Research Article

A Gate-and-Kernel Framework for Baryogenesis on Emergent Backgrounds

Aric Dunn[1]

1. Independent researcher

Three companion case studies on the Hamada mechanism, causal fermion systems, and loop quantum cosmology suggest a common structural question: when has an asymmetry formula earned the descendant language it uses? This note formulates a gate-and-kernel answer. In a first descendant regime where linear evolution is meaningful, the observable asymmetry is written as a gated transport law,

The transport identity itself is standard variation of parameters; the useful claim is the decomposition around it. Gates are built from calibrated witness maps for ordering, geometry, current basis, thermal support, EFT control, and, when needed, transport-state admissibility. The thresholds are inherited from the control inequalities of the chosen reduction rather than fitted to the desired baryon yield, which makes the framework prospective as well as retrospective. A baryon-channel support diagnostic then yields a three-category classification of published formulas. The kernel representation is placed inside a five-layer architecture linking emergence criteria, background realization, primitive source and transfer, descendant transport, and final readout. Tests include the Hamada/CFS/LQC triad, standard thermal leptogenesis, and a harder resonant-leptogenesis stress test in which witness maps select between a diagonal Boltzmann reduction and a density-matrix description. The result is a reusable audit and model-building scaffold for asymmetry generation on emergent backgrounds.

Correspondence: papers@team.qeios.com — Qeios will forward to the authors

1. Introduction

The companion case studies all ask the same question in different settings: when has an asymmetry formula earned the source language and transport structure it uses? In the Hamada mechanism, the issue

is source identification: one anomaly channel sources baryon-plus-lepton charge directly while the other first sources lepton number and only later yields baryon number through sphalerons [1][2]. In the causal-fermion-systems line, the source paper already uses the correct primitive observable—a fermionic spectral imbalance in a smooth globally hyperbolic descendant spacetime—and the framework's role is confirmatory [3][4][5]. In loop quantum cosmology (LQC), the sharpest issue is regime control: a displayed baryogenesis formula is written in an effective FRW descendant regime with thermal and current structure already active, and one benchmark uses the interaction EFT beyond the control window stated by the source paper itself [6][7][8].

What survives that comparison is not a special FRW scaling law but a reusable decomposition. Every successful reduction has four ingredients:

1. a primitive source variable in the earliest admissible descendant basis;
2. an admissibility gate that records whether ordering, geometry, matter/current, thermal, and EFT-control assumptions are active in the stated regime;
3. a transport generator for dilution, washout, diffusion, oscillation, or charge processing;
4. a final readout that projects the transported descendant state onto the observable baryon asymmetry.

This note turns those ingredients into a formal framework. Its mathematical core is modest. The transport solution is standard variation of parameters. The nontrivial move is the bookkeeping around it: identifying the primitive source, calibrating the relevant gates, and separating source, transport, and readout in a way that can be audited across mechanisms. That decomposition sits naturally inside a broader stack. Geometry-only programs such as spacetime functionalism and layered emergence address when geometric description is licensed [9][10][11][12]. Boltzmann, density-matrix, and Kadanoff–Baym reductions already evolve asymmetries once the source and background are fixed [13][14][15][16]. The present framework sits between those layers.

The paper has four aims. First, it formalizes witness-defined channel gates and a calibrated threshold rule that can be used prospectively, not only read off from completed papers. Second, it defines a baryon-channel support diagnostic and a corresponding three-category classification. Third, it places the gate-and-kernel identity inside a five-layer architecture. Fourth, it tests the resulting framework on the Hamada/CFS/LQC triad, on standard thermal leptogenesis, and on a harder resonant-leptogenesis case in which the admissible descendant state itself is witness dependent.

2. Definitions and setup

2.1. The first descendant regime

The framework does not derive transport from a microscopic quantum–gravity theory. It begins at the first descendant stage where a linear asymmetry equation is meaningful in the published or proposed calculation.

Definition 2.1 (first descendant regime). *A reduced description R is a first descendant regime for an asymmetry calculation if it supplies:*

1. *an ordering variable τ with respect to which retarded evolution is meaningful;*
2. *a finite or effectively finite asymmetry state $Y(\tau)$ in a chosen descendant basis;*
3. *a linearized transport law for Y on the interval of interest.*

This definition is intentionally operational. In the Hamada note the first descendant regime is the low-energy FRW-effective phase below the spacetime transition. In the CFS note it is the smooth globally hyperbolic spacetime used in the source paper itself. In the LQC note it is the effective FRW-plus-thermal regime in which $\dot{R}, T_D$, and the current interaction are actually computed.

2.2. Primitive sources, witness maps, and calibrated gates

Definition 2.2 (primitive source). *Let R be a first descendant regime. The primitive source $S_{prim}(\tau)$ is the minimal source vector in the earliest admissible descendant basis that feeds the transport equation. A published source S_{pub} may coincide with S_{prim} or may represent a premature projection, specialization, or channel average of it.*

Definition 2.3 (witness map and calibrated threshold). *For each descendant requirement α—for example ordering, geometry, current basis, thermal support, EFT control, or coherence control for a chosen reduction—let*

$$W_\alpha[R, T; \tau]$$

be an operational witness functional extracted from the theory/benchmark pair (R, T). Let $W_\alpha^\star > 0$ be the control threshold supplied by the reduced formalism itself, and define the normalized witness

$$\hat{W}_\alpha(\tau) \equiv \frac{W_\alpha[R, T; \tau]}{W_\alpha^\star}.$$

The corresponding elementary gate is

$$g_\alpha(\tau) = \Theta\left(\hat{W}_\alpha(\tau) - 1\right), g_\alpha(\tau) \in \{0, 1\}.$$

In practice the thresholds are fixed by the control inequalities of the descendant reduction, not by fitting the desired final asymmetry. Retrospectively one reads them from the regime of validity stated by the source paper. Prospectively one fixes them *before* solving transport by declaring the reduced formalism to be used and then adopting its control window.

A convenient normalization makes this explicit:

$$\hat{W}_{th} = \frac{\Gamma_{int}}{H} \text{ for thermal support,} \tag{1}$$

$$\hat{W}_{EFT} = \frac{\Lambda_{EFT}}{Q_{\max}} \text{ for EFT control,} \tag{2}$$

$$\hat{W}_{decoh}^{diag} = \frac{\Delta M}{\overline{\Gamma}} \text{ for a diagonal Boltzmann reduction of a quasi-degenerate system.} \tag{3}$$

Here $Q_{\max}$ is chosen from the reduced operator itself rather than guessed after the fact. If the descendant EFT contains an operator with up to k derivatives acting on a field Φ, define the associated inverse-variation scale

$$Q_\Phi^{(k)}(\tau) \equiv \left|\frac{\nabla^k \Phi}{\Phi_{amp}}\right|^{1/k},$$

where Φ_{amp} is the nonzero amplitude, mass, or frequency scale already used to normalize that field in the chosen reduction. Then one may take

$$Q_{\max}(\tau) = \max\left\{p_{hard}(\tau), T(\tau), H(\tau), \sqrt{|R(\tau)|}, Q_{\Phi_1}^{(k_1)}(\tau), \ldots, Q_{\Phi_n}^{(k_n)}(\tau)\right\},$$

with p_{hard} the hard momentum scale carried by the sourced modes. For a curvature–sourced operator with no additional matter derivatives, this reduces to $Q_{\max} = \max\{T, H, \sqrt{|R|}\}$. When derivative operators are present, the user must declare the relevant $Q_\Phi^{(k)}$ before transport is solved; different operators can therefore carry different EFT witnesses even on the same background. Equation (3) is not unique; it is a conservative control quantity for deciding whether coherent mixing can be neglected in a flavour–diagonal resonant-leptogenesis reduction.

Remark 2.4 (prospective witness design). The same witness language can be used prospectively. One first specifies the descendant reduction R and the transport formalism T, then writes down the required witness set $\{W_\alpha\}$ for each source channel, and only afterwards asks whether any channel gate is open on the support of the proposed primitive source. This removes the circularity worry: the thresholds are inherited from the reduction, not tuned to the answer.

Definition 2.5 (channel gate). *Let r label source channels and let A_r be the set of descendant requirements needed for channel r. The channel gate is*

$$g_r(\tau) = \prod_{\alpha \in A_r} g_\alpha(\tau).$$

If Π_r projects onto the corresponding source channel, the full gate operator is

$$G(\tau) = \sum_r g_r(\tau)\Pi_r. \tag{4}$$

Restricted to geometry alone, this witness language overlaps with spacetime functionalism and layered-emergence work [9][10][11][12]. The present framework keeps that lesson but extends it to source and charge language.

2.3. Transport and baryon readout

Definition 2.6 (transport generator and readout). *The transport generator $L(\tau)$ is the linear operator governing descendant evolution once the source and gate are specified. It may contain expansion, washout, diffusion, oscillation, flavor conversion, or other processing. The final baryon asymmetry is extracted by a readout vector e_B, so that*

$$\eta_B(\tau) = e_B^\top Y(\tau).$$

The point of separating e_B from S_{prim} is precisely what the Hamada and CFS cases made concrete: the primitive source need not already be baryon number.

3. Kernel representation and diagnostic criteria

3.1. Master equation and retarded solution

The universal predictive object is the gated transport law

$$\frac{dY}{d\tau} = L(\tau)Y + G(\tau)P(\tau)S_{prim}(\tau), \eta_B(\tau_f) = e_B^\top Y(\tau_f), \tag{5}$$

where P is the source–basis/transfer map into the descendant state. Let the propagator U satisfy

$$\partial_\tau U(\tau, \tau') = L(\tau)U(\tau, \tau'), U(\tau', \tau') = I, U(\tau, \tau') = 0 \text{ for } \tau < \tau', \tag{6}$$

so that the evolution is explicitly retarded. Then the full descendant solution is

$$Y(\tau_f) = U(\tau_f, \tau_i)Y(\tau_i) + \int_{\tau_i}^{\tau_f} d\tau U(\tau_f, \tau)G(\tau)P(\tau)S_{prim}(\tau), \tag{7}$$

and therefore

$$\eta_B(\tau_f) = e_B^\top U(\tau_f, \tau_i) Y(\tau_i) + \int_{\tau_i}^{\tau_f} d\tau\, e_B^\top U(\tau_f, \tau) G(\tau) P(\tau) S_{prim}(\tau). \tag{8}$$

When no pre-existing descendant asymmetry is assumed, one sets $Y(\tau_i) = 0$ and retains only the sourced term. It is convenient to define the baryon readout kernel

$$K_B(\tau_f, \tau) \equiv e_B^\top U(\tau_f, \tau) G(\tau) P(\tau). \tag{9}$$

Equation (8) is standard variation of parameters. The framework's content is not the ODE step itself, but the disciplined decomposition of primitive source, witness-defined gates, transport, and readout around that identity.

3.2. Baryon-channel support and classification

The key support quantity must be defined in the baryon readout channel, not merely at the level of some gated source. To make the necessary condition mathematically clean, define it with an absolute value in a one-dimensional readout channel, or more generally with any compatible norm.

Definition 3.1 (baryon-channel support quantity). *In a scalar baryon readout channel define*

$$\Delta_B^{supp} \equiv \int_{\tau_i}^{\tau_f} d\tau\, |K_B(\tau_f, \tau) S_{prim}(\tau)|. \tag{10}$$

For a vector readout channel one may replace the absolute value by any fixed norm.

Proposition 3.2 (necessary support-overlap condition). *If the sourced contribution to $\eta_B(\tau_f)$ is nonzero, then $\Delta_B^{supp} > 0$. Equivalently, $\Delta_B^{supp} = 0$ implies that the sourced contribution to the baryon readout vanishes almost everywhere on $[\tau_i, \tau_f]$.*

Proof. If $\Delta_B^{supp} = 0$, then the nonnegative integrand in Eq. (10) vanishes almost everywhere, so $K_B(\tau_f, \tau) S_{prim}(\tau) = 0$ almost everywhere. The sourced integral in Eq. (8) is therefore zero. $\square$

Corollary 3.3 (sufficiency under sign definiteness). *Suppose $K_B(\tau_f, \tau) S_{prim}(\tau)$ has a definite sign on $[\tau_i, \tau_f]$ in the one-dimensional baryon readout channel. Then $\Delta_B^{supp} > 0$ if and only if the sourced contribution to $\eta_B(\tau_f)$ is nonzero.*

Proof. Under sign definiteness, the sourced integral in Eq. (8) cannot cancel against itself, so positivity of Eq. (10) is equivalent to nonvanishing of the sourced integral. $\square$

This corollary should not be applied blindly in quasi-degenerate resonant systems. In resonant leptogenesis, mixing and oscillatory contributions can change the sign of the projected source across

time or parameter space once off-diagonal correlations are retained. In that regime the support quantity remains a necessary audit, but the net asymmetry must be computed from Eq. (8) or its density-matrix analogue unless sign definiteness has been checked explicitly in the chosen reduction.

The support criterion becomes useful when combined with one further question: is the published source already the primitive one?

Definition 3.4 (three-category rule). *A published asymmetry formula in a first descendant regime is assigned:*

1. *Category I (inadmissible or regime-control failure) if $\Delta_B^{supp} = 0$, or if a required witness factor for the published reduction is unearned on the support needed to source baryon number;*

2. *Category II (corrective source identification) if $\Delta_B^{supp} > 0$, but $S_{pub} \neq PS_{prim}$ at source stage;*

3. *Category III (confirmatory) if $\Delta_B^{supp} > 0$ and $S_{pub} = PS_{prim}$ at source stage.*

This partition is bookkeeping rather than deep mathematics. Its value lies in making explicit which of two structurally different questions failed: admissibility/support, or source identification.

4. Domain of validity and relation to neighboring frameworks

4.1. What the framework assumes

The framework applies when the following conditions are met:

1. There exists a descendant ordering variable τ on the interval of interest. This may be cosmic time, conformal time, a global time function on a smooth spacetime, or another model-specific ordering parameter. The framework does not require that τ be fundamental; it only requires that the published asymmetry calculation uses it consistently.

2. The asymmetry degrees of freedom can be represented by a linear state Y over the interval considered. This includes ordinary Boltzmann vectors, flavor-density vectors, and vectorized density matrices after standard linearization or gradient expansion [13][14][15][16].

3. A primitive source, a descendant charge basis, and a workable witness set $\{W_\alpha\}$ can be identified.

4. The effects of background evolution, washout, diffusion, and transfer can be absorbed into a linear generator L, possibly time dependent and matrix valued.

4.2. What the framework does not assume

The framework does not assume a specific FRW background, a particular power–law decoupling relation, a single source mechanism, or a unique microscopic origin for the gate factors. It also does not claim that clock and geometry disappear whenever a foundational theory becomes strongly quantum. In particular, the LQC case study shows that the relevant issue is not a blanket "no clock at the bounce" claim but a narrower regime–control question about which descendant sectors the published effective benchmark has actually activated [6][8].

4.3. Where the framework does not apply

The framework is not a microscopic derivation. It should not be used where no first descendant regime exists, where the state evolution is essentially nonlinear in the sourced asymmetry, or where the very notion of a retarded source term is unavailable. Likewise, the framework does not by itself fix UV completions. If an EFT gate is opened above a stated cutoff, an additional UV–completion statement is needed to justify that move.

4.4. Relation to other frameworks

Restricted to geometry alone, the gate language overlaps with spacetime functionalism and layered–emergence discussions of when geometric description becomes legitimate [9][10][11][12]. The present framework does not compete with those programs; it extends the admissibility question to ordering structure and to matter/current–basis questions that baryogenesis formulas actually need.

On the transport side, the framework does not replace Boltzmann, density–matrix, or Kadanoff–Baym methods. Those frameworks compute L and solve Eq. (5) once the source and background are fixed [13][14][15][16]. The gate-and-kernel contribution is earlier: which source variable, charge basis, and descendant transport language are licensed in the first place?

A recent 2026 paper by Mandel proposes a *unified framework* within gravitational baryogenesis itself, combining an entropy–production clock source with CP violation from a gravitational $\theta R\tilde{R}$ term and deriving a universal transfer factor for that mechanism family [17]. In the present stack, those ingredients would sit inside the primitive–source/transfer layer of a single gravitational–baryogenesis line. The question asked here is orthogonal: which primitive sources, witness sets, and descendant transport

modules are admissible *across* different mechanisms at all? The two projects are therefore complementary rather than competitive.

5. The stacked version of the framework

The gate-and-kernel representation can be placed inside a broader architecture without changing its mathematics. The point of the stack is not to fuse every neighboring formalism into one microscopic theory. It is to show where each one belongs and how the gate-and-kernel identity plugs into them.

5.1. Five layers

A convenient ordered tuple for the stacked version is

$$\mathcal{F} = (E, B, S_{prim}, P, L, e_B), \tag{11}$$

with the following interpretation:

1. **Emergence / admissibility layer** E: witnesses or arguments that license ordering, geometry, matter/current, thermal, EFT-control, and any additional reduction-specific language.

2. **Background layer** B: the descendant background realization on which the asymmetry calculation is written, for example Einstein/FRW, effective LQC FRW, or a smooth globally hyperbolic CFS descendant spacetime.

3. **Primitive source and transfer layer** (S_{prim}, P): the earliest admissible source variable and the map from that source basis into the descendant state basis.

4. **Transport layer** L: washout, dilution, diffusion, oscillation, scattering, flavor conversion, or other linear processing.

5. **Readout layer** e_B: the projection used to extract the baryon asymmetry or another final observable.

With this notation, the master equation may be written suggestively as

$$\frac{dY}{d\tau} = L[B](\tau)Y + G[E,B](\tau)P(\tau)S_{prim}(\tau), \eta_B(\tau_f) = e_B^\top Y(\tau_f). \tag{12}$$

Equation (12) is not new mathematics beyond Eq. (5); it is a bookkeeping statement that makes the modularity visible. Changing the background module changes L and possibly some gate factors. Changing the primitive source or transfer layer changes S_{prim} and P. The formal structure stays fixed.

5.2. How neighboring frameworks fit

This layered picture clarifies what it means to combine the neighboring frameworks already in play:

- Spacetime functionalism and layered-emergence work sit naturally in E. They answer when geometric description is earned, but they do not by themselves identify a primitive asymmetry source.
- Einstein/FRW, effective LQC, and CFS descendant spacetimes are alternative background modules B.
- Boltzmann, density-matrix, and Kadanoff–Baym methods inhabit L. They remain the correct tools once gates are open and sources are fixed.
- The present framework sits between those layers, connecting admissibility, primitive-source choice, transfer maps, and readout.

6. Tests of the stacked framework

Examples are where a larger framework either earns trust or loses it. The first two tests are interoperability checks. The next three are the more demanding ones: continuity with the original triad, a standard external mechanism, and a genuinely gate-sensitive contested reduction.

6.1. Test A: geometry-only reduction

Switch off source language by taking $S_{prim} = 0$ and ignore the readout map. Then the only nontrivial question left is whether the geometry part of the gate has been licensed in the descendant description. In that limit the stacked framework adds little to spacetime functionalism, and it should not. Restricted to geometry alone, the framework overlaps with the functional-role question already studied by Knox, Lam–Wüthrich, Oriti, and Huggett–Wüthrich [9][10][11][12]. The point of the test is interoperability, not novelty inflation.

6.2. Test B: fully opened reduction to standard transport

Now suppose all relevant gate factors are open and the primitive source/transfer map have already been identified. Then G may be replaced by the identity on the active descendant sector and Eq. (12) reduces to an ordinary linear transport problem. That is precisely the regime in which Boltzmann, density-matrix, and Kadanoff–Baym methods are known to apply [13][14][15][16]. Again, this is a positive interoperability test.

6.3. Test C: continuity across the three case studies

The third test is the one that motivated the framework: can the same stacked tuple $\mathcal{F}$ host three genuinely different case–study verdicts without changing the master equation? Table 1 shows that it can.

Case	Background module B	Primitive source / transfer layer	Critical witnesses	Verdict
Hamada	Post-transition FRW-effective descendant background	$S^{\chi}_{prim} = S_{\chi}P_{\chi}$ with $P_{gauge} = (1,1)^{\top}$ and $P_{grav} = (0,1)^{\top}$	Ordering, geometry, matter/current open	Category II: one channel needs a more primitive source than Y_B
CFS	Smooth globally hyperbolic descendant spacetime with global time function	$S_{prim}(t) = \mathcal{B}(t)e_F$ with $\mathcal{B}(t)$ $= \frac{d}{dt}Tr\left(\eta\Lambda\Pi(t)\right)$	Source-paper gate open; no premature baryonic projection	Category III: the source paper already uses the primitive variable
LQC	Effective FRW descendant background with thermal-current structure	$S_{prim}(\tau)$ $, with S_D $= \delta\left(\tau - \tau_D\right)S_D e_B$ $\propto -\dot{R}_{eff}/\left(M_*^2 T_D\right)$	Thermal/current, EFT control	Category I at high T_D; lower benchmark admissible

Table 1. Continuity of the stacked framework across the Hamada, CFS, and LQC case studies. The background module, source layer, and gate profile change from case to case; the master equation does not.

The important point is methodological. The stack does not merely classify the three papers after the fact.

It shows exactly which layer had to change to obtain a different verdict. In Hamada the crucial change is in the source/transfer layer. In CFS it is a confirmation that the source layer was already right. In LQC the sharpest issue sits in the gate profile and benchmark control window.

The witness formalism can also be made prospective in the hardest of these cases. For an LQC-inspired baryogenesis operator suppressed by a cutoff M_*, a conservative EFT witness is

$$\hat{W}_{EFT}^{LQC}(\tau) = \frac{M_*}{Q_{\max}(\tau)}, Q_{\max}(\tau) = \max\left\{T(\tau), H(\tau), \sqrt{|R(\tau)|}, Q_{der}(\tau)\right\}.$$

Here Q_{der} is the operator-specific inverse-variation scale defined above. For the curvature-sourced benchmark with no additional matter derivatives one simply has $Q_{der} = 0$ and hence $Q_{\max} = \max\left\{T, H, \sqrt{|R|}\right\}$. For higher-derivative operators the user must state the relevant $Q_\Phi^{(k)}$ entering Q_{der} before transport is solved. That choice does not solve the microscopic theory, but it does say prospectively what the reduced description must check before a benchmark is treated as admissible.

6.4. Test D: standard thermal leptogenesis

A first external check is standard thermal leptogenesis. In one common reduction, one writes descendant equations of the schematic form

$$\frac{dY_{N_1}}{dz} = -D(z)\left(Y_{N_1} - Y_{N_1}^{eq}\right), \frac{dY_{B-L}}{dz} = \epsilon_1 D(z)\left(Y_{N_1} - Y_{N_1}^{eq}\right) - W(z)Y_{B-L}, \tag{13}$$

with $z = M_1/T$ and a final baryon readout

$$\eta_B = c_{sph}Y_{B-L}(z_f), c_{sph} = \frac{28}{79} \tag{14}$$

for the minimal Standard Model in the symmetric phase [2][18]. In the stacked notation, this corresponds to a thermal FRW descendant background, primitive source

$$S_{prim}(z) = \epsilon_1 D(z)\left(Y_{N_1} - Y_{N_1}^{eq}\right)e_{B-L},$$

readout $e_B = c_{sph}e_{B-L}$, and a transport generator containing the usual washout term $W(z)$. The relevant ordering, geometry, thermal, and current-basis witnesses are all open in the published descendant regime. The framework therefore classifies standard thermal leptogenesis as Category III: the primitive source is already a lepton/$B - L$ asymmetry, and baryon number appears only at readout through sphaleron transfer.

6.5. Test E: resonant leptogenesis as a hard witness test

A harder external test is resonant leptogenesis, precisely because the admissible descendant state depends on the regime. Pilaftsis and Underwood showed that resonant enhancement occurs when the heavy-neutrino mass splittings are comparable to their decay widths, with the boundary layer $\Delta M \sim \bar{\Gamma}/2$ especially delicate [19]. Garny, Kartavtsev, and Hohenegger then showed that in the maximally resonant regime a Kadanoff–Baym treatment gives a smaller asymmetry than the naive Boltzmann result because coherent transitions between the Majorana species contribute [20]. Dev, Millington, Pilaftsis, and Teresi further developed a flavour-covariant transport description that simultaneously captures resonant mixing, coherent oscillations, and charged-lepton decoherence, and they showed that the total asymmetry can differ by up to an order of magnitude from flavour–diagonal or only partially off-diagonal rate equations [21].

This is exactly the kind of case in which witness maps must select the descendant reduction *before* transport is solved. A flavour-diagonal Boltzmann truncation R_{diag} is admissible only if a decoherence witness such as Eq. (3) is open:

$$\hat{W}^{diag}_{decoh} = \frac{\Delta M}{\bar{\Gamma}} \gtrsim 1.$$

The symbol $\gtrsim 1$ should be read as a conservative diagonal–admissibility condition rather than a sharp physical phase boundary. Near the resonant layer emphasized by [19], the safe choice is to keep the enlarged descendant state unless one has an additional reduction-specific argument for diagonalization. When the conservative condition fails, a diagonal source basis is not a licensed first-descendant description of the quasi-degenerate system. One must enlarge the descendant state to include off-diagonal correlations, for example schematically

$$\frac{d\rho_N}{dz} = -i[H_N, \rho_N] - \frac{1}{2}\{\Gamma_N, \rho_N - \rho_N^{eq}\}, \frac{dY_{B-L}}{dz} = Tr\left(\epsilon\Gamma_N(\rho_N - \rho_N^{eq})\right) - W(z)Y_{B-L}, \quad (15)$$

as in density-matrix or flavour-covariant reductions [16][21].

The framework's verdict is therefore reduction dependent in a controlled way:

1. If a flavour-diagonal Boltzmann formula is used in a regime where the diagonal-decoherence witness is closed, the formula is Category I *as written*: the proposed descendant truncation is not admissible.

2. If one uses the enlarged density-matrix/flavour-covariant descendant state, the same mechanism can be Category III: the primitive source and transport language are then matched to the regime.

Because oscillatory and mixing contributions need not be sign definite in this regime, the sign–definite corollary above should not be assumed automatically here; the support quantity remains a necessary audit, but the net asymmetry must be extracted from the full descendant evolution unless sign definiteness has been checked in the chosen reduction. This is the hardest test in the paper because the witness machinery does not merely change a coefficient; it chooses the admissible state space. It also shows how simplified reductions can re-emerge after the witness check. In strong washout, for example, late-time effective asymmetries can approximate the full resonant dynamics over wide parameter regions once their own validity criteria are satisfied [22].

7. Illustrative forward use: a two-threshold spectral-transfer mechanism

The framework can also be used forward. This section records one example mechanism family suggested by the comparison of the case studies. It is not a completed UV model and it does not claim uniqueness. The point is narrower: once primitive-source identification, gates, and transfer maps are separated, a class of delayed-conversion architectures becomes natural to write down.

Related two-stage ideas already appear in darkogenesis and axiogenesis, where a hidden–sector asymmetry or a PQ-charge asymmetry is generated first and only later converted into visible baryon number [23][24]. The illustrative difference here is the primitive source. Instead of beginning from a conventional dark charge or PQ charge, the source stage generates a non-baryonic spectral or occupation asymmetry X of the same broad structural type highlighted by the CFS case study, and a later threshold opens the visible transfer channel.

7.1. A toy microscopic seed

One can also write a concrete toy microscopic origin for the stored variable X. Consider a hidden fermion χ coupled to a rolling pseudoscalar a and to a late-opening sterile-neutrino portal,

$$\mathcal{L}_{toy} = \overline{\chi}\left(i\partial - m_\chi - \frac{\partial_\mu a}{f_a}\gamma^\mu\gamma_5\right)\chi + \left(y_{\chi i}\Phi\overline{\chi}N_i + y_{Li}\tilde{H}^\dagger\overline{L}N_i + \frac{1}{2}M_{ij}N_iN_j + \text{h.c.}\right), \qquad (16)$$

with at least two sterile states N_i so that the late transfer stage can carry a genuine CP-odd phase. Time-dependent pseudoscalar backgrounds of this type are known to bias fermion-helicity production and can leave unequal populations in the two quasiparticle branches [25][26]. The primitive source variable is then not a dark conserved charge but an occupation/spectral imbalance,

$$X \equiv n_{\chi,+} - n_{\chi,-}, \tag{17}$$

or the corresponding momentum-integrated spectral-weight difference in the first descendant basis.

The late gate is provided by the portal field Φ. At early times one may have $\langle \Phi \rangle = 0$ or $m_\Phi^2(T) > 0$, so the $\chi \leftrightarrow N_i$ transfer channel is effectively closed even if the hidden-sector source is already active. After a later threshold, for example when $m_\Phi^2(T)$ changes sign or a background expectation value turns on, the portal opens and the stored X imbalance can feed the sterile-neutrino sector. A minimal baryogenic completion then uses the phases in $y_{\chi i} y_{Lj}^*$ and in the Majorana mass matrix M_{ij} to generate a nonzero CP-odd conversion efficiency ϵ_B in the late transfer/decay stage, in the same broad interference sense familiar from ordinary leptogenesis [18][19]. If those phases vanish, the mechanism is only asymmetry redistribution: the rolling pseudoscalar can still bias the stored helicity variable X, but no net visible $B - L$ is produced. With the CP-odd portal active, standard $N_i \leftrightarrow LH$ processing can convert the stored imbalance into $B - L$ and finally into baryon number. This is only a toy seed, not a completed UV completion, but it shows concretely how the forward-use mechanism can be structurally distinct from darkogenesis: the earliest source variable is a quasiparticle occupation imbalance rather than a hidden Noether charge.

7.2. Minimal descendant model

Take the descendant state to be

$$Y = (Y_X, Y_B)^\top, \tag{18}$$

with an early source in the X direction and a later transfer gate into baryon number. A minimal reduction of Eq. (5) is

$$\frac{d}{d\tau} \begin{pmatrix} Y_X \\ Y_B \end{pmatrix} = \begin{pmatrix} -\Gamma_X(\tau) & 0 \\ \epsilon_B \Gamma_{tr}(\tau) I_{tr}(\tau) & -W_B(\tau) \end{pmatrix} \begin{pmatrix} Y_X \\ Y_B \end{pmatrix} + \begin{pmatrix} I_X(\tau) S_X(\tau) \\ 0 \end{pmatrix}. \tag{19}$$

Here S_X is the primitive source, I_X records whether the source stage is admissible, I_{tr} records whether the later transfer channel is open, Γ_X controls depletion of the stored X asymmetry, Γ_{tr} is the conversion rate into the visible sector, W_B is the visible-sector washout rate, and ϵ_B is a dimensionless CP-odd

baryonic yield per transferred unit of X, inherited from the complex portal/sterile-neutrino sector and vanishing if the late conversion is CP symmetric.

The formal solution for the baryon component is

$$Y_B(\tau_f) = \int_{\tau_i}^{\tau_f} d\tau e^{-\int_\tau^{\tau_f} du W_B(u)} \epsilon_B \Gamma_{tr}(\tau) I_{tr}(\tau) Y_X(\tau), \tag{20}$$

where

$$Y_X(\tau) = e^{-\int_{\tau_i}^\tau du \Gamma_X(u)} Y_X(\tau_i) + \int_{\tau_i}^\tau du e^{-\int_u^\tau dv \Gamma_X(v)} I_X(u) S_X(u). \tag{21}$$

If the primitive source is localized near τ_X and the transfer gate opens later near $\tau_{tr} > \tau_X$, Eqs. (20) and (21) give the schematic prediction

$$Y_B(\tau_f) \approx \epsilon_B T_{BX} \exp\left[-\int_{\tau_X}^{\tau_{tr}} du \Gamma_X(u)\right] Y_X(\tau_X), T_{BX} \equiv \int_{\tau_{tr}}^{\tau_f} d\tau e^{-\int_\tau^{\tau_f} du W_B(u)} \Gamma_{tr}(\tau). \tag{22}$$

A successful realization therefore needs:

1. an early primitive source S_X for a long-lived non-baryonic asymmetry;

2. survival of Y_X until the transfer threshold τ_{tr};

3. opening of the transfer gate while visible washout is weak enough for the converted asymmetry to survive;

4. a transfer channel that maps the stored X asymmetry into $B - L$ or directly into B.

7.3. Why this is distinct from ordinary darkogenesis

This example interpolates between two lessons from the case studies. From the CFS line it borrows the idea that the primitive source need not be baryon number, lepton number, or even an ordinary Noether charge at the first descendant stage; a spectral or occupation imbalance can be the correct source variable. From the Hamada and LQC notes it borrows the insistence that later transfer and visible-sector readout must be written only after their gates are open.

Darkogenesis usually generates a dark-sector Noether asymmetry during a CP-violating first-order transition and only later transfers it. The present template instead begins with a primitive spectral or occupation imbalance that need not be a Noether charge even at the first descendant stage. The distinction is therefore in the source structure, not just in the labels.

The taxonomy applies immediately. If I_{tr} never opens, or if the survival factor in Eq. (22) is exponentially tiny on the interval required for transfer, the proposal is Category I. If a future model writes the source as baryon number from the outset even though the first descendant stage is really X, the result is Category II. If a model is written from the start in the $X \to B$ language of Eqs. (19)–(22), it lands in Category III.

8. Capability map and present uses

The gate-and-kernel package is strongest when presented as a tool for screening, organizing, and generating admissible descendant formulas. Table 2 summarizes the present capability map.

Mechanism family	Primitive source emphasis	Decisive witnesses	Framework output
Hamada gravitational baryogenesis	Visible asymmetry source must be channel-resolved	Current-basis and source–projection witnesses	Category II correction of the published source variable
CFS baryogenesis	Spectral imbalance already primitive	Ordering and geometry already open in the source paper	Category III confirmation
LQC gravitational baryogenesis	Curvature–derivative source in an effective FRW descendant	EFT-control and thermal-current witnesses	Benchmark screening: admissible vs. over–cutoff reductions
Thermal leptogenesis	$B - L$ source from heavy–neutrino decay	Thermal support and standard transport witnesses	Category III interoperability check
Resonant leptogenesis	Quasi–degenerate heavy–neutrino system may require off-diagonal state variables	Coherence witness selects the admissible descendant state space	Reduction selection: diagonal Boltzmann when justified, density matrix otherwise
Two–threshold spectral transfer	Quasiparticle occupation imbalance X precedes visible asymmetry	Source/survival witnesses together with the portal gate I_{tr} and CP-odd transfer efficiency ϵ_B	Forward template: Category III when written from the start in the $X \to B - L \to B$ language; Category I or II if the gate stays closed or the primitive source is misidentified

Table 2. Compact capability map. The framework's job is to decide which source basis, witness set, and transport module are licensed in the stated regime.

The same map yields several immediate uses:

1. **Pre-transport audits.** Before solving any transport equation, one can ask whether the published source variable, current basis, thermal description, or EFT control window has actually been earned in the stated regime.

2. **Benchmark screening.** Different displayed benchmarks in the same paper can land in different categories, separating mechanism-level ideas from benchmark-level control problems.

3. **Descendant-line comparisons.** Naming drift can occur in descendants even when the source paper was already safe. The layered bookkeeping can compare papers inside a research line rather than only across unrelated mechanisms.

4. **Reduction selection in contested regimes.** The resonant-leptogenesis test shows that witness maps can choose between competing descendant reductions before transport is solved.

5. **Organized mechanism generation.** Once the stack is explicit, new mechanism templates can be generated by changing one layer at a time: keep the same background but alter the primitive source; keep the same source but delay the transfer gate; keep the same source and transfer map but move to a different background module.

These uses are deliberately modest. The framework is a screening and scaffolding device, not a final microscopic theory.

9. Conclusion

The reusable part of the case-study comparison is now clear. A universal FRW scaling law is not the right object. What is reusable is one operator formula, one honest domain of validity, and one explicit account of how primitive source, witness-defined gates, transport, and readout fit together.

That is enough to unify the Hamada, CFS, and LQC analyses without forcing them into an algebraic mold they do not share. It is also enough to show three broader things a framework paper should show: a geometry-only reduction that does not compete with spacetime functionalism, a fully opened reduction that hands the problem to standard transport theory, and a prospective witness rule that can choose the admissible descendant state in a harder case such as resonant leptogenesis.

The resulting framework is most useful when treated as a common interface between emergence criteria, background modules, primitive-source choices, and transport methods. Its ambition is therefore limited but real: it provides a repeatable way to audit descendant asymmetry formulas and a disciplined way to design new ones.

Notes

This note isolates reusable gate–and–kernel content behind several companion case studies and places it in a layered architecture. It does not claim a microscopic derivation of transport from quantum gravity and does not claim a unique final theory of emergence.

References

1. ^*Hamada K, Minamizaki A, Sugamoto A (2008). "Baryogenesis by Quantum Gravity." Mod Phys Lett A. 23: 237–244. doi:10.1142/S021773230802639X.*

2. a, b*Harvey J, Turner M (1990). "Cosmological Baryon and Lepton Number in the Presence of Electroweak Fermion Number Violation." Phys Rev D. 42:3344–3349. doi:10.1103/PhysRevD.42.3344.*

3. ^*Finster F, Jokel M, Paganini C (2022). "A Mechanism of Baryogenesis for Causal Fermion Systems." Class Quantum Grav. 39(16):165005. doi:10.1088/1361-6382/ac7a86.*

4. ^*Finster F, van den Beld-Serrano M (2025). "Baryogenesis in Minkowski Spacetime." J Geom Phys. 207:105346. doi:10.1016/j.geomphys.2024.105346.*

5. ^*Fischer P, Paganini C (2026). "Causal Fermion Systems: Spacetime as the Web of Correlations of a Many-Body Quantum System." Class Quantum Grav. 43(4):045005. doi:10.1088/1361-6382/ae2379.*

6. a, b*Odintsov S, Oikonomou V (2016). "Loop Quantum Cosmology Gravitational Baryogenesis." EPL (Europhys Lett). 116(4):49001. doi:10.1209/0295-5075/116/49001.*

7. ^*Davoudiasl H, Kitano R, Kribs G, Murayama H, Steinhardt P (2004). "Gravitational Baryogenesis." Phys Rev Lett. 93:201301. doi:10.1103/PhysRevLett.93.201301.*

8. a, b*Ashtekar A, Singh P (2011). "Loop Quantum Cosmology: A Status Report." Class Quantum Grav. 28:213001. doi:10.1088/0264-9381/28/21/213001.*

9. a, b, c, d*Knox E (2013). "Effective Spacetime Geometry." Stud Hist Philos Mod Phys. 44(3):346–356. doi:10.1016/j.shpsb.2013.04.002.*

10. a, b, c, d*Lam V, Wüthrich C (2018). "Spacetime Is as Spacetime Does." Stud Hist Philos Mod Phys. 64:39–51. doi:10.1016/j.shpsb.2018.04.003.*

11. a, b, c, d*Oriti D (2020). Levels of Spacetime Emergence in Quantum Gravity. Oxford University Press. arXiv:1807.04875.*

12. a, b, c, d*Huggett N, Wüthrich C (2025). Out of Nowhere: The Emergence of Spacetime in Quantum Theories of Gravity. Oxford University Press. doi:10.1093/oso/9780198758501.001.0001.*

13. a, b, c, d*Buchmüller W, Fredenhagen S (2000). "Quantum Mechanics of Baryogenesis." Phys Lett B. **483**:217–224. doi:10.1016/S0370-2693(00)00573-6.*

14. a, b, c, d*Hohenegger A, Kartavtsev A, Lindner M (2008). "Deriving Boltzmann Equations from Kadanoff–Baym Equations in Curved Space-Time." Phys Rev D. **78**:085027. doi:10.1103/PhysRevD.78.085027.*

15. a, b, c, d*Konstandin T (2013). "Quantum Transport and Electroweak Baryogenesis." Phys Usp. **56**(8):747–771. doi:10.3367/UFNe.0183.201308a.0785.*

16. a, b, c, d, e*Kainulainen K, Parkkinen H (2024). "Quantum Transport Theory for Neutrinos with Flavor and Particle-Antiparticle Mixing." J High Energy Phys. **2024**(2):217. doi:10.1007/JHEP02(2024)217.*

17. ^*Mandel Y (2026). "Gravitational Baryogenesis from Entropy Production and Parity Violation: A Unified Framework." arXiv:2601.11690.*

18. a, b*Bödeker D, Buchmüller W (2021). "Baryogenesis from the Weak Scale to the Grand Unification Scale." Rev Mod Phys. **93**:035004. doi:10.1103/RevModPhys.93.035004.*

19. a, b, c*Pilaftsis A, Underwood T (2004). "Resonant Leptogenesis." Nucl Phys B. **692**:303–345. doi:10.1016/j.nuclphysb.2004.05.029.*

20. ^*Garny M, Kartavtsev A, Hohenegger A (2013). "Leptogenesis from First Principles in the Resonant Regime." Ann Phys. **328**:26–63. doi:10.1016/j.aop.2012.10.007.*

21. a, b*Dev P, Millington P, Pilaftsis A, Teresi D (2014). "Flavour Covariant Transport Equations: An Application to Resonant Leptogenesis." Nucl Phys B. **886**:569–664. doi:10.1016/j.nuclphysb.2014.06.020.*

22. ^*Garbrecht B, Gautier F, Klaric J (2014). "Strong Washout Approximation to Resonant Leptogenesis." J Cosmol Astropart Phys. **2014**(09):033. doi:10.1088/1475-7516/2014/09/033.*

23. ^*Shelton J, Zurek K (2010). "Darkogenesis." Phys Rev D. **82**:123512. doi:10.1103/PhysRevD.82.123512.*

24. ^*Co R, Harigaya K (2020). "Axiogenesis." Phys Rev Lett. **124**:111602. doi:10.1103/PhysRevLett.124.111602.*

25. ^*Adshead P, Sfakianakis E (2015). "Fermion Production during and after Axion Inflation." J Cosmol Astropart Phys. **2015**(11):021. doi:10.1088/1475-7516/2015/11/021.*

26. ^*Adshead P, Sfakianakis E (2016). "Leptogenesis from Left-Handed Neutrino Production during Axion Inflation." Phys Rev Lett. **116**(9):091301. doi:10.1103/PhysRevLett.116.091301.*

Supplementary data: available at https://doi.org/10.32388/KTFZST

Declarations