

# Socio-Traffic Thermodynamics:

Integrated Model and Risk-Budgeted Information-Release  
Audit Synthesis

Keiji Yoshimura  
Independent Researcher

May 25, 2026

Version: v0.3.1-integrated-revision-source-license-checked

GitHub repository: [github.com/yokken0907/socio-traffic-thermodynamics](https://github.com/yokken0907/socio-traffic-thermodynamics)

## Abstract

Socio-Traffic Thermodynamics (STT) is a claim-bounded theoretical and numerical framework for interpreting congestion recurrence as a non-equilibrium pattern generated by socially motivated motion on finite-capacity, finite-bandwidth transportation networks. The original manuscript introduced a mesoscopic formulation combining social potential, finite perceptual bandwidth, stochastic car-following dynamics, information compression, and dissipative stop-and-go wave formation in a periodic Optimal Velocity ring-road surrogate. The subsequent v0.2.0–v0.2.9 audit sequence narrowed and strengthened one operational subclaim: in finite-capacity toy networks, route information can reduce mean social cost only when it does not induce excessive synchronized reaction and overload severity. This v0.3.1 integrated revision combines the original model development with the later risk-budgeted information-release audits. The locked synthesis is conservative: risk-budgeted low-rate, coarse, staggered, or guarded information release reduced mean cost relative to a no-information null in the frozen holdout while not increasing overload severity, whereas common precise and delayed common information increased synchronization, overload, and cost. The manuscript does not claim real-world traffic prediction, traffic-policy validation, city-scale deployment readiness, or a universal congestion solution. It should be read as a hypothesis-generating toy-model synthesis about synchronization risk, information release, and buffering architecture under finite capacity.

**Reader guidance and claim boundary.** This manuscript is a claim-bounded integrated revision. The older STT formulation is retained as a model-building scaffold, while the v0.2.x audits restrict the positive claim to tested toy networks and frozen holdout designs. The work is adjacent to classical car-following instability, stop-and-go wave formation, traffic-flow theory, routing-game inefficiency, and information-design effects [1–6]. It is not a calibrated urban simulation, a route-guidance product, a traffic-policy recommendation, or a claim that congestion can be universally eliminated. The allowed claims concern finite-capacity toy networks, risk-budgeted information release, action synchronization, overload severity, and reduced surrogate diagnostics.

## 1 Introduction

Traffic congestion is often framed as an engineering inefficiency to be minimized by capacity expansion, signal control, route guidance, or demand management. That framing is useful but

incomplete. Congestion can emerge endogenously even in the absence of an explicit bottleneck: fluctuations in an otherwise homogeneous stream can grow, destabilize free flow, and form backward-propagating jam waves once the system approaches a critical density or loses string stability [1–3].

Network-level routing adds another complication. Selfish routing can fail to minimize aggregate latency, as formalized by price-of-anarchy results [4]. More information is not automatically better: informational Braess-type effects can worsen aggregate outcomes when better-informed users react in correlated ways [5], whereas carefully designed private or garbled signals can sometimes move equilibrium flows toward lower social cost [6].

STT addresses this combined problem by treating congestion as a recurrent non-equilibrium structure on finite-bandwidth social-traffic networks. The present version does not expand this into a universal social theory. Instead, it integrates the original theoretical formulation with a later audit program that asks a narrower design question: when can information release improve mean efficiency without increasing overload risk relative to a conservative null?

## 2 Model Formulation

### 2.1 Network and state variables

Let  $G = (V, E)$  be a directed transportation network. Let  $\Omega$  denote a large finite or continuum population of agents. Each agent  $i \in \Omega$  is characterized by position  $x_i(t)$ , velocity  $v_i(t)$ , internal state  $s_i(t) \in [-1, 1]$ , route  $p_i(t) \in \mathcal{P}$ , and destination-dependent social potential  $\Phi_i = \Phi(x_i, t)$ .

Define macroscopic fields

$$\rho(x, t), \quad u(x, t), \quad m(x, t) := \mathbb{E}[s_i \mid x_i = x], \quad (1)$$

where  $\rho$  is density,  $u$  is mean velocity, and  $m$  is an order parameter representing local synchronization of internal states.

### 2.2 Microscopic longitudinal dynamics

For each agent,

$$\dot{x}_i = v_i, \quad (2)$$

$$\dot{v}_i = F(h_i, \Delta v_i, s_i; \rho_i, \Phi_i) + \varepsilon_i(t) + \sqrt{2D_v} \xi_i(t), \quad (3)$$

where  $h_i$  is headway,  $\Delta v_i$  is relative velocity,  $F$  is a smooth car-following response,  $\varepsilon_i(t)$  is decentralized residual control, and  $\xi_i(t)$  is white noise.

The internal-state dynamics is modeled as

$$\tau_s \dot{s}_i = -s_i + \tanh(\beta[J(\mathcal{K} * s)_i + \lambda \Phi_i - \kappa C_i]) + \sqrt{2D_s^{\text{eff}}} \zeta_i(t), \quad (4)$$

with perceived congestion cost  $C_i$  and effective internal noise

$$D_s^{\text{eff}} = D_{s,0} + \chi_{\text{sov}} \sigma_{\text{sov}}^2 + \chi_W W_i^{-1}. \quad (5)$$

Here  $W_i$  denotes effective perceptual bandwidth. The term  $\sigma_{\text{sov}}$  is not a normative judgment

about individual freedom; it is a toy-model variance term representing decentralized residual control or idiosyncratic response.

### 2.3 Route choice and information signal

Given information  $\mathcal{I}_i$ , agent  $i$  chooses a path  $p \in \mathcal{P}$  with probability

$$\mathbb{P}(p_i = p \mid \mathcal{I}_i) = \frac{\exp\{-\theta[\mathcal{C}_p - \lambda_\Phi \Phi_p]\}}{\sum_{q \in \mathcal{P}} \exp\{-\theta[\mathcal{C}_q - \lambda_\Phi \Phi_q]\}}, \quad (6)$$

where  $\mathcal{C}_p$  is generalized travel cost and  $\Phi_p$  is destination-adjusted social potential. In the audit sequence, this formulation is specialized into finite-capacity route-choice toy networks with information policies ranging from no-information nulls to common precise, delayed common, coarse, staggered, guarded, and risk-budgeted releases.

### 2.4 Macroscopic closure

Under mean-field closure, a schematic macroscopic form is

$$\partial_t \rho + \nabla \cdot (\rho u) = 0, \quad (7)$$

$$\partial_t (\rho u) + \nabla \cdot (\rho u \otimes u) = -\nabla P(\rho, m) - \frac{\rho}{\tau} [u - U(\rho, m, \Phi)] + \nu \Delta u + \sqrt{2D_{u\rho}} \eta, \quad (8)$$

$$\tau_m \partial_t m = A(\rho - \rho_c)m - Bm^3 + D_m \Delta m + \lambda \Phi - \kappa C(\rho) - m + \sqrt{2D_m^{\text{eff}}} \zeta. \quad (9)$$

This closure is a conceptual reduced model. It is not asserted to be a calibrated continuum theory for any particular city.

## 3 Metastability and Information Buffering

### 3.1 Metastability and jam nucleation

Let  $(\rho_0, u_0, m_0)$  be a homogeneous free-flow equilibrium. Fourier linearization yields

$$\dot{\hat{z}}_k = A_k \hat{z}_k + B_k \hat{\omega}_k + G_k \hat{\xi}_k. \quad (10)$$

The system is linearly stable if  $\Re \lambda(A_k) < 0$  for all  $k$ . It is metastable if the stability margin

$$\alpha := \min_k \{-\Re \lambda_{\max}(A_k)\} \quad (11)$$

is positive but small. In that regime, perturbations decay slowly and rare transitions become more likely. Near a critical stop-and-go mode with amplitude  $a$ ,

$$\dot{a} = -\partial_a \mathcal{V}(a; \rho, \Phi) + \sqrt{2D_{\text{eff}}} \eta(t), \quad (12)$$

with

$$D_{\text{eff}} = D_0 + \chi_{\text{sov}} \sigma_{\text{sov}}^2 + \chi_W W^{-1}. \quad (13)$$

The nucleation-rate heuristic

$$r_{\text{nuc}} \asymp \exp\left[-\frac{\Delta\mathcal{V}(\rho, \Phi)}{D_{\text{eff}}}\right] \quad (14)$$

expresses the claim that systems near criticality become more vulnerable to fluctuation-driven jam formation.

### 3.2 Variance principle for information buffering

Let  $X$  denote network state and  $Y_i$  the signal received by agent  $i$ . A signal policy is a stochastic kernel  $K(y | x)$ . Path choices induced by  $\{Y_i\}$  generate edge loads

$$F_e = \sum_{i=1}^N \mathbf{1}\{e \in A_i\}. \quad (15)$$

If edge social cost  $SC_e(f) = fc_e(f)$  is convex, then

$$\mathbb{E}[SC(F)] = SC(\bar{F}) + \frac{1}{2} \sum_e SC_e''(\bar{F}_e) \text{Var}(F_e) + o(\text{Var}(F_e)), \quad (16)$$

and

$$\text{Var}(F_e) = \sum_i \text{Var}(A_{ie}) + 2 \sum_{i < j} \text{Cov}(A_{ie}, A_{je}). \quad (17)$$

The key mechanism is therefore not information quantity alone but action covariance. Public high-resolution information can reduce uncertainty while increasing conditional action correlation. Coarse, private, staggered, or rate-limited information can preserve enough adaptation while suppressing synchronized load injection.

## 4 Original Ring-Road Diagnostic

### 4.1 Simulation setting

To test whether the STT transition mechanism appears in a minimal microscopic system, the original manuscript used a one-lane periodic ring with an Optimal Velocity (OV) car-following law [1] and additive residual stochastic input:


$$\dot{x}_i = v_i, \quad (18)$$

$$\dot{v}_i = a[V(h_i) - v_i] + \sigma_{\text{sov}}\xi_i(t), \quad (19)$$

where  $h_i = x_{i+1} - x_i$  is headway on the ring and  $\xi_i(t)$  is a zero-mean unit Gaussian process.

### 4.2 Space-time observation

Figure 1 shows the resulting space-time diagram. The trajectories begin near parallel, indicating approximate uniform flow. As the system evolves, localized bunching emerges and forms a persistent high-density band that propagates backward relative to vehicle motion. This is the canonical signature of a self-sustained stop-and-go wave [2].



**Figure 1:** Space-time diagram of the stochastic OV ring-road simulation. The figure is used as a minimal reduced diagnostic of endogenous stop-and-go wave formation, not as a calibrated city-scale traffic simulation.

### 4.3 Dissipative-structure interpretation

The backward-propagating wave is not interpreted as a conservative traveling mode. It is a dissipative structure sustained by throughput, relaxation delay, and irreversible adaptation. A schematic diagnostic is

$$\Sigma_{\text{irr}} = \int_{\Omega_x} \left[ \alpha_1 |\partial_x u|^2 + \alpha_2 |\partial_t u|^2 + \alpha_3 \text{Var}(u | x) \right] dx. \quad (20)$$

The phrase “dissipative structure” is therefore used in an operationally bounded sense: it denotes an organized non-equilibrium pattern whose existence depends on ongoing flux and dissipation [3].

## 5 Risk-Budgeted Information-Release Audit

### 5.1 Audit program

The v0.2.x program consisted of ten logged phases. Each phase preserved the toy-model boundary and avoided real-world policy interpretation.

**Table 1:** STT v0.2.x audit sequence. The status tokens are displayed in shortened form to avoid layout overclaiming and to keep the table readable.

Phase	Purpose	Status
v0.2.0	Phase diagram and information-buffering preflight	PASS preflight
v0.2.1	Information-synchronization stress test	PASS run
v0.2.2	Desynchronization audit	PASS run
v0.2.3	Capacity-margin audit	PASS run
v0.2.4	Buffer-architecture audit	PASS run
v0.2.5	Mechanism-isolation audit	PASS run
v0.2.6	Topology and demand robustness audit	PASS run
v0.2.7	Risk-weighted Pareto audit	PASS run
v0.2.8	Risk-budgeted information-release audit	PASS run
v0.2.9	Frozen-policy holdout audit	PASS run

The terminal synthesis found all ten expected phase outputs and fixed the decision as the following synthesis-lock status:

PASS-STT-V030-  
SYNTHESIS-LOCKED

## 5.2 Risk-budget formulation

Let  $C$  denote mean social cost and  $S$  denote overload severity. The no-information policy is used as a conservative null baseline with metrics  $(C_0, S_0)$ . A candidate information policy  $\pi$  is admissible under severity budget  $\epsilon_{\text{sev}}$  if

$$S(\pi) - S_0 \leq \epsilon_{\text{sev}}. \quad (21)$$

Among admissible policies, the design objective is

$$\Delta C(\pi) = C_0 - C(\pi) > 0. \quad (22)$$

The strict zero-budget condition  $\epsilon_{\text{sev}} = 0$  requires mean-cost improvement without increasing overload severity relative to the no-information null.

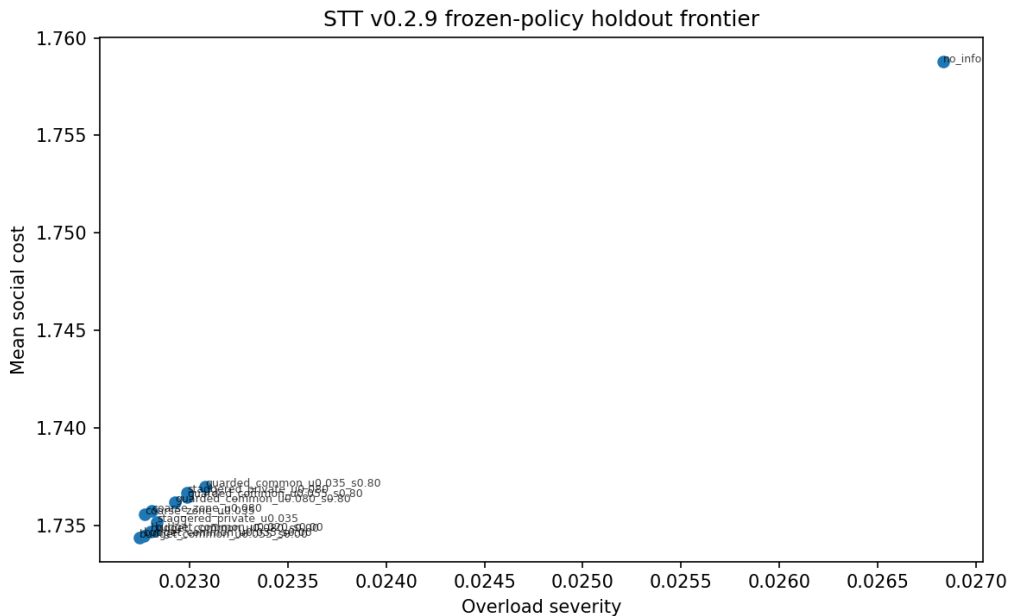
## 5.3 Frozen holdout results

The v0.2.9 audit froze candidate policy families from prior phases and evaluated them on holdout topologies, pulse modes, utilizations, and seeds.

**Table 2:** Representative v0.2.9 frozen-holdout policy results. Positive improvement means lower mean cost than the no-information baseline.

Policy	Mean cost	Severity	Sync	Improvement
budget_common_u0.055_s0.00	1.734	0.0227	0.0028	0.024
budget_common_u0.035_s0.00	1.734	0.0228	0.0019	0.024
coarse_zone_u0.035	1.736	0.0228	0.0041	0.023
staggered_private_u0.035	1.735	0.0228	0.0069	0.024
no_info	1.759	0.0268	0.0000	0.000
common_precise	3.568	0.2177	0.9244	-1.809
delayed_common	6.716	0.4381	0.9391	-4.957

The best mean-cost frozen policy was `budget_common_u0.055_s0.00`. Its mean cost was 1.734372 compared with 1.758804 for no information, giving an improvement of 0.024432. Its overload severity was 0.022745, which was 0.004088 lower than the no-information baseline. By contrast, common precise and delayed common information substantially increased synchronization, overload, and cost.



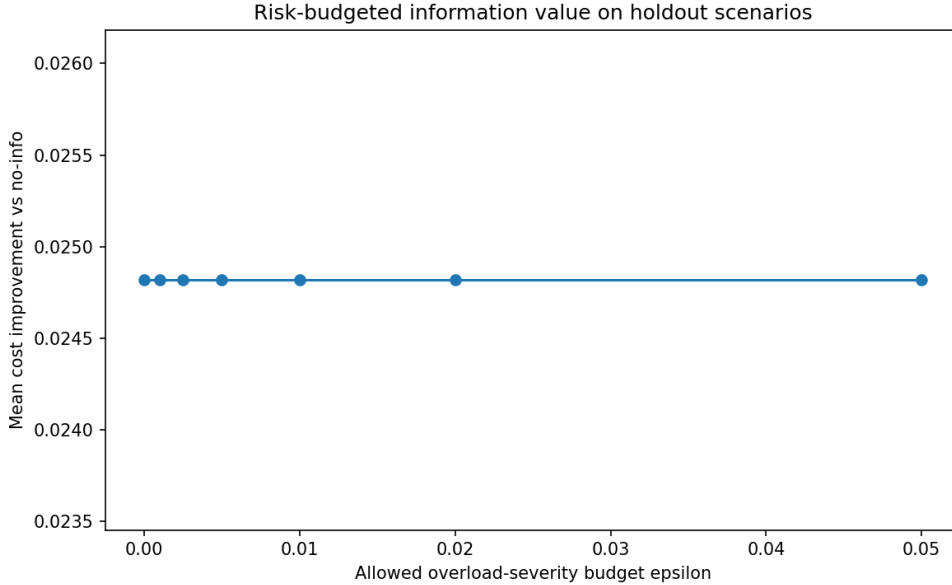
**Figure 2:** Frozen-policy holdout frontier from STT v0.2.9. Candidate risk-budgeted policies lie close to the lower-left region, while the no-information baseline is conservative but slightly less efficient.

## 5.4 Zero-budget finding

Under the strict zero-budget condition, all 27 holdout conditions had at least one policy that improved mean cost without increasing overload severity relative to no information. The winning families were `budget_common` and `coarse_zone`.

**Table 3:** Zero-budget holdout winners.

Winner family	Win count	Win fraction	Mean improvement	Mean sync index
<code>budget_common</code>	20	0.741	0.0273	0.0039
<code>coarse_zone</code>	7	0.259	0.0176	0.0040



**Figure 3:** Risk-budgeted information value in the frozen holdout. In this holdout, the best available candidate already improves mean cost at zero additional overload-severity budget. The flat curve is a property of the tested candidate set and holdout design, not a universal traffic-policy claim.

## 6 Architectural Interpretation

The integrated result narrows the original architectural interpretation. STT does not imply that all information should be restricted, nor that high-precision information is intrinsically harmful. The observed mechanism is synchronization risk under finite capacity. Shared precise information can reduce individual uncertainty while increasing conditional action correlation. If many agents respond in the same direction at the same time, finite-capacity links can be pushed into overload.

This suggests four buffering directions, each still at the hypothesis-generating level:

1. **Temporal buffering:** reduce abrupt synchronization gradients rather than only lowering average demand.
2. **Spatial buffering:** use transfer basins, reservoirs, or staging regions to reduce sudden injection into metastable corridors.
3. **Informational buffering:** treat routing guidance as a variance-control channel rather than a universal shortest-path oracle.
4. **Behavioral buffering:** use damping-capable agents or protocols to reduce oscillatory amplification.

## 7 Limitations

The strongest limitation is model class. The ring-road simulation is a minimal stochastic OV surrogate. The v0.2.x network audits are toy networks, not calibrated urban simulations. The policies are algorithmic abstractions rather than deployable route-guidance mechanisms. The

metrics do not include equity, compliance heterogeneity, public trust, legal constraints, safety, multimodal behavior, or real sensor latency. Therefore, the result is a hypothesis-generating diagnostic claim rather than an operational traffic-policy claim.

A second limitation is baseline definition. The no-information baseline is conservative in this toy design, whereas real systems already contain heterogeneous information sources, institutional controls, learned route preferences, and external routing platforms. Any empirical extension must define the baseline carefully.

## 8 Next Validation Tier

The next validation tier should not be another small toy audit. It should use externally defined synthetic networks or public benchmark traffic scenarios with frozen candidate policies and predeclared metrics. The minimum extension should preserve: (i) a no-information or status-quo baseline, (ii) an overload-risk metric, (iii) a mean-efficiency metric, (iv) a synchronization metric, and (v) a frozen holdout protocol.

## 9 Conclusion

The integrated STT result is narrower but stronger than the original broad framing. Congestion can be modeled as a non-equilibrium dissipative pattern in socially attractive, finite-capacity networks, and the original ring-road surrogate illustrates endogenous stop-and-go wave formation without a fixed bottleneck. The later risk-budgeted audit adds a more operational lesson: in finite-capacity toy networks, information release should be treated as a synchronization-control problem under overload-risk constraints, not merely as a route-precision maximization problem.

## Data and Code Availability

The manuscript is accompanied by a repository package containing the integrated manuscript source, selected figures, original ring-road scripts, v0.3.0 synthesis outputs, manifest files, claim-boundary documentation, and source-defined evaluation-only license files. The GitHub repository is available at [github.com/yokken0907/socio-traffic-thermodynamics](https://github.com/yokken0907/socio-traffic-thermodynamics). Any future GitHub-Zenodo archive should be generated automatically from the release rather than manually assigning an external identifier.

## Acknowledgments

The author conceptually directed the theory-development and validation workflow. Large language models were used as computational, drafting, and software-assistance tools under direct human supervision. All claim-boundary decisions and release decisions remain the author's responsibility.

## References

- [1] M. Bando, K. Hasebe, A. Nakayama, A. Shibata, and Y. Sugiyama, “Dynamical model of traffic congestion and numerical simulation,” *Physical Review E* **51**, 1035–1042 (1995).
- [2] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, “Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam,” *New Journal of Physics* **10**, 033001 (2008).
- [3] D. Helbing, “Traffic and related self-driven many-particle systems,” *Reviews of Modern Physics* **73**, 1067–1141 (2001).
- [4] T. Roughgarden, “The price of anarchy is independent of the network topology,” *Journal of Computer and System Sciences* **67**(2), 341–364 (2003).
- [5] D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar, “Informational Braess’ paradox: The effect of information on traffic congestion,” *Operations Research* **66**(4), 893–917 (2018).
- [6] S. Das, E. Kamenica, and R. Mirka, “Reducing congestion through information design,” working paper (2017).