

Heaven & Hell II: Scale Laws and Robustness in One-Step Heaven-Hell Consensus*

Nnamdi Aghanya Romain Leemans
nnamdi.aghanya@cranfield.ac.uk romain.leemans@protonmail.com

October 28, 2025

Abstract

We study the Heaven-Hell dynamics, a model for network consensus. A known result establishes an exact one-step convergence threshold for systems with a single uniform hub: the per-node inbound hub weight W suffices if and only if $W \geq \text{maxrest}$, where maxrest is the maximum non-hub inbound mass. In this work, we develop *scale laws* and *operational refinements* that make this threshold robust to tie-breaking policies, node-specific tolerances, targeted seeding, multiple hubs, and asynchronous update schedules. Our main contributions are: (i) a conservation-law perspective that yields a majority-form characterisation and monotonicity principles; (ii) a parameterised tie policy with exact characterisations for TIEGLORY/TIEGNASH/TIESTAY; (iii) tighter *pointwise deg-max* bounds that strictly improve the classical $\text{indeg} \cdot \text{wmax}$ worst-case guarantee; (iv) a one-pass fairness result for asynchronous updates; and (v) practically checkable sufficient conditions for seeded one- and two-step convergence, including multi-hub budget splits. All proofs are mechanised in Coq, and experiments on rings, grids, BA/scale-free graphs, and heterogeneous weighted graphs validate the tightness and show substantial gap closures over prior bounds¹.

1 Introduction

The Heaven-Hell (HH) model describes consensus dynamics on weighted directed graphs, which is a binary State Consensus on weighted directed graphs. In the HH model, each node updates its state based upon a weighted majority vote of all neighbours; however, there exists a single designated hub node that will enforce a Glory state and cause convergence of all other nodes. The paper builds on previously established work that described exact one-step convergence thresholds for uniform-hub systems[1].

This paper establishes scale laws that increase robustness against realistic variations in the system, such as node-specific tolerances, parameterised tie-breaking policies, multi-hub seeding strategies and asynchronous update schedules. Using a conservation-law perspective, we derived majority form characterisation and monotonicity principles that provide tighter pointwise bounds, exact thresholds, and conditions for both seeded and multi-step convergence. We provide empirical validations of the results using diverse families of graphs, including rings, grids, scale-free networks, and heterogeneous weighted graphs. These results confirmed the tightness of our bounds and demonstrated significant improvements over classical worst-case guarantees.

*GUDraBIT Research Technical Report TR-2025-02.

¹<https://github.com/gudfit/HeavenHellII>

2 Preliminaries: The Heaven-Hell Model

We work on a finite vertex set V with a distinguished hub $g \in V$. Let $w : V \times V \rightarrow \mathbb{N}$ be nonnegative weights representing directed edges, and let $\tau : V \rightarrow \mathbb{N}$ be an optional tolerance per node. A configuration (state) is a function $s : V \rightarrow \{\text{GLORY}, \text{GNASH}\}$. We define the inbound weight to a node v from the hub as $\text{hub_weight}(v) = \sum_{u=g} w(u, v)$ and from all other nodes as *rest weight*, $\text{rest_weight}(v) = \sum_{u \neq g} w(u, v)$. In the uniform-hub setting, we assume $\text{hub_weight}(v) = W$ for all $v \neq g$.

A single update step first forces the hub g to GLORY. Then, each node $v \neq g$ adopts the GLORY state if and only if its score from the GLORY-state neighbours dominates its score from the GNASH-state neighbours, with the handling of ties dependent on the active policy.

Notation. We denote by $\text{score}_G(s, v)$ and $\text{score}_N(s, v)$ the weighted inbound score to v from the GLORY and GNASH nodes, respectively, after forcing g to GLORY. We write $\text{total_in}(v) = \sum_u w(u, v)$. The one-step next state under tolerance τ and a TIEGLORY policy is denoted $\text{next_heaven_tau}(s, v)$; see §5 for details on policies.

3 A Conservation Law and a Majority-Form View

A simple but powerful identity underpins our analysis. The total inbound weight to a node is conserved, perfectly partitioning between the scores from nodes in GLORY and GNASH.

Lemma 1 (Score Conservation). *For any state s and node v , the scores after forcing the hub are related by:*

$$\text{score}_G(s, v) + \text{score}_N(s, v) = \text{total_in}(v).$$

Proof. By definition, $\text{total_in}(v) = \sum_{u \in V} w(u, v)$. The scores are given by weighted sums over indicators: $\text{score}_G(s, v) = \sum_{u \in V} w(u, v) \cdot \mathbb{I}[s'(u) = \text{GLORY}]$ and $\text{score}_N(s, v) = \sum_{u \in V} w(u, v) \cdot \mathbb{I}[s'(u) = \text{GNASH}]$, where s' is the state after forcing the hub. Since any node is either in GLORY or GNASH, the indicators satisfy $\mathbb{I}[s'(u) = \text{GLORY}] + \mathbb{I}[s'(u) = \text{GNASH}] = 1$ for all u . The result follows from distributing the sum. $\square$

This conservation law allows us to re-characterise the update rule in a majority form.

Theorem 2 (Majority Form with Tolerance). *For any node $v \neq g$, its next state is GLORY if and only if:*

$$\text{next_heaven_tau}(s, v) = \text{GLORY} \iff 2 \cdot \text{score}_G(s, v) + \tau(v) \geq \text{total_in}(v).$$

Proof. For a non-hub node v , the TIEGLORY update rule is $\text{next_heaven_tau}(s, v) = \text{GLORY} \iff \text{score}_N(s, v) \leq \text{score}_G(s, v) + \tau(v)$. We substitute $\text{score}_N(s, v)$ using Lemma 1: $\text{total_in}(v) - \text{score}_G(s, v) \leq \text{score}_G(s, v) + \tau(v)$. Rearranging the terms yields the equivalent condition $\text{total_in}(v) \leq 2 \cdot \text{score}_G(s, v) + \tau(v)$. $\square$

4 Monotonicity Principles

The majority-form characterisation makes monotonicity properties immediate.

State Monotonicity. If a state t has weakly more GLORY nodes than a state s , then for any node v , its GLORY-score weakly increases and its GNASH-score weakly decreases. Consequently, any node that transitions to GLORY from state s will also do so from state t . The formal proof is deferred to Lemma 9.

Tolerance Monotonicity. If a tolerance function τ' pointwise dominates another, $\tau'(v) \geq \tau(v)$ for all v , then the required hub weight for convergence can only decrease: $\text{maxneed}(\tau') \leq \text{maxneed}(\tau)$, where $\text{maxneed}(\tau) = \max_{v \neq g} (\text{rest_weight}(v) - \tau(v))$. Thus, increasing tolerances helps achieve consensus.

5 Tie Policies and Exact Characterisations

We parameterise the tie-breaking policy by three rules: TIEGLORY, TIEGNASH, and TIESTAY. Let $T(v) = \text{score}_G(s, v) + \tau(v)$ and $R(v) = \text{score}_N(s, v)$ be the total GLORY and GNASH scores for a node $v \neq g$ after forcing the hub.

Lemma 3 (Tie Policies). *The next state of a node $v \neq g$ is determined as follows:*

- TIEGLORY: $\text{next} = \text{GLORY}$ iff $R(v) \leq T(v)$.
- TIEGNASH: $\text{next} = \text{GLORY}$ iff $R(v) < T(v)$.
- TIESTAY: $\text{next} = \text{GNASH}$ if $T(v) < R(v)$, GLORY if $R(v) < T(v)$, otherwise stay $s(v)$.

These distinctions are critical in knife-edge cases and when stability (TIESTAY) is preferred over a bias toward GLORY.

6 Scale Laws: Exact Thresholds and Tighter Bounds

We now establish the main threshold results for one-step convergence.

Theorem 4 (Uniform Hub with Tolerance (Exact)). *Assume $\text{hub_weight}(v) = W$ for all $v \neq g$. One-step convergence to all-GLORY from any initial state is guaranteed if and only if:*

$$W \geq \text{maxneed}(\tau) := \max_{v \neq g} (\text{rest_weight}(v) - \tau(v)).$$

Proof. The system converges in one step if and only if for every state s and every node $v \neq g$, $\text{next_heaven_tau}(s, v) = \text{GLORY}$. This is equivalent to requiring that the condition holds even in the worst-case initial state, where all non-hub nodes start in GNASH. In this state, the GLORY score for a node v is exactly its hub weight, $\text{score}_G(s, v) = \text{hub_weight}(v) = W$, and its GNASH score is its rest weight, $\text{score}_N(s, v) = \text{rest_weight}(v)$. The update rule $\text{score}_N \leq \text{score}_G + \tau(v)$ becomes $\text{rest_weight}(v) \leq W + \tau(v)$ for all $v \neq g$. This per-node condition is equivalent to the aggregated threshold $W \geq \max_{v \neq g} (\text{rest_weight}(v) - \tau(v))$, as formally shown in Lemma 8. State monotonicity ensures that if convergence occurs from the worst-case state, it also occurs from any other state. $\square$

Corollary 5 (Uniform Hub, No Tolerance). *Setting $\tau(v) = 0$ for all v in Theorem 4, one-step all-GLORY is guaranteed if and only if $W \geq \max_{v \neq g} \text{rest_weight}(v)$.*

Pointwise Deg–Max Bound. Let $\text{indeg}_{-g}(v)$ be the number of non-hub in-neighbours of v , and $\text{max_in}_{-g}(v)$ be the maximum inbound edge weight to v from a non-hub. We have $\text{rest_weight}(v) \leq \text{indeg}_{-g}(v) \cdot \text{max_in}_{-g}(v)$. This gives a strictly tighter worst-case bound than the classical global product:

$$\text{maxrest} \leq \max_{v \neq g} \left(\text{indeg}_{-g}(v) \cdot \text{max_in}_{-g}(v) \right) \leq \text{indeg}_{\text{global}} \cdot \text{wmax}_{\text{global}}.$$

7 Seeding and Multi-Hub Variants

Let $H \subseteq V$ be a seed set of nodes initially forced to GLORY. Define your collective influence as $\text{hub_H}(v) := \sum_{u \in H} w(u, v)$ and your opposition as $\text{rest_H}(v) := \sum_{u \notin H} w(u, v)$.

Theorem 6 (Exact One-Step Criterion under Seeded Forcing). *For any seed set H , one-step convergence to all-GLORY is guaranteed if and only if for all $v \notin H$:*

$$\text{hub_H}(v) + \tau(v) \geq \text{rest_H}(v).$$

Proof. ($\Leftarrow$) Assume that the condition holds. For any node $v \notin H$ and any state s , its score from GLORY nodes after forcing H is $\text{score}_G(s, v) \geq \text{hub_H}(v)$, while its score from GNASH is $\text{score}_N(s, v) \leq \text{rest_H}(v)$. The update condition $\text{score}_N \leq \text{score}_G + \tau(v)$ is satisfied because $\text{rest_H}(v) \leq \text{hub_H}(v) + \tau(v) \leq \text{score}_G(s, v) + \tau(v)$.

($\Rightarrow$) Assume that one-step convergence holds. Consider the initial state where all nodes $u \notin H$ are GNASH. For any $v \notin H$, the scores are exactly $\text{score}_G(s, v) = \text{hub_H}(v)$ and $\text{score}_N(s, v) = \text{rest_H}(v)$ (see Lemma 10). Since v must become GLORY, the update condition $\text{score}_N \leq \text{score}_G + \tau(v)$ implies $\text{rest_H}(v) \leq \text{hub_H}(v) + \tau(v)$. □

8 Asynchronous Schedules: One-Pass Fairness

Let $\text{run_sched.tau}(s, \text{sched})$ denote a single sweep in which the nodes are updated sequentially.

Theorem 7 (One-Pass Fairness). *If the one-step domination condition $\text{hub_weight}(v) + \tau(v) \geq \text{rest_weight}(v)$ holds for all $v \neq g$, and a finite schedule visits every non-hub node at least once, then after one pass every $v \neq g$ will be in the GLORY state, regardless of the starting configuration.*

Proof. The domination condition implies that for any state s , $\text{next_heaven.tau}(s, v) = \text{GLORY}$ for all $v \neq g$. When a node v is updated according to the schedule, its state becomes GLORY. By state monotonicity, once a node becomes GLORY, it cannot revert to GNASH in subsequent updates within the same pass, as the set of GLORY nodes only grows. Since the schedule covers all non-hub nodes, each will be updated to GLORY at its turn and will remain so. □

9 Experiments

We summarise the experiments that validate the theoretical results; see Figure 1.

Regular Families. Rings with k -nearest neighbours: The exact threshold is $W^* = 2k$. 2D grid: With degree $d = 4$, the theoretical threshold $W^* = 4$ is matched experimentally (Fig 1A).




Figure 1: Dynamics and scale effects. (A) One-step success probability rises sharply at $W/W^* = 1$ on rings and grids. (B) On heterogeneous graphs, the classical deg-wmax bound diverges with indegree m while the pointwise bound stays tight. (C) Equal multi-hub splits reduce per-hub weight as $1/h$; annotated examples match the text. (D) Under domination, asynchronous one-pass schedules monotonically drive the fraction of Glory to 1.

Heterogeneous Weighted Graphs. On an adversarially constructed graph, our pointwise deg-max bound is exact, while the classical bound is off by a factor of 200 (Fig 1B):

$$\text{maxrest} = 800, \quad \text{pointwise deg-max} = 800 \text{ (exact)}, \quad \text{classical} = 160,000 (\times 200).$$

Seeding and Multi-Hubs. On a ring with $W^* = 6$, a single hub at $W = 5$ fails. However, splitting a budget of $W = 6$ between two hubs ($3 + 3$) or three hubs ($2 + 2 + 2$) and seeding them achieves a one-step all-GLORY, consistent with Theorem 6 (Fig 1C).

Asynchronous Fairness. On a ring and a grid satisfying the one-step condition, schedules that cover all non-hubs in one pass yield all-GLORY across random orderings, matching Theorem 7 (Fig 1D).

10 Related Work and Limitations

The Heaven-Hell model aligns with the broader literature on threshold-based dynamics in networks, including weighted majority dynamics [2] and bootstrap percolation [3, 4]. Our conservation law

and majority-form views extend monotone system analyses [5], providing exact per-node thresholds that contrast with probabilistic bounds in influence maximisation [6]. Unlike heuristics in complex contagion [7], our analysis provides exact conditions for deterministic one-step convergence, bridging theoretical guarantees with practical refinements absent in prior hub-based consensus models [1].

Our one-step guarantees, however, have their bounds and limitations. In particular, one open and important area for further archaeological work is extending these exact results to optimal multi-step policies, or extensible-time varying weights or stochastic observations. The per-node thresholds, although exact in the case of our hub-forcing semantics, would be overly conservative if the hub itself were forced in some way. We offer sufficient conditions for optimal seeded splits; however, the challenge of discovering optimal seed budgets remains an important open combinatorial challenge. Finally, the experiments used in this paper assess canonical topologies. Stress tests on the algorithm, such as those on dynamically varying networks or with real social data, would be helpful for future study in these areas.

11 Conclusion

This work shows that the one-step threshold for Heaven-Hell consensus scales robustly to tolerances, tie policies, seeding, multiple hubs, and asynchronous schedules. A simple conservation-law perspective yields both exact thresholds and sharp, interpretable bounds. Experiments demonstrate their tightness and practical utility.

Artefact. All results are machine-checked in Coq. The accompanying source code includes the formal proofs (e.g., `ScaleLaw`, `Monotonicity`) and Python scripts (`showcase.py`) to reproduce all experiments.

References

- [1] Aghanya N. Heaven & Hell: One-Step Hub Consensus. [Internet]. 2025. [cited 2025 Oct 24]. Available from: <https://arxiv.org/abs/2509.19630>.
- [2] Peleg D. Size bounds for dynamic monopolies. [Internet]. 1998. [cited 2025 Oct 24]. Available from: [https://doi.org/10.1016/S0166-218X\(98\)00043-2](https://doi.org/10.1016/S0166-218X(98)00043-2).
- [3] Chalupa J, Leath PL, Reich GR. Bootstrap percolation on a Bethe lattice. [Internet]. 1979. [cited 2025 Oct 24]. Available from: <https://doi.org/10.1088/0022-3719/12/1/008>.
- [4] Adler J, Lev U. Bootstrap percolation: a renormalisation group approach. [Internet]. 1991. [cited 2025 Oct 24]. Available from: [https://doi.org/10.1016/0378-4371\(91\)90295-N](https://doi.org/10.1016/0378-4371(91)90295-N).
- [5] Smith HL. Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems. [Internet]. 1995. [cited 2025 Oct 24]. Available from: <https://bookstore.ams.org/surv-41-s/>.
- [6] Kempe D, Kleinberg J, Tardos É. Maximizing the spread of influence through a social network. [Internet]. 2003. [cited 2025 Oct 24]. Available from: <https://www.cs.cornell.edu/home/kleinber/kdd03-inf.pdf>.
- [7] Centola D, Macy M. Complex contagions and the weakness of long ties. [Internet]. 2007. [cited 2025 Oct 24]. Available from: <https://www.jstor.org/stable/10.1086/521848>.

A Proofs of Supporting Lemmas

This appendix provides formal proofs for technical lemmas referenced in the main paper.

Lemma 8 (Equivalence of Per-Node and Aggregated Thresholds). *Let W be the uniform hub weight. The per-node condition ($\forall v \neq g, W + \tau(v) \geq \text{rest_weight}(v)$) is equivalent to the aggregated condition $W \geq \max_{v \neq g}(\text{rest_weight}(v) - \tau(v))$.*

Proof. ($\Rightarrow$) Assume $\forall v \neq g, W + \tau(v) \geq \text{rest_weight}(v)$. This is equivalent to $\forall v \neq g, W \geq \text{rest_weight}(v) - \tau(v)$. Since W is greater than or equal to every element in the set $\{\text{rest_weight}(v) - \tau(v)\}_{v \neq g}$, it must be greater than or equal to the maximum element of that set. Thus, $W \geq \max_{v \neq g}(\text{rest_weight}(v) - \tau(v))$.

($\Leftarrow$) Assume $W \geq \max_{v \neq g}(\text{rest_weight}(v) - \tau(v))$. By the definition of maximum, for any specific $v_0 \neq g$, we have $\max_{v \neq g}(\text{rest_weight}(v) - \tau(v)) \geq \text{rest_weight}(v_0) - \tau(v_0)$. By transitivity, $W \geq \text{rest_weight}(v_0) - \tau(v_0)$, which rearranges to $W + \tau(v_0) \geq \text{rest_weight}(v_0)$. Since this holds for any arbitrary v_0 , the per-node condition is satisfied. $\square$

Lemma 9 (State Monotonicity of Scores). *Let states s, t be such that for all u , if $s(u) = \text{GLORY}$ then $t(u) = \text{GLORY}$. Then for any node v , $\text{score}_G(s, v) \leq \text{score}_G(t, v)$ and $\text{score}_N(s, v) \geq \text{score}_N(t, v)$.*

Proof. Let s' and t' be the states after forcing the hub to GLORY. The premise implies that the set of GLORY nodes under s' is a subset of those under t' . The score $\text{score}_G(s, v) = \sum_u w(u, v) \mathbb{I}[s'(u) = \text{GLORY}]$ is a sum of non-negative terms. Since the sum for $\text{score}_G(t, v)$ is over a superset of terms (or the same set), we have $\text{score}_G(s, v) \leq \text{score}_G(t, v)$. The second inequality, $\text{score}_N(s, v) \geq \text{score}_N(t, v)$, follows directly from this and the conservation law (Lemma 1): $\text{total_in}(v) - \text{score}_G(s, v) \geq \text{total_in}(v) - \text{score}_G(t, v)$. $\square$

Lemma 10 (Scores under All-GNASH Initialization). *Let H be a seed set. In an initial state where $s(u) = \text{GNASH}$ for all $u \notin H$, the scores for a node $v \notin H$ after forcing H to GLORY are exactly $\text{score}_G(s, v) = \text{hub_H}(v)$ and $\text{score}_N(s, v) = \text{rest_H}(v)$.*

Proof. After forcing nodes in H to GLORY, the state s' is such that $s'(u) = \text{GLORY}$ if $u \in H$ and $s'(u) = \text{GNASH}$ if $u \notin H$. By definition, the scores for node v are computed from this state s' .

$$\begin{aligned} \text{score}_G(s, v) &= \sum_{u \in V} w(u, v) \mathbb{I}[s'(u) = \text{GLORY}] = \sum_{u \in H} w(u, v) = \text{hub_H}(v). \\ \text{score}_N(s, v) &= \sum_{u \in V} w(u, v) \mathbb{I}[s'(u) = \text{GNASH}] = \sum_{u \notin H} w(u, v) = \text{rest_H}(v). \end{aligned} \quad \square$$