)/d_\tau, \quad 0 < \alpha_\tau < 1,$$

so

$$R_C = \Pi_e + \Pi_\mu + \alpha_\tau \Pi_\tau$$

and

$$f_\tau(C) = \alpha_\tau / (2 + \alpha_\tau) < \tfrac{1}{3}.$$

The deficit is exactly

$$\tfrac{1}{3} - f_\tau(C) = 2(1 - \alpha_\tau) / (3(2 + \alpha_\tau)),$$

with the two-sided capacity control

$$2 (C_{full} - C) / (9 d_\tau) \le \tfrac{1}{3} - f_\tau(C) \le (C_{full} - C) / (3 d_\tau),$$

which is exactly Theorem 6 specialised to $g = 1$, $h = 2$, $d_\tau^- = d_\tau^+ = d_\tau$; it also follows directly from the closed form using $2 \le 2 + \alpha_\tau \le 3$ and $1 - \alpha_\tau = (C_{full} - C)/d_\tau$.

7.8 Suppression factor

Define the census-normalised suppression factor

$$S_\tau(C) = f_\tau(C) / f_\tau^{census}.$$

For the physical three-branch case,

$$S_\tau(C) = 3 \alpha_\tau / (2 + \alpha_\tau),$$

so

$$0 \le S_\tau(C) < 1$$

in the saturated regime, with $S_\tau(C) \to 1$ only as $\alpha_\tau \to 1$.

---

8. Centred Tau Contrast

8.1 Normalised maintenance operator

Define

$\rho_C = R_C / \mathrm{Tr}\, R_C.$

Then

$\rho_C \geq 0, \mathrm{Tr}\, \rho_C = 1,$

and the tau fraction is

$f_\tau(C) = \mathrm{Tr}(\Pi_\tau \rho_C).$

8.2 Centred tau projector

Define

$Q_\tau = \Pi_\tau - (g/(g+h)) I_\ell.$

It is traceless:

$\mathrm{Tr}\, Q_\tau = g - (g/(g+h))(g+h) = 0.$

Its eigenvalues are

$h/(g+h)$ on $\mathcal{H}_\tau$, $-g/(g+h)$ on $\mathcal{H}_0$.

8.3 Theorem 7 — Exact Tau-Contrast Theorem

Statement.

$f_\tau(C) - g/(g+h) = \mathrm{Tr}(Q_\tau \rho_C),$

and in the saturated regime

$\mathrm{Tr}(Q_\tau \rho_C) < 0.$

Proof.

$\mathrm{Tr}(Q_\tau \rho_C) = \mathrm{Tr}(\Pi_\tau \rho_C) - (g/(g+h)) \mathrm{Tr}\, \rho_C = f_\tau(C) - g/(g+h),$

since $\mathrm{Tr}\, \rho_C = 1$. Strict negativity follows from Theorem 5. ■

8.4 Exact partial-regime contrast

When the complement is full,

$f_\tau(C) = q_\tau / (h + q_\tau)$,

so

$\kappa_\tau(C) \coloneqq f_\tau(C) - g/(g+h) = h\,(q_\tau - g) / (\,(g+h)(h + q_\tau)\,)$.

Since $q_\tau < g$,

$\kappa_\tau(C) < 0$.

8.5 Physical rank-one contrast

For $g = 1, h = 2$,

$\kappa_\tau(C) = 2(\alpha_\tau - 1) / (\,3(2 + \alpha_\tau)\,)$,

with saturated range

$-\tfrac{1}{3} \le \kappa_\tau(C) < 0$.

8.6 One visible contrast direction

On the trace-one affine space, decompose

$\rho_C = (1/(g+h))\, I_\ell + \beta_\tau Q_\tau + \rho_C^\perp$,

with

$\mathrm{Tr}\, \rho_C^\perp = 0$, $\mathrm{Tr}(Q_\tau \rho_C^\perp) = 0$.

Only the Q_τ component affects f_τ.

Thus the saturated-maintenance optimisation produces precisely the contrast direction required by the participation formalism — and Theorem 7 fixes its sign.

---

9. Hard-Projection and Soft-Saturation Forms

9.1 Hard capacity threshold

Define the hard spectral projector

$P_\Lambda = \mathbf{1}_{[0, \Lambda]}(D_M)$.

9.2 Theorem 8 — Hard Tau-Exclusion Projection

Statement. If

$d_0^+ \leq \Lambda < d_\tau^-$,

then

$P_\Lambda \Pi_\tau = 0$,

and if $P_\Lambda \neq 0$ its tau-maintenance fraction is

$f_\tau[P_\Lambda] = 0$.

Proof. The spectral support of Π_τ lies strictly above Λ. The spectral projector onto demands at or below Λ therefore has no overlap with the tau sector. ■

9.3 Interpretation of exact zero

The result $f_\tau = 0$ does not mean:

- tau states do not exist;
- the tau field is removed from the theory;
- tau particles cannot be produced;
- the tau mass is infinite;
- tau amplitudes vanish in every process.

It means that the declared saturated-maintenance response allocates no continuing support to the tau branch at that capacity.

9.4 Soft maintenance profile

A sharp threshold may be regularised.

Let

$s_\Lambda : [0, \infty) \to (0, 1]$

be non-increasing, and define

$R_\Lambda^{\mathrm{soft}} = s_\Lambda(D_M)$

by spectral functional calculus. Examples include rational, exponential, logistic, and entropy-regularised profiles (Appendix D derives the last from a regularised optimisation rather than by selection).

9.5 Sector mean occupancies

Define

$$\bar{s}_\tau = (1/g)\,\mathrm{Tr}(\,\Pi_\tau\, s_\Lambda(D_M)\,),\quad \bar{s}_0 = (1/h)\,\mathrm{Tr}(\,\Pi_0\, s_\Lambda(D_M)\,).$$

The soft tau fraction is

$$f_\tau^{\mathrm{soft}} = g\,\bar{s}_\tau \,/\, (\,g\,\bar{s}_\tau + h\,\bar{s}_0\,).$$

9.6 Theorem 9 — Soft Tau-Maintenance-Suppression Theorem

Statement. If s_Λ is non-increasing and $d_0^+ < d_\tau^-$, then

$$\bar{s}_\tau \le \bar{s}_0 \text{ and } f_\tau^{\mathrm{soft}} \le g/(g+h).$$

If s_Λ decreases strictly somewhere across the closed interval $[d_0^+, d_\tau^-]$ separating the two spectra, then

$$\bar{s}_\tau < \bar{s}_0 \text{ and } f_\tau^{\mathrm{soft}} < g/(g+h).$$

Proof. Every eigenvalue of D_M in the tau sector strictly exceeds every eigenvalue in the complement. A non-increasing profile therefore assigns to every tau eigenvalue an occupancy no larger than that of any complement eigenvalue; strict decrease across the separating interval makes every tau occupancy strictly smaller than every complement occupancy. Sector averaging preserves the inequality. Finally,

$$f_\tau^{\mathrm{soft}} < g/(g+h) \Leftrightarrow (g+h)\,g\,\bar{s}_\tau < g\,(g\,\bar{s}_\tau + h\,\bar{s}_0) \Leftrightarrow \bar{s}_\tau < \bar{s}_0,$$

and likewise with non-strict inequalities. ■

9.7 Sign robustness across saturation profiles

The exact magnitude depends on the chosen profile.

The sign does not, provided:

1. the profile is independently fixed;
2. it decreases with maintenance demand;

3. the tau spectrum lies strictly above the complement spectrum.

---

10. General Benefit-to-Demand Extension

10.1 Why a benefit operator may be needed

The neutral objective Tr R assigns equal closure value to each maintained direction.

A more general gate may derive a positive maintenance-benefit operator

B_M ≥ 0.

The objective then becomes

Tr(B_M R).

10.2 Commuting benefit and demand

Assume

[B_M, D_M] = 0.

In their common eigenbasis,

$D_M |i\rangle = d_i |i\rangle$, $B_M |i\rangle = b_i |i\rangle$,

and the optimisation is

$\max \Sigma_i b_i r_i$ subject to $\Sigma_i d_i r_i \le C$, $0 \le r_i \le 1$.

Assume additionally

B_M > 0

on the carrier. This mirrors the strict-positivity convention for D_M in Section 3.5, and for the same reason: on ker B_M an occupancy increment raises the objective by nothing, so an optimum need not saturate the capacity constraint and the threshold characterisation degrades there. If B_M is only positive semidefinite, the efficiency theorem applies on supp B_M. Extending the allocation onto ker B_M requires a secondary rule: free allocation there is harmless for the benefit objective but changes the unweighted tau fraction, so without such a rule the full-carrier unweighted tau fraction need not be unique. Under the standing assumption B_M > 0 the issue does not arise, and every statement below applies to the full carrier.

Lemma 1 applies verbatim; Lemma 3 applies with the pinching taken in the common eigenbasis; and Lemma 2 applies with the mixing perturbation $R_\varepsilon = (1-\varepsilon)R + \varepsilon I$, whose objective increase $\mathrm{Tr}(B_M(I-R))$ is strictly positive for $R \neq I$ precisely because $B_M > 0$.

The non-commuting case $[B_M, D_M] \neq 0$ is out of scope for SMPτ-1: without a common eigenbasis the efficiency ordering of Section 10.3 is not spectral, and the corner argument of Theorem 3 no longer decouples benefit from demand. A generalised gate would require its own reduction theorem. Falsifiers F24 and F25 are to be read with this commutator condition explicit.

10.3 Maintenance efficiency

Define the benefit-to-demand efficiency

$\eta_i = b_i / d_i$.

10.4 Theorem 10 — Generalised Efficiency-Projection Theorem

Statement. For $0 < C < C_{full}$, an optimum fills modes in non-increasing order of $\eta_i = b_i/d_i$, with partial occupancy possible only at the threshold efficiency η_C. Every optimum has the form

$R = P_{>\eta_C} + X,\ 0 \leq X \leq P_{=\eta_C}$,

where $P_{>\eta_C}$ and $P_{=\eta_C}$ are the spectral projectors of the efficiency operator $B_M D_M^{-1}$, and the threshold block is constrained by the residual capacity through

$\mathrm{Tr}(D_M X) = C - \mathrm{Tr}(D_M P_{>\eta_C})$.

Unlike the neutral case, the capacity fixes $\mathrm{Tr}(D_M X)$ rather than $\mathrm{Tr}\, X$; the trace of X itself is determined only when D_M is proportional to the identity on the threshold-efficiency subspace.

Proof. Suppose $r_i < 1$ and $r_j > 0$ with $\eta_i > \eta_j$. Transfer maintenance cost $\delta C > 0$ from j to i:

$\delta r_i = \delta C/d_i,\ \delta r_j = -\delta C/d_j$,

with δC small enough to keep both occupancies in $[0, 1]$. The cost is unchanged; the objective changes by

$\delta C\ (b_i/d_i - b_j/d_j) > 0$,

contradicting optimality. The corner argument of Theorem 3 applies unchanged, since it uses only $0 \leq R \leq I$. Finally, the analogue of Lemma 2 (which uses $B_M > 0$) forces $\mathrm{Tr}(D_M R) = C$ at every optimum, and subtracting the cost of the fully occupied block gives $\mathrm{Tr}(D_M X) = C - \mathrm{Tr}(D_M P_{>\eta_C})$. ■

10.5 General tau-efficiency separation

Define, using the commuting assumption to make the ratio spectral,

$\eta_\tau^+ = \sup \sigma(B_M D_M^{-1} |\mathcal{H}_\tau)$, $\eta_0^- = \inf \sigma(B_M D_M^{-1} |\mathcal{H}_0)$.

If

$\eta_\tau^+ < \eta_0^-$,

the tau sector is still filled after the entire complement, and Theorems 4–7 hold with the demand ordering replaced by the efficiency ordering.

The tau-maintenance-suppression theorem therefore survives unequal branch benefits whenever the tau branch has strictly lower maintenance efficiency.

One uniqueness caveat is required. Under strict cross-sector efficiency separation, the suppression sign holds for every optimum. A unique numerical function $f_\tau(C)$ additionally requires one of: D_M scalar on the threshold-efficiency block, a one-dimensional threshold, or a secondary rule fixing X. Without such a condition, the general benefit case fixes the sign of the tau contrast but not a single-valued tau fraction.

10.6 Neutral case

For $B_M = I$,

$\eta_i = 1/d_i$,

so lower demand means higher efficiency, recovering Theorem 2.

10.7 Importance of this extension

Tau suppression does not fundamentally require equal benefit.

It requires the stronger and more physical condition:

The tau branch yields less maintained closure per unit maintenance capacity than the electron–muon complement.

---

11. Capacity Evolution, Continuity, and Recovery Toward Census

11.1 The value function

Define the maintained-support value function

$V(C) = \mathrm{Tr}\, R_C = \max \{ \mathrm{Tr}\, R : R \in \mathcal{K}_C \}$.

11.2 Theorem 11 — Capacity Monotonicity, Concavity, and Continuity

Statement. Under Premises P1–P8:

1. $V(C)$ is non-decreasing, concave, continuous, and piecewise linear on $[0, C_{full}]$, with slope $1/\lambda$ on each interval where the threshold sits at demand value λ; the slopes decrease as C increases.
2. $q_\tau(C)$ is continuous, non-decreasing, and piecewise linear: zero on $(0, C_0]$, equal to the tau-sector threshold filling on (C_0, C_{full}), and g at C_{full}.
3. $f_\tau(C)$ is defined for $C > 0$ (strict positivity of D_M gives $R_C = 0$ at $C = 0$, so $f_\tau(0)$ is undefined), is continuous and non-decreasing on $(0, C_{full}]$, satisfies $\lim_{C\downarrow 0} f_\tau(C) = 0$, and satisfies $f_\tau(C) < g/(g+h)$ for all $0 < C < C_{full}$.

Proof. Monotonicity: $\mathcal{K}_C \subseteq \mathcal{K}_{C'}$ for $C \leq C'$. Concavity: if $R_1 \in \mathcal{K}_{C_1}$ and $R_2 \in \mathcal{K}_{C_2}$, then for $t \in [0, 1]$ the convex combination $tR_1 + (1-t)R_2$ lies in $\mathcal{K}_{tC_1+(1-t)C_2}$, since the operator interval $[0, I]$ is convex and the cost is linear; hence $V(tC_1+(1-t)C_2) \geq tV(C_1) + (1-t)V(C_2)$. Piecewise linearity and the slope law follow from the explicit threshold solution: while the threshold occupies demand value λ, each additional unit of capacity purchases $1/\lambda$ additional occupancy, and the demand values encountered are non-decreasing in C, so the slopes $1/\lambda$ are non-increasing — consistent with concavity. A finite concave function on an interval is continuous on its interior, and the explicit formula extends continuity to the endpoints.

For q_τ: Theorem 4 gives $q_\tau = 0$ on $(0, C_0]$. On (C_0, C_{full}) the complement is full at cost C_0 and the remaining capacity $C - C_0$ fills the tau spectrum from below, so $q_\tau(C)$ is the same threshold-filling function restricted to $\mathscr{H}_\tau$: continuous, piecewise linear, non-decreasing, with $q_\tau(C_0^+) = 0$ and $q_\tau(C_{full}^-) = g$. In particular no jump occurs at C_0.

For f_τ: on (C_0, C_{full}), $f_\tau = q_\tau/(h + q_\tau)$ is a continuous increasing function of q_τ, hence continuous and non-decreasing in C; it vanishes continuously as $C \downarrow C_0$, matching the low-capacity regime. Strict inequality with census is Theorem 5. ■

11.3 Rank-one recovery curve

For one tau direction of demand d_τ,

$q_\tau(C) = (C - C_0)/d_\tau$

throughout the partial-tau interval, so

$$f_\tau(C) = (C - C_0) / (2 d_\tau + C - C_0),$$

with derivative

$$df_\tau/dC = 2 d_\tau / (2 d_\tau + C - C_0)^2 > 0.$$

11.4 Census as the full-capacity limit

The raw census is not wrong. It is the limiting participation fraction when the maintenance operator becomes uniform:

$$R_C \to I,$$

so

$$\lim_{C \uparrow C_{full}} f_\tau(C) = \tfrac{1}{3}.$$

The theorem identifies the condition under which the census becomes physically realised: complete rather than saturated maintenance.

11.5 No overshoot

Under strict tau-last ordering and the neutral objective,

$$f_\tau(C) \leq f_\tau^{census}$$

for every capacity $C > 0$.

The tau fraction cannot overshoot one third merely by increasing the same maintenance capacity. An overshoot would require a change of objective, benefit structure, demand ordering, carrier, or participation type.

---

12. Basis Covariance and Branch Symmetry

12.1 Theorem 12 — Basis-Covariance Theorem

Statement. Under a unitary change of carrier basis,

$$D_M \mapsto U D_M U^\dagger, \quad R_C \mapsto U R_C U^\dagger, \quad \Pi_\tau \mapsto U \Pi_\tau U^\dagger,$$

the capacity, objective, optimum set, and tau fraction are unchanged.

Proof. Unitary conjugation preserves positivity, operator order, trace, and the trace pairing:

$$\mathrm{Tr}(U D_M U^\dagger \cdot U R_C U^\dagger) = \mathrm{Tr}(D_M R_C),\quad \mathrm{Tr}(U \Pi_\tau U^\dagger \cdot U R_C U^\dagger) = \mathrm{Tr}(\Pi_\tau R_C).$$

The feasible set and objective are therefore carried onto themselves, so the optimum set is conjugated rather than altered. ■

12.2 Physical branch versus basis label

Transforming the operator while leaving Π_τ fixed is not a passive basis change. It changes the physical relationship between the demand and the declared tau branch, and is covered by falsifier F26.

12.3 Degenerate tau multiplicity

If $\mathcal{H}_\tau$ has rank $g > 1$ and the tau demand is internally degenerate, unitary averaging within the demand-degenerate tau block supplies the canonical representative

$$R_\tau = \alpha_\tau I_g.$$

This averaging is a mathematical convention unless an independent tau-block symmetry is derived: demand degeneracy alone does not establish a physical symmetry. By Section 5.8 the convention changes nothing aggregate — the tau fraction is constant across the optimum set — so the choice of representative carries no physical weight at this gate.

12.4 Flavour-mixing caveat

If the derived demand operator fails to commute with the inherited branch projectors, the meaning of "tau-sector demand" requires further work.

Possible resolutions include:

- deriving a physical spectral branch projector;
- proving approximate block diagonalisation;
- deriving a charged-lepton mass-frame bridge;
- defining a basis-invariant tau observable.

SMPτ-1 does not license manual projection onto a preferred flavour coordinate when the dynamics mixes that coordinate.

---

13. Perturbative Stability

13.1 Demand-gap stability

Let

$\tilde{D}_M = D_M + E, E = E^\dagger,$

and assume the correction preserves the branch decomposition:

$[E, \Pi_\tau] = 0.$

13.2 Theorem 13 — Tau-Gap Stability Theorem

Statement. If

$\|E\|_\infty < \Delta_\tau / 2,$

then the corrected demand operator retains strict tau-demand separation, with corrected gap at least $\Delta_\tau - 2\|E\|_\infty > 0$.

Proof. Weyl bounds give

$\sup \sigma(\tilde{D}_M |\mathscr{H}_0) \le d_0^+ + \|E\|_\infty, \inf \sigma(\tilde{D}_M |\mathscr{H}_\tau) \ge d_\tau^- - \|E\|_\infty.$

Therefore

$\inf \sigma(\tilde{D}_\tau) - \sup \sigma(\tilde{D}_0) \ge \Delta_\tau - 2\|E\|_\infty > 0.$ ■

13.3 Corollary — Suppression-Sign Stability

Under the conditions of Theorem 13, and for any capacity remaining strictly inside $(0, \tilde{C}_{full})$ with $\tilde{C}_{full} = \mathrm{Tr}\, \tilde{D}_M$, the tau-last ordering, the strict census inequality, and the negative tau contrast all survive the correction, since Theorems 4, 5, and 7 use only positivity, branch covariance, and the strict gap.

13.4 Non-branch-preserving corrections

If E mixes the tau and complement sectors, norm smallness alone does not preserve the original branch projector.

A complete treatment then requires:

- spectral-subspace perturbation bounds;
- a stable physical identification of the tau branch;
- control of the angle between old and corrected projectors.

This is an additional bridge, not silently included in SMPτ-1.

13.5 Capacity uncertainty

If C is uncertain but remains inside

$0 < C < C_{\text{full}}$,

the suppression sign remains negative; and if C is bounded away from C_{full} by an independently known margin $\delta > 0$, Theorem 6 supplies the uniform magnitude floor

$g/(g+h) - f_\tau(C) \geq h\, \delta / ((g+h)^2\, d_\tau^+)$.

This separates a robust sign prediction, a capacity-margin magnitude floor, and a fully numerical magnitude prediction into three distinct evidential grades.

---

14. Numerical Non-Insertion Theorems

14.1 Structural sign versus numerical magnitude

The sign

$f_\tau - \frac{1}{3} < 0$

follows from:

1. finite capacity;
2. a saturated optimisation;
3. tau-last demand or efficiency ordering.

The numerical value of f_τ additionally requires:

- the complete demand spectrum;
- the capacity;
- the maintenance objective;
- threshold degeneracies;
- the state and scale;
- any softening profile.

14.2 Theorem 14 — Constructive Magnitude Underdetermination

Statement. In the rank-one three-branch model, every target value

$f\star \in [0, \frac{1}{3})$

can be produced by some positive demands satisfying $d\_e \leq d\_\mu < d\_\tau$ and some saturated capacity $C < C\_full$.

Proof. For $f\star = 0$, choose any capacity $0 < C \leq d\_e + d\_\mu$.

For $0 < f\star < \frac{1}{3}$, solve

$f\star = \alpha/(2+\alpha) \Longrightarrow \alpha = 2f\star/(1-f\star) \in (0, 1)$.

Choose any positive ordered demands and define

$C = d\_e + d\_\mu + \alpha\ d\_\tau$.

Then $C_0 < C < C\_full$, and the saturated solution has tau occupancy α, giving $f\_\tau = f\star$. ■

14.3 Corollary — No numerical magnitude from saturation alone

The existence of a saturated projection and tau-last ordering fixes the sign but not the magnitude.

A numerical value is VERSF-derived only if the demand spectrum and capacity are independently fixed. (Theorem 6 sharpens this: a capacity margin fixes a floor, not a value.)

14.4 No-target-fitting requirement

A valid prediction must satisfy

$\{ m\_\tau, \tau\_\tau, \Gamma\_\tau, y\_\tau, f\_\tau\verb|^|observed \} \cap Inputs(D\_M, C, B\_M) = \emptyset$

unless the use of those quantities is explicitly declared as empirical calibration rather than derivation.

14.5 Inverse reconstruction is not prediction

Given a selected fraction, one may infer the required partial occupancy:

$\alpha_\tau = 2f_\tau/(1-f_\tau)$,

and then infer

$C - C_0 = \alpha_\tau d_\tau$.

This is a legitimate inverse audit. It is not a first-principles derivation if the observed fraction supplied the starting value.

14.6 Theorem 15 — Independent-Contrast Requirement

Statement. A numerical tau-maintenance fraction is VERSF-predictive only if the normalised tau contrast

$\kappa_\tau = \mathrm{Tr}(Q_\tau \rho_C)$

is fixed without using the intended tau result as an input.

Deriving the complete operator R_C is sufficient but not necessary; independently deriving the relevant demand ordering, spectrum, and capacity is enough.

---

15. Mass, Lifetime, and Observable Firewalls

15.1 Maintenance suppression is not mass suppression

The theorem proves a reduction in maintenance participation:

$f_\tau^{\mathrm{maint}} < \frac{1}{3}$.

It does not prove $m_\tau < m_\mu$ or $m_\tau > m_\mu$.

A mass result requires a separate relation between maintained support and the physical pole mass or Yukawa eigenvalue.

15.2 Conditional inverse-support mass corollary

Suppose a later VERSF gate independently proves

$m_a = \mathfrak{M}(\bar{r}_a)$, $d\mathfrak{M}/d\bar{r} < 0$,

where $\bar{r}_a$ is the branch maintenance occupancy. Then

$\bar{r}_\tau < \bar{r}_{e\mu} \Longrightarrow m_\tau > m_{e\mu}$.

This is a conditional corollary only. SMPτ-1 does not prove the function $\mathfrak{M}$.

15.3 Maintenance suppression is not a lifetime prediction

A lower sustainable-maintenance share may suggest a possible connection to reduced persistence. But the physical tau lifetime requires:

- the electroweak interaction Hamiltonian;
- available decay channels;
- matrix elements;
- phase space;
- radiative corrections;
- the relation between maintenance deficit and decay dynamics.

No lifetime follows from the projection theorem alone.

15.4 Maintenance suppression is not a Yukawa prediction

The physical tau Yukawa coupling requires:

- the Higgs-radial bridge;
- left–right chiral pairing;
- field normalisation;
- renormalisation scale and scheme;
- diagonalisation of the charged-lepton Yukawa matrix.

The maintenance fraction is not itself a Yukawa coupling.

15.5 No disappearance theorem

Even an exact hard-maintenance result

$\Pi_\tau R_C = 0$

does not remove $\mathscr{H}_\tau$ from the carrier.

The tau branch remains part of the theory but lies outside the continuously maintained support at that capacity.

15.6 Electroweak representation firewall

Left-handed and right-handed charged leptons transform differently under the electroweak gauge group.

A carrier combining them must specify:

- whether chirality is counted explicitly;
- whether the maintenance operator acts before or after electroweak symmetry breaking;
- how the tau projector spans the corresponding left and right sectors;
- how gauge covariance is preserved.

15.7 Observable extraction

An empirical tau-maintenance fraction requires a gauge-invariant observable $\mathcal{O}$ with a derived decomposition

$$\mathcal{O} = \mathcal{O}_\tau + \mathcal{O}_0,$$

where both terms use the same state, scale, scheme, and measurement functional.

Without this bridge, $f_\tau(C)$ is an internal VERSF maintenance quantity.

15.8 What SMPτ-1 genuinely advances

The paper advances the programme from:

a tau contrast is mathematically required

to:

an independently derived tau-last efficiency ordering, combined with finite maintenance capacity, forces that contrast to be negative, with a capacity-controlled window on its magnitude.

It does not yet advance from the contrast to a measured tau observable.

---

16. Premise-Dependency Map

The named premises are:

Premise	Content
P1	Carrier consistency: one counting convention across e, μ, τ
P2	Branch-projector consistency: Π_τ physical, not a coordinate label

Premise	Content
P3	Positive carrier metric
P4	Independent demand construction: (L_M, G_M) fixed by prior VERSF structure
P5	Branch covariance: [D_M, Π_a] = 0
P6	Finite capacity: 0 < C < C_full, independently fixed
P7	Neutral branch value (or derived B_M with efficiency separation, Section 10)
P8	Strict tau-demand separation: Δ_τ > 0

Dependencies of the main results:

Result	P1	P2	P3	P4	P5	P6	P7	P8
Theorem 1 (positive demand)	●		●	●				
Lemma 1 (existence)	●		●	●				
Lemma 2 (saturation)	●		●	●		●		
Lemma 3 (reduction)	●		●	●				
Theorem 2 (projection)	●		●	●		●	●	
Theorem 3 (optimum set)	●		●	●		●	●	
Theorem 3b (dual certificate)	●		●	●		●	●	
Theorem 4 (tau-last)	●	●	●	●	●	●	●	●
Theorem 5 (suppression)	●	●	●	●	●	●	●	●
Theorem 6 (deficit bound)	●	●	●	●	●	●	●	●
Theorem 7 (contrast)	●	●	●	●	●	●	●	●
Theorem 8 (hard exclusion)	●	●	●	●	●			●
Theorem 9 (soft suppression)	●	●	●	●	●			●
Theorem 10 (efficiency)	●	●	●	●	●	●	replaces P7	replaced by η-gap
Theorem 11 (value function)	●	●	●	●	●	●	●	●
Theorem 12 (covariance)	●		●	●				
Theorem 13 (stability)	●	●	●	●	●			●

Result	P1	P2	P3	P4	P5	P6	P7	P8
Theorems 14–15 (non-insertion)	audit-level; consume no physical premise beyond the model class							

The map makes explicit that the load-bearing physical premises are P4 (independence of the demand construction), P6 (independently fixed finite capacity), P7 (the neutral objective, or an independently derived benefit replacement — Appendix B.8 shows that an unconstrained benefit operator can reverse the ordering entirely), and P8 (the strict tau gap). P2 and P5 are additionally indispensable to interpreting the selected spectral sector as the physical tau branch. The remaining premises are structural bookkeeping.

17. Falsification Conditions

SMPτ-1 fails, or a claimed application fails, under any of the following conditions.

#	Failure condition	Consequence
F1	The charged-lepton carrier is not explicitly defined	No valid denominator
F2	Electron, muon, and tau sectors use different hidden multiplicities	Census invalid
F3	The tau projector is a preferred basis label rather than a physical branch projector	Basis dependence
F4	The carrier contains unresolved negative- or null-norm directions	Positive optimisation invalid
F5	The maintenance-load map is not specified	Demand operator unexplained
F6	The response metric G_M is not positive	Demand interpretation fails
F7	D_M is fitted directly from the desired tau fraction	Hidden insertion
F8	Tau mass or lifetime is used to create the demand ordering and then "predicted"	Circularity
F9	The capacity C is selected to reproduce a target fraction	Magnitude insertion
F10	The capacity is infinite or exceeds full-maintenance cost	No saturated suppression
F11	The allocation operator is allowed to exceed I without interpretation	Occupancy inconsistency
F12	The maintenance cost is not Tr(D_M R) or an independently derived replacement	Variational mismatch
F13	The objective is changed after seeing the desired ordering	Objective insertion
F14	A higher-demand mode is filled while a lower-demand mode remains available under the neutral objective	Not an optimum

#	Failure condition	Consequence
F15	Tau demand overlaps or falls below complement demand	Strict tau-last theorem fails
F16	Equal census rank is called equal maintenance	Census/dynamics confusion
F17	Exact zero maintenance is interpreted as tau nonexistence	Category error
F18	Tau suppression is called a mass prediction without a mass bridge	Mass overclaim
F19	Tau suppression is called a lifetime prediction without decay dynamics	Lifetime overclaim
F20	Tau suppression is identified with the Yukawa coupling	Operator-type confusion
F21	The hard projector is inserted without deriving a finite-capacity principle	Projection insertion
F22	A soft profile is selected from the desired result	Profile insertion
F23	A soft profile increases with demand while still being called maintenance saturation	Physical inconsistency
F24	Benefit differences are ignored despite being dynamically present	Objective incompleteness
F25	Tau benefit-to-demand efficiency is not lower in the generalised model	General suppression fails
F26	A basis transformation is applied to D_M but not to Π_τ	Inconsistent covariance
F27	Flavour-mixing terms are ignored without a block-diagonalisation theorem	Branch ambiguity
F28	A branch-mixing perturbation is treated as covered by the branch-preserving stability theorem	Stability overclaim
F29	Threshold degeneracy is ignored when claiming uniqueness of the detailed operator	Uniqueness overclaim
F30	The numerical magnitude is claimed from ordering alone	Underdetermination failure
F31	A capacity-dependent internal fraction is promoted to a universal constant	Universalisation error
F32	The carrier changes across a scale or phase threshold without rematching	Threshold mismatch
F33	Left and right chiral multiplicities are combined inconsistently	Electroweak counting failure
F34	An empirical percentage is claimed without a measurement functional	Observable failure
F35	Passing SMPτ-1 is presented as experimental confirmation	Admissibility/truth confusion
F36	The tau demand gap is smaller than uncontrolled corrections	Sign instability

#	Failure condition	Consequence
F37	The optimisation's maintenance objective has no VERSF physical justification	Physical-origin debt
F38	The tau branch is identified from observed mass ordering rather than derived VERSF structure	Retrospective labelling
F39	Full maintenance still yields a non-census fraction under the neutral carrier	Normalisation inconsistency
F40	A fitted maintenance operator is renamed a first-principles saturation projector	Hidden construction
F41	The deficit window of Theorem 6 is quoted without independently fixed C_full − C and tau spectral edges	Window overclaim
F42	Sector aggregates are claimed X-independent while a demand eigenvalue is shared between $\mathcal{H}_0$ and ℋ_τ, or claimed X-dependent under strict spectral separation	Optimum-set misuse
F43	The static optimisation is described as maintenance dynamics without a derived relaxation equation	Dynamics overclaim

---

18. SMPτ-1 Closure Certificate

SMPτ-1 requirement	Result
Charged-lepton maintenance carrier typed	Closed conditionally
Tau projector distinguished from a basis label	Closed at audit level
Raw tau census identified	Closed
Maintenance-load map architecture supplied	Closed structurally
Positive demand operator proved	Closed
Finite-capacity allocation space defined	Closed
Positive-contraction interpretation installed	Closed
Neutral maintenance objective defined	Closed conditionally
Saturated variational problem formulated	Closed
Existence of an optimum proved	Closed
Capacity-constraint saturation proved for $0 < C < C\_full$	Closed
Demand-basis reduction proved	Closed
Spectral threshold form proved	Closed
Complete optimum-set characterisation proved	Closed
Dual threshold certificate and complementary slackness proved	Closed

SMPτ-1 requirement	Result
Lower-demand-first ordering proved	Closed
Degenerate-threshold freedom identified and shown aggregate-irrelevant under strict cross-branch spectral separation	Closed
Canonical symmetric threshold representative supplied	Closed
Strict tau-demand gap defined as a named premise	Closed
Tau-last ordering proved	Closed
Zero tau support below complement saturation proved	Closed
Partial tau support after complement saturation proved	Closed
Tau fraction below census throughout saturation proved	Closed
Equality only at full maintenance proved	Closed
Two-sided capacity-controlled deficit bound proved	Closed
Centred tau contrast defined	Closed
Negative tau-contrast theorem proved	Closed
Hard tau-exclusion projector proved	Closed
Soft monotone-saturation theorem proved	Closed
Benefit-to-demand extension proved	Closed
Value-function concavity, continuity, and piecewise linearity proved	Closed
Capacity monotonicity proved	Closed
Census recovery at full capacity proved	Closed
Basis covariance proved	Closed
Branch-preserving perturbative stability proved	Closed
Numerical underdetermination proved	Closed
No-target-fitting requirement installed	Closed
Premise-dependency map supplied	Closed
Maintenance/mass distinction installed	Closed
Maintenance/lifetime distinction installed	Closed
Maintenance/Yukawa distinction installed	Closed
Electroweak representation firewall installed	Closed
Falsification table supplied	Closed
Numerical D_M from VERSF primitives	Open unless inherited
Numerical capacity C	Open
Numerical suppression magnitude	Open
Charged-lepton mass bridge	Open
Tau-lifetime bridge	Open
Observable extraction	Open
Empirical confirmation	Open

Closure grade

SMPτ-1 is mathematically closed as a conditional neutral-objective finite-capacity projection theorem, complete optimum-set characterisation, primal–dual threshold certificate, tau-last ordering theorem, tau-maintenance-suppression theorem, coarse two-sided capacity-control result, and numerical non-insertion theorem. Its commuting positive-benefit extension is conditionally closed subject to the explicit threshold-uniqueness qualifications of Section 10. The suppression sign becomes physically predictive only when the carrier, physical tau projector, maintenance demand or efficiency ordering, objective, and finite capacity are independently derived from VERSF without using the intended tau result — with the lower window bound then supplying the predictive floor of Section 13.5. The paper is neither a dynamical evolution theorem nor a mass, lifetime, Yukawa, decay, or observable theorem.

---

19. Conclusion

The charged-lepton census contains one tau branch among three flavour branches. Its exact raw census fraction is

$\frac{1}{3}$.

That count does not determine a physical maintenance share.

SMPτ-1 introduces an independently constructed positive maintenance-demand operator,

$D_M = L_M^\dagger G_M L_M \geq 0$,

and a positive maintenance allocation,

$0 \leq R \leq I$.

For finite maintenance capacity C, the system maximises maintained support subject to

$\mathrm{Tr}(D_M R) \leq C$.

The resulting optimum is not an arbitrary fitted diagonal matrix. It is a threshold spectral projection onto the lowest-demand modes, with at most a partial occupancy at the saturation boundary:

$R_C = P_{<\lambda_C} + X_C$.

More than this: **every** optimum has this form — optimality pins the low- and high-demand occupancies at their extremes, the operator inequalities $0 \leq R \leq I$ then force every off-block component to vanish, and the entire residual freedom is confined to a threshold block whose trace the capacity fixes and on which, under strict cross-branch spectral separation, no aggregate quantity of the gate depends. A dual certificate confirms the same structure from complementary slackness.

If the tau-demand spectrum lies strictly above the electron–muon spectrum, the optimiser necessarily maintains every available non-tau mode before allocating support to the tau branch.

This produces an exact three-stage structure:

complement filling → partial tau admission → full census recovery.

Below the complement-maintenance capacity,

$f_\tau(C) = 0.$

After the complement is full but before complete maintenance,

$f_\tau(C) = q_\tau(C) / (h + q_\tau(C)), 0 < q_\tau(C) < g,$

so

$f_\tau(C) < g/(g+h).$

For one tau direction among three charged-lepton directions,

$f_\tau(C) < \frac{1}{3}$

throughout every genuinely saturated regime, with the capacity-controlled two-sided bracket

$2(C_{full} - C)/(9\ d_\tau) \leq \frac{1}{3} - f_\tau(C) \leq (C_{full} - C)/(3\ d_\tau)$

in the rank-one partial-tau regime.

The exact departure from the census is the normalised expectation of the centred tau projector:

$f_\tau(C) - \frac{1}{3} = \mathrm{Tr}[\ (\Pi_\tau - \frac{1}{3} I_3)\ \rho_C\ ] < 0.$

This closes the tau-maintenance-suppression sign theorem.

The result survives replacement of the hard projector by any independently fixed maintenance profile that decreases with demand. It also survives unequal branch benefits whenever a commuting positive benefit operator gives the tau branch a strictly lower benefit-to-demand ratio.

The sign is robust under branch-preserving perturbations smaller than half the independently derived demand gap.

The numerical magnitude is not fixed by the theorem. Every value below the census fraction can be reconstructed by a suitable choice of demand spectrum and capacity. A number becomes predictive only when those quantities are independently derived; a floor becomes predictive as soon as the capacity margin and upper tau demand are.

Nor does maintenance suppression by itself determine the tau mass, lifetime, decay width, or Yukawa coupling. Those require distinct bridges.

The genuine programme advance is therefore:

positive participation requirement → finite-capacity maintenance optimisation → spectral projection → tau-last ordering → negative tau contrast within a capacity-controlled window.

In one sentence:

Within a consistently typed VERSF charged-lepton maintenance carrier, if an independently derived positive maintenance-demand operator places every tau mode at lower maintenance efficiency than every electron–muon mode, then maximisation of sustainable support under finite capacity produces a lowest-demand spectral projector as the unique optimum shape, excludes or partially attenuates the tau branch until the lower-demand complement is fully maintained, and forces the normalised tau-maintenance fraction strictly below its raw census value with a deficit confined to the two-sided window set by the capacity shortfall and the tau spectral edges, the entire suppression being the negative expectation of the centred tau-sector projector, and with no numerical mass, lifetime, Yukawa, or observable claim licensed until the corresponding bridges are independently derived.

Appendix A — Minimal Algebraic Skeleton

Carrier:

$\mathscr{H}\_\ell = \mathscr{H}_0 \oplus \mathscr{H}\_\tau$, dim $\mathscr{H}_0 = h$, dim $\mathscr{H}\_\tau = g$.

Projector:

$\Pi\_\tau^2 = \Pi\_\tau$, Tr $\Pi\_\tau = g$.

Raw census:

$f\_\tau\text{\^{}census} = g/(g+h)$.

Positive demand:

$D_M = L_M^\dagger G_M L_M \geq 0.$

Maintenance allocation:

$0 \leq R \leq I.$

Capacity:

$\mathrm{Tr}(D_M R) \leq C.$

Saturated optimisation:

$R_C \in \arg\max \{ \mathrm{Tr}\, R : 0 \leq R \leq I, \mathrm{Tr}(D_M R) \leq C \}.$

Complete optimum characterisation:

$R_C = P_{<\lambda_C} + X_C,\ 0 \leq X_C \leq P_{=\lambda_C},\ \mathrm{Tr}\, X_C = (C - \mathrm{Tr}(D_M P_{<\lambda_C})) / \lambda_C,$
$P_{>\lambda_C} R_C = 0,\ P_{<\lambda_C} R_C = P_{<\lambda_C}.$

Tau-demand gap:

$\Delta_\tau = \inf \sigma(D_\tau) - \sup \sigma(D_0) > 0.$

Complement capacity:

$C_0 = \mathrm{Tr}(D_M \Pi_0).$

Full capacity:

$C_{full} = \mathrm{Tr}\, D_M.$

Tau exclusion:

$0 < C \leq C_0 \Longrightarrow \Pi_\tau R_C = 0.$

Partial-tau regime:

$C_0 < C < C_{full} \Longrightarrow \Pi_0 R_C = \Pi_0,\ 0 < q_\tau < g.$

Tau fraction:

$f_\tau(C) = q_\tau / (h + q_\tau).$

Suppression:

$f\_\tau(C) < g/(g+h)$.

Deficit window (partial-tau regime, $\kappa = C\_full - C$):

$h \kappa / ((g+h)^2 d\_\tau^+) \leq g/(g+h) - f\_\tau(C) \leq \min\{ \kappa / ((g+h) d\_\tau^-), g/(g+h) \}$.

Normalised operator:

$\rho\_C = R\_C / \mathrm{Tr}\, R\_C$.

Centred contrast:

$Q\_\tau = \Pi\_\tau - (g/(g+h)) I$.

Exact contrast:

$f\_\tau(C) - g/(g+h) = \mathrm{Tr}(Q\_\tau \rho\_C) < 0$.

Soft saturation:

$R\text{\^{}soft} = s\_\Lambda(D\_M)$, s_Λ non-increasing.

General efficiency:

$\eta_i = b_i/d_i$,

with tau suppression persisting if

$\sup \eta\_\tau < \inf \eta_0$.

Appendix B — Worked Models

These models illustrate the theorem. They are not numerical VERSF predictions.

B.1 Three equal demands

Let

$D_1 = I_3$.

No tau-demand separation exists. The optimiser has no physical reason to distinguish the branches. A tau-maintenance-suppression theorem cannot be obtained.

B.2 Ordered demands below tau admission

Let

$D_2 = 1\cdot\Pi_e + 2\cdot\Pi_\mu + 5\cdot\Pi_\tau$.

Take $C = 2$. The optimum fills the electron mode and half of the muon mode:

$R_C = \Pi_e + \frac{1}{2}\Pi_\mu$,

so $f_\tau = 0$.

B.3 Exact complement saturation

With the same demand operator, take $C = 3$. Then

$R_C = \Pi_e + \Pi_\mu$,

and again $f_\tau = 0$.

B.4 Partial tau admission

Take $C = 5$. The complement cost is 3, leaving 2 units of capacity for the tau mode of demand 5. Hence

$\alpha_\tau = \frac{2}{5}$, $R_C = \Pi_e + \Pi_\mu + \frac{2}{5}\Pi_\tau$,

and the tau fraction is

$f_\tau = (2/5) / (2 + 2/5) = \frac{1}{6}$,

below the census value $\frac{1}{3}$. The deficit is $\frac{1}{3} - \frac{1}{6} = \frac{1}{6}$; the Theorem 6 floor gives $2(8-5)/(9\cdot 5) = 2/15 \leq \frac{1}{6}$, and the upper control gives $(8-5)/(3\cdot 5) = \frac{1}{5} \geq \frac{1}{6}$, bracketing the exact value.

B.5 Near-full maintenance

Take $C = 7.5$. Then

$\alpha_\tau = (7.5 - 3)/5 = 0.9$,

so

$f_\tau = 0.9/2.9 \approx 0.3103$.

The fraction approaches one third but remains below it.

B.6 Full maintenance

Take $C = 8$. Then $R_C = I_3$ and $f_\tau = \tfrac{1}{3}$.

B.7 Soft exponential saturation

Let

$R_s = e^{-\beta D_2}, \beta > 0.$

Then

$f_\tau = e^{-5\beta} / (e^{-\beta} + e^{-2\beta} + e^{-5\beta}).$

Because $5 > 2 \geq 1$,

$e^{-5\beta} < \tfrac{1}{2} (e^{-\beta} + e^{-2\beta}),$

so $f_\tau < \tfrac{1}{3}$, in accordance with Theorem 9.

B.8 Benefit-to-demand reversal

Suppose

$D = 1 \cdot \Pi_e + 2 \cdot \Pi_\mu + 5 \cdot \Pi_\tau, B = 1 \cdot \Pi_e + 1 \cdot \Pi_\mu + 10 \cdot \Pi_\tau.$

The efficiencies are

$\eta_e = 1, \eta_\mu = \tfrac{1}{2}, \eta_\tau = 2.$

The tau branch is now the most efficient and is admitted first.

This does not contradict SMPτ-1. It violates the tau-efficiency-separation premise, and the example demonstrates why the objective must be physically specified.

Appendix C — General $g \oplus h$ Theorem

Let the selected high-demand sector have rank g, and the lower-demand complement have rank h.

If the complement is fully maintained and the selected sector has total occupancy q, then

$$f_s = q / (h + q).$$

The census fraction is

$$f_s^{census} = g/(g+h).$$

Their difference is

$$f_s - f_s^{census} = h (q - g) / ((g+h)(h + q)).$$

Thus

$$q < g \Longrightarrow f_s < f_s^{census}.$$

The centred projector is

$$Q_s = \Pi_s - (g/(g+h)) I.$$

For the normalised allocation ρ,

$$f_s - f_s^{census} = \mathrm{Tr}(Q_s \rho).$$

The squared Hilbert–Schmidt norm is

$$\mathrm{Tr}(Q_s^2) = gh/(g+h).$$

The theorem therefore applies to any finite carrier in which one sector is strictly last in maintenance efficiency, and the deficit identity above is the algebraic source of Theorem 6.

Appendix D — Entropy-Regularised Soft Saturation

A smooth saturation profile can arise from a regularised optimisation rather than being selected by hand.

Let

$$S_{bin}(R) = -\mathrm{Tr}[\, R \ln R + (I - R) \ln(I - R) \,]$$

be the binary operator entropy for $0 < R < I$.

Consider maximising

$$\mathcal{J}_\varepsilon(R) = \mathrm{Tr}\, R - \eta\, \mathrm{Tr}(D_M R) + \varepsilon\, S_{bin}(R),$$

where $\eta > 0$ is the capacity multiplier and $\varepsilon > 0$ the regularisation strength.

The stationarity condition, taken in the demand eigenbasis where the problem decouples, is

$$1 - \eta\, d_i + \varepsilon \ln((1 - r_i)/r_i) = 0,$$

giving

$$r_i = [\, 1 + \exp((\eta\, d_i - 1)/\varepsilon)\,]^{-1},$$

that is,

$$R_{\{\varepsilon,\eta\}} = [\, I + \exp((\eta\, D_M - I)/\varepsilon)\,]^{-1}.$$

This is a strictly decreasing function of D_M. The objective $\mathcal{J}_\varepsilon$ is strictly concave on the open operator interval (the linear terms are affine and S_{bin} is strictly concave there), so the stationary point is the unique maximiser.

Consequently, strict tau-demand separation gives

$\bar{r}_\tau < \bar{r}_0$ and hence $f_\tau < g/(g+h)$,

by Theorem 9 applied to the profile $s(d) = [1 + \exp((\eta d - 1)/\varepsilon)]^{-1}$.

As $\varepsilon \to 0^+$ at fixed η, the smooth profile approaches the sharp threshold projector at demand $1/\eta$, with occupancy exactly ½ at the threshold itself. To recover a prescribed hard-capacity solution with a threshold occupancy other than ½, the multiplier must be selected as $\eta = \eta(\varepsilon)$ so that $\mathrm{Tr}(D_M R_{\{\varepsilon,\eta\}}) = C$ holds before the limit is taken; the limiting allocation then reproduces the saturated optimum of Theorem 3, including its partial threshold block.

This provides a mathematically controlled bridge between hard saturation and a softened physical response. The parameters η and ε must still be independently derived before a numerical prediction is claimed.

Appendix E — Open Bridge-Debt Ledger

Debt	Required content	Status
D1 — Carrier-origin debt	Derive the charged-lepton maintenance carrier from VERSF primitives	Prior dependency

Debt	Required content	Status
D2 — Branch-projector debt	Derive electron, muon, and tau projectors without observational relabelling	Open/prior
D3 — Maintenance-map debt	Derive L_M	Open unless inherited
D4 — Response-metric debt	Derive G_M	Open unless inherited
D5 — Demand-ordering debt	Prove strict tau-last demand or efficiency ordering	Central physical dependency
D6 — Capacity-origin debt	Derive C from VERSF closure dynamics	Open
D7 — Objective debt	Prove neutral benefit or derive B_M	Open
D8 — Saturation-profile debt	Derive hard or soft saturation law	Partially closed structurally
D9 — Scale debt	Determine the scale at which D_M and C are defined	Open
D10 — Chiral bridge debt	Embed left- and right-handed tau sectors consistently	Open
D11 — Higgs bridge debt	Connect maintenance support to Yukawa formation	Open
D12 — Mass bridge debt	Derive physical charged-lepton masses	Open
D13 — Lifetime bridge debt	Derive tau decay dynamics	Open
D14 — Observable debt	Construct a measurable maintenance-sensitive quantity	Open
D15 — Empirical debt	Compare with data after all conventions are fixed	Open
D16 — Capacity-margin debt	Independently bound C_full − C and the tau spectral edges d_τ^{-}, d_τ^{+} to activate the Theorem 6 window	Open
D17 — Relaxation-dynamics debt	Derive an evolution equation whose stable stationary points are the saturated optimisers	Open

SMPτ-1 does not discharge these debts by naming them. It identifies the exact dependencies required for the next gate.

Appendix F — References

1. K. Taylor, "The Coloured Participation Fraction Beyond 3/12 in VERSF," VERSF Theoretical Physics Programme, CPF-1, July 2026. [Inherited participation discipline, Section 1.1.]

2. R. Bhatia, *Matrix Analysis*, Springer, New York, 1997. [Pinching maps and spectral calculus, Sections 5.4–5.5.]
3. E. H. Lieb and M. Loss, *Analysis*, second edition, American Mathematical Society, Providence, 2001. [Trace inequalities and positivity, Sections 3 and 5.]
4. T. Kato, *Perturbation Theory for Linear Operators*, second edition, Classics in Mathematics reprint, Springer, Berlin, 1995. [Weyl bounds, Section 13.]
5. M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2010. [Unital completely positive trace-preserving structure of the dephasing map, Section 5.4.]
6. S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004. [Lagrangian duality and complementary slackness, Section 5.10.]
7. A. W. Marshall, I. Olkin, and B. C. Arnold, *Inequalities: Theory of Majorization and Its Applications*, second edition, Springer, 2011. [Ordering and rearrangement arguments, Sections 5.6 and 9.]
8. M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Addison-Wesley, 1995. [Electroweak interpretation firewalls only, Section 15.]
9. S. Weinberg, *The Quantum Theory of Fields, Volume II: Modern Applications*, Cambridge University Press, 1996. [Electroweak interpretation firewalls only, Section 15.]

Final One-Sentence Theorem

Within the SMPτ-1 gate, let a consistently typed finite VERSF charged-lepton maintenance carrier decompose into a rank-g tau sector and rank-h complement, let an independently derived positive maintenance-load map and positive response metric define a strictly positive demand operator whose entire tau spectrum lies above the complement spectrum, and let sustainable maintenance be represented by a positive contraction selected under a finite capacity either by the neutral objective Tr R or by an independently derived strictly positive benefit operator B_M satisfying $[B_M, D_M] = 0$; then every optimum is a lowest-demand or highest-efficiency spectral projection with at most partial threshold occupancy and vanishing off-block components — a structure certified independently by convex duality and complementary slackness — every complement mode is maintained before any tau mode, the tau support vanishes below complement saturation and remains incomplete until full maintenance, the normalised tau-maintenance fraction is therefore strictly smaller than its census value $g/(g+h)$ throughout the saturated regime with a deficit confined to the two-sided window between $h(C_{full} - C)/((g+h)^2 d_\tau^+)$ and $(C_{full} - C)/((g+h) d_\tau^-)$, its complete departure from census is the negative expectation of the centred projector $Q_\tau = \Pi_\tau - gI/(g+h)$, the suppression sign survives every independently fixed monotone softening and every sufficiently small branch-preserving correction that leaves the demand gap open, and no numerical mass, lifetime, Yukawa, decay, or observable result is VERSF-derived until the demand spectrum, capacity, benefit functional, electroweak embedding, and measurement bridge are fixed without using the intended tau result as input.