Thermodynamics of Cooperation: Necessary Constraints

Keith Lostracco

Abstract

We derive the constraints that necessarily govern resource allocation among multiple self-maintaining open thermodynamic systems sharing a finite-resource environment. Starting from the Second Law of Thermodynamics and the condition that each system actively maintains its structural boundary, we prove that mutual boundary constraints are necessary for any feasible multi-agent allocation to exist. Each constraint carries a computable shadow price quantifying its energetic cost. Using the Generalized Nash Equilibrium Problem (GNEP) framework with thermodynamic cost functions, we show that the constrained system admits a variational equilibrium with uniform shadow prices. We prove that removing constraints under resource collision admits no stable allocation, that the constrained equilibrium is welfare-superior to unconstrained alternatives, and that constraint value scales monotonically with scarcity. Worked examples with asymmetric agents demonstrate that the variational equilibrium outperforms both greedy and equal-split allocations, with welfare gains scaling with agent heterogeneity and resource scarcity.

Contents

1 Introduction

Consider $N \geq 2$ open thermodynamic systems, each actively maintaining a low-entropy internal state against the universal trend toward disorder, that share a finite pool of localized free energy. Each system imports energy from the environment, exports entropy, and allocates resources to maintain its structural boundary. When aggregate demand exceeds supply, the systems' optimization problems become coupled: each agent's feasible set depends on the others' strategies.

This paper derives the mathematical structure of that coupling. We prove that mutual boundary constraints are not optional design choices but necessities (Theorem 5), that each constraint carries a computable shadow price quantifying its energetic cost (Theorem 2), that the constrained system admits a variational equilibrium with uniform shadow prices (Theorem 7), that this equilibrium uniquely maximizes total social welfare over the feasible set (Theorem 9), and that constraint value scales monotonically with scarcity (Corollary 11).

The derivation rests on standard physics and standard optimization theory. The agents are open thermodynamic systems: structures that persist only by continuously importing free energy and exporting entropy. The resource environment is finite (a consequence of the Second Law). Given agents with concave utility functions sharing finite resources, the theorems follow from KKT theory and the GNEP framework.

The results contribute to optimization theory: we show that the Generalized Nash Equilibrium Problem (GNEP) framework, when applied to agents with thermodynamic cost functions, produces a variational equilibrium with uniform shadow prices, a mathematical analog of uniform resource pricing. The constraint necessity result (Theorem 5) and its economic consequences (shadow prices, scarcity scaling) provide the foundation for *Thermodynamic Friction* [1], *Strategic Entropy Injection* [2], *Accumulated Negentropy* [3], and *Cooperative Equilibrium* [4].

The specific contributions of this work are:

1. **Constraint necessity** (Theorem 5). We prove that no feasible joint allocation exists in which all agents simultaneously achieve their unconstrained optima. Boundary constraints are therefore necessary for multi-agent coexistence, not optional design choices.

2. **Computable shadow prices** (Theorem 2). Each boundary constraint carries a Lagrange multiplier measuring, in energy units, its marginal cost to the constrained agent. This provides a real-time, auditable signal of constraint tension.

3. **Variational equilibrium with uniform prices** (Theorem 7). The GNEP admits a variational equilibrium in which all agents face the same shadow price for each shared resource, a structural analog of uniform market pricing derived from the optimization problem rather than imposed by mechanism.

4. **Welfare optimality** (Theorem 9). The variational equilibrium uniquely maximizes total social welfare over the feasible set. Every other feasible allocation, including greedy and equal-split, is strictly dominated.

5. **Scarcity scaling** (Corollary 11). The welfare value of the constraint system is monotonically increasing in scarcity: under abundance, constraints are costless; under scarcity, they become the dominant determinant of system welfare.

All formal results are stated as numbered theorems with full proofs.

The paper proceeds as follows: model and quadratic energy specification, unconstrained and constrained problems, KKT conditions and shadow price, impossibility of unconstrained coexistence, generalized Nash equilibrium, worked example, transition to conflict, discussion, and conclusion.

1.1 Related Work

The mathematical machinery of this paper (constrained optimization, Lagrange multipliers, and Nash equilibria) is classical. Our contribution is the application of that machinery to a specific physical setting (dissipative multi-agent systems under finite resources) and the demonstration that the resulting constraints are necessary rather than conventional.

Generalized Nash equilibrium problems. Von Neumann and Morgenstern [5] created the mathematics of strategic interaction, axiomatizing rational preference as expected utility maximization. Nash [1950; 1951] proved that every finite game has at least one equilibrium in mixed strategies. Rosen [8] extended the framework to games with coupled constraints, the Generalized Nash Equilibrium Problem (GNEP), and proved existence and uniqueness of the Normalized Nash Equilibrium (variational equilibrium) with shared Lagrange multipliers for concave N-person games. Facchinei and Kanzow [9] provide a comprehensive survey of GNEP theory and algorithms. The entire classical framework rests on a foundational abstraction: utility is measured in dimensionless, subjective "utils," a quantity that does not correspond to any physical observable and is not interpersonally comparable. Our GNEP formulation follows standard lines; the novelty is the constraint structure (boundary constraints derived from thermodynamic persistence rather than imposed by market or regulatory design) and the demonstration that the resulting variational equilibrium has a physical interpretation as uniform resource pricing in energy units.

Commons and shared-resource problems. Hardin [10] argued that common-pool resources are inevitably degraded absent external regulation, the "tragedy of the commons." Ostrom [11] refuted this empirically, documenting hundreds of cases where communities successfully self-governed shared resources and distilling design principles for durable commons institutions. But Ostrom's evidence is institutional and empirical, not mathematical. Our results provide a formal optimization-theoretic foundation for her observation: the constraint systems she documented are instances of the boundary constraints that this work proves are necessary for feasible multi-agent allocation. The distinction is that our result is a mathematical necessity (no feasible solution exists without constraints) rather than an empirical regularity.

Ecological economics and energy denomination. A parallel tradition recognized what game theorists did not: energy is the fundamental currency of all biological and economic activity. Lotka [12] proposed the maximum power principle: natural selection favors systems that maximize energy throughput. Georgescu-Roegen [13] argued that the Second Law of Thermodynamics imposes fundamental constraints on economic activity, and that neoclassical production functions written "without regard to dimensions or to other physical constraints" [14] are physically meaningless. Our framework inherits this insistence on physical units: all utility functions, shadow prices, and welfare comparisons are denominated in energy (Joules), making the results interpersonally comparable and physically measurable.

Thermodynamic self-maintenance. The physical model of agents as open thermodynamic systems maintaining low-entropy internal states draws on Schrödinger [15], who identified that living systems persist by "feeding on negative entropy," importing free energy and exporting entropy. Prigogine [16] demonstrated that systems far from thermodynamic equilibrium can spontaneously generate ordered structures through symmetry-breaking instabilities, but did not formalize the boundaries that maintain organized entities as distinct agents or the interactions between them. Friston [17] formalized persistence via Markov blankets (statistical boundaries delineating internal from external states) under the Free Energy Principle, but addressed only the single-agent problem: what one entity does to maintain itself, not what happens when multiple such entities interact with conflicting resource needs. England [18] derived a statistical-mechanical lower bound on heat production during self-replication, establishing that self-replicating systems are thermodynamically expected under driven conditions. Maturana and Varela's concept of *autopoiesis* (the operational closure of self-producing systems) describes the topology of self-maintenance that our boundary integrity parameter B_i quantifies energetically [19]. We use this physical picture to motivate the utility functions and constraint structure, but the mathematical results require only the regularity conditions of Assumption 1 and finite resource endowments; no thermodynamic premises enter the proofs.

2 Model

2.1 Physical Setting

The environment is governed by the Second Law of Thermodynamics: the total entropy of an isolated system never decreases ($\Delta S_{\text{universe}} \geq 0$). Maintaining highly ordered structures requires continuous importation of free energy and exportation of entropy [20; 21; 22]. Any localized structure that maintains internal order must therefore perform continuous thermodynamic work; ceasing that work results in entropy accumulation, boundary dissolution, and irreversible loss of the organized state.[1]

The class of systems we consider is defined by a single observable property: each system persists through active boundary maintenance. This includes biological organisms, ecosystems, institutions, and economies. Any such system faces the same optimization problem: acquire sufficient free energy to maintain boundary integrity in a shared, finite-resource environment.

2.2 Self-Maintaining Open Systems

An entity i is an open thermodynamic system maintaining a boundary, a *Markov blanket* [23; 17], between its internal low-entropy state and the external high-entropy environment. It is characterized by an internal state $\mathbf{s}_i \in \mathcal{S}_i$, a boundary with integrity $B_i > 0$ (a scalar measured in energy units representing the total energetic investment in boundary maintenance), and a metabolism that continuously imports free energy and exports entropy.

Maintaining boundary integrity against environmental entropy requires continuous work: $C_{\text{maintain},i} = \gamma_i B_i$, where $\gamma_i > 0$ is the entropy leakage rate. This baseline metabolic cost of persistence is independent of any external interaction.

[1] A reader versed in non-equilibrium thermodynamics may object that self-maintaining systems *accelerate* entropy rather than resist it: the biosphere degrades solar free energy into thermal radiation far more efficiently than bare rock. This perspective (developed by Prigogine, refined by England, and implicit in Lotka's maximum power principle) is entirely compatible with our model. Locally, the entity must perform anti-entropic work to maintain its boundary ($\gamma_i B_i > 0$); globally, that work exports more entropy than it prevents ($\Delta S_{\text{universe}} > 0$). Our model operates at the local scale, where the entity's optimization problem is identical regardless of whether its existence is *permitted* or *favored* by global entropy dynamics. The global perspective explains *why* self-maintaining entities exist; our model determines *how* they must interact once they do.

Each entity's persistence objective translates into a constrained optimization problem: acquire sufficient energy to pay maintenance costs, repair external damage, and retain a surplus buffer. The utility function satisfies standard regularity conditions:

Assumption 1 (Regularity)**.** For each agent i:

(a) U_i is twice continuously differentiable (C^2) on $\mathbb{R}^n_{\geq 0}$.

(b) U_i is strictly concave in $\mathbf{x}_i$: $\nabla^2_{\mathbf{x}_i} U_i \prec 0$ (negative definite Hessian).

(c) U_i is monotonically increasing at the origin: $\nabla_{\mathbf{x}_i} U_i(\mathbf{0}) \succ \mathbf{0}$.

(d) $U_i(\mathbf{x}_i) \to -\infty$ as $\|\mathbf{x}_i\| \to \infty$ (coercivity).

Condition (b) reflects diminishing thermodynamic returns: processing increasing quantities of any single resource incurs escalating metabolic costs. Condition (c) ensures that a zero-resource state is never optimal for a persisting entity: an agent with no resources cannot maintain its boundary. Condition (d) guarantees that the unconstrained optimum exists at a finite point: sufficiently extreme consumption of any resource eventually incurs net thermodynamic harm (toxicity, metabolic overload, storage costs), so utility cannot grow without bound.

2.3 Resource Environment and Strategy Space

Consider $N \geq 2$ agents sharing an environment with n resource dimensions and finite endowment $\mathbf{R} = (R_1, \ldots, R_n) \in \mathbb{R}^n_{>0}$. Finiteness is physically grounded: localized, usable free energy is bounded in any finite spatial region [24]. Each agent i selects a strategy vector $\mathbf{x}_i \in \mathbb{R}^n_{\geq 0}$, and the physically realizable joint strategy must satisfy resource conservation: $\sum_{i=1}^N x_{ij} \leq R_j$ for all j.

Collision is not an edge case; it is the default condition of multi-agent existence under finite resources.

3 The Quadratic Energy Model

The physical setting, notation, and regularity conditions are established in §2. For the derivations and worked examples that follow, we adopt a concrete utility specification satisfying Assumption 1, the **quadratic energy model**:

$$U_i(\mathbf{x}_i) = \sum_{j=1}^{n} \left(\alpha_{ij}\, x_{ij} - \frac{\beta_{ij}}{2}\, x_{ij}^2 \right)$$

where $\alpha_{ij} > 0$ is the marginal energy yield of resource j to agent i and $\beta_{ij} > 0$ governs the diminishing-returns rate. More general concave forms (CES, Cobb-Douglas) work identically for the structural results below.

4 The Unconstrained Problem (No Boundary Constraints)

4.1 Formulation

If no other agents exist (or, equivalently, if no mutual constraints are recognized), each agent solves independently:

$$\max_{\mathbf{x}_i \geq \mathbf{0}} U_i(\mathbf{x}_i)$$

Under Assumption 1, this has a unique interior solution at:

$$\nabla_{\mathbf{x}_i} U_i(\mathbf{x}_i^\circ) = \mathbf{0}$$

For the quadratic model:

$$x_{ij}^\circ = \frac{\alpha_{ij}}{\beta_{ij}} \quad \forall j$$

This is the agent's **unconstrained optimum**: the strategy it would pursue if it were alone in the universe.

4.2 The Collision Condition

When $N \geq 2$ agents share the environment, the physically realizable joint strategy must satisfy the **resource conservation constraint** (scarcity):

$$\sum_{i=1}^{N} x_{ij} \leq R_j \quad \forall j \in \{1, \dots, n\}$$

Definition 2 (Resource Collision)**.** A resource dimension j is in *collision* if the sum of unconstrained optima exceeds supply:

$$\sum_{i=1}^{N} x_{ij}^\circ > R_j$$

Lemma 1 (Inevitability of Collision)**.** *Given a finite resource endowment* $\mathbf{R}$ *shared by* $N \geq 2$ *agents with overlapping resource needs, collision is guaranteed whenever* N *is sufficiently large. Specifically, for any resource* j *essential to all agents, collision occurs whenever* $N \geq N_j^* := \lfloor R_j / \min_i x_{ij}^\circ \rfloor + 1$*.*

Proof. By Assumption 1(b)–(d), each agent i has a unique finite unconstrained optimum with $x_{ij}^\circ > 0$ for every essential resource j. Let $\delta_j := \min_i x_{ij}^\circ > 0$. Then for any N agents sharing resource j:

$$\sum_{i=1}^{N} x_{ij}^\circ \geq N \cdot \delta_j$$

Setting $N \cdot \delta_j > R_j$ gives $N > R_j/\delta_j$, so $N \geq \lfloor R_j/\delta_j \rfloor + 1 =: N_j^*$ suffices. For $N \geq N_j^*$, the aggregate unconstrained demand exceeds supply and resource j is in collision by Definition 2. In any biologically realistic scenario, where per-capita endowment R_j/N is small relative to individual metabolic requirements x°_{ij}, this bound is satisfied for many or all resource dimensions. $\square$

4.3 The Unconstrained Regime

When collision occurs and no constraints exist, each agent pursues $\mathbf{x}^\circ_i$ regardless of the physical feasibility of the aggregate. The result is **simultaneous attempted over-extraction**: multiple agents' strategies claim the same physical resource units.

Since two agents cannot exclusively consume the same unit of energy, the collision is resolved by **conflict**, the direct expenditure of stored energy to contest the resource. This generates the interaction costs (thermodynamic friction) quantified in the thermodynamic friction derivation.

The unconstrained case is therefore a high-friction regime where every agent's strategy imposes an unmitigated externality on every other agent's feasible set.

5 The Constrained Problem (Inter-Agent Constraints as Boundary Conditions)

5.1 Defining a Boundary Constraint

We now introduce the central construction.

Definition 3 (Boundary Constraint)**.** A *boundary constraint* $\mathcal{R}_{B,k}$ of agent B is a constraint of the form

$$g_{Bk}(\mathbf{x}_A) \leq 0 \qquad (k = 1, \ldots, m_B)$$

imposed on every other agent $A \neq B$'s optimization problem, where $g_{Bk} : \mathbb{R}^n_{\geq 0} \to \mathbb{R}$ is a continuously differentiable, convex function encoding a specific aspect of agent B's protected boundary, satisfying $g_{Bk}(\mathbf{0}) < 0$ (the zero-action strategy violates no agent's boundary constraints).

Scope note. The condition $g_{Bk}(\mathbf{0}) < 0$ restricts Definition 3 to *negative boundary constraints*, i.e., non-interference constraints (agent A cannot extract agent B's allocated resources, cannot penetrate B's boundary). *Positive boundary constraints*, i.e., provision constraints (agent A must supply B with at least τ_B units), have the opposite sign structure: $g_{Bk}(\mathbf{0}) = \tau_B > 0$, so the null strategy itself violates the constraint. The present framework derives negative boundary constraints as the first-order consequence of resource collision among persistence-directed agents; positive boundary constraints arise as stability conditions on an already-established variational equilibrium (see Future Work).

Interpretation: The function g_{Bk} defines a **forbidden region** in agent A's action space. If $g_{Bk}(\mathbf{x}_A) > 0$, then strategy $\mathbf{x}_A$ violates agent B's k-th boundary constraint. The constraint $g_{Bk}(\mathbf{x}_A) \leq 0$ forces agent A to remain outside B's protected domain.

Examples of constraint functions:

Boundary Constraint (plain language)	Constraint function $g_{Bk}(\mathbf{x}_A) \le 0$	Meaning
B has a guaranteed minimum share $\bar{x}_{Bj}$ of resource j	$g_{Bk}(\mathbf{x}_A) = x_{Aj} - (R_j - \bar{x}_{Bj})$	A cannot take more than $R_j - \bar{x}_{Bj}$ of resource j
B's total resource intake cannot be reduced below threshold τ_B	$g_{Bk}(\mathbf{x}_A) = \sum_j x_{Aj} - (\lvert\mathbf{R}\rvert - \tau_B)$, where $\lvert\mathbf{R}\rvert = \sum_j R_j$	A's total consumption is bounded
B's spatial boundary Ω_B is inviolable	$g_{Bk}(\mathbf{x}_A) = \mathbb{1}[\mathbf{x}_A \cap \Omega_B \neq \emptyset]$	A's action vector cannot penetrate B's physical/informational boundary

> **Note.** The spatial boundary example uses an indicator function for conceptual clarity. Since $\mathbb{1}[\cdot]$ is neither continuously differentiable nor convex, in practice it is replaced by a smooth convex approximation (e.g., a distance-based penalty for convex Ω_B) satisfying Definition 3's regularity requirements.

The framework is deliberately general: **any** restriction on one agent's strategy that protects another agent's domain is a boundary constraint. This includes resource allocation bounds, spatial exclusion zones, and information channel constraints, all encoded as inequality constraints.

5.2 The Constrained Optimization Problem

With boundary constraints in place, agent A's optimization problem becomes:

$$\max_{\mathbf{x}_A \geq \mathbf{0}} \; U_A(\mathbf{x}_A) \qquad \text{subject to:}$$

$$\text{(i) Resource scarcity:} \quad x_{Aj} + \hat{x}_{-Aj} \leq R_j \quad \forall j \qquad [\lambda_j \geq 0]$$

$$\text{(ii) Boundary constraints of other agents:} \quad g_{Bk}(\mathbf{x}_A) \leq 0 \quad \forall B \neq A,\; k = 1, \ldots, m_B \qquad [\mu_{Bk} \geq 0]$$

where $\hat{x}_{-Aj} = \sum_{i \neq A} x_{ij}$ is the aggregate consumption of resource j by all other agents (taken as given from A's perspective), and λ_j, μ_{Bk} are the associated Lagrange multipliers.

5.3 The Lagrangian

Assembling the Lagrangian for agent A:

$$\mathcal{L}_A(\mathbf{x}_A, \boldsymbol{\lambda}, \boldsymbol{\mu}) = U_A(\mathbf{x}_A) - \sum_{j=1}^{n} \lambda_j \left(x_{Aj} + \hat{x}_{-Aj} - R_j\right) - \sum_{B \neq A} \sum_{k=1}^{m_B} \mu_{Bk}\, g_{Bk}(\mathbf{x}_A)$$

where $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_n)$ and $\boldsymbol{\mu} = \{\mu_{Bk}\}_{B \neq A,\, k}$.

6 The Karush-Kuhn-Tucker (KKT) Conditions

6.1 Necessary and Sufficient Conditions

Under Assumption 1 and standard constraint qualification, the KKT conditions are both necessary and sufficient for optimality (concave objective, convex constraints; see [25, §5.5.3]). Slater's condition holds for agent A's subproblem at any feasible joint profile: $g_{Bk}(\mathbf{0}) < 0$ for all B, k by Definition 3, and continuity then gives $g_{Bk}(\epsilon\mathbf{1}/n) < 0$ for sufficiently small $\epsilon > 0$; the shared resource constraint $x_{Aj} + \hat{x}_{-Aj} \leq R_j$ admits a strictly feasible interior point whenever $\hat{x}_{-Aj} < R_j$, which holds at any feasible profile of the GNEP. At the optimum $\mathbf{x}_A^*$:

(i) Stationarity:

$$\left.\frac{\partial U_A}{\partial x_{Aj}}\right|_{\mathbf{x}_A^*} = \lambda_j + \sum_{B \neq A} \sum_{k=1}^{m_B} \mu_{Bk} \left.\frac{\partial g_{Bk}}{\partial x_{Aj}}\right|_{\mathbf{x}_A^*} \qquad \forall j$$

(ii) Primal feasibility:

$$\begin{aligned} x_{Aj}^* + \hat{x}_{-Aj} \leq R_j & \qquad \forall j \\ g_{Bk}(\mathbf{x}_A^*) \leq 0 & \qquad \forall B \neq A,\ \forall k \end{aligned}$$

(iii) Dual feasibility:

$$\lambda_j \geq 0 \quad \forall j, \qquad \mu_{Bk} \geq 0 \quad \forall B, k$$

(iv) Complementary slackness:

$$\begin{aligned} \lambda_j \left(x_{Aj}^* + \hat{x}_{-Aj} - R_j\right) = 0 & \qquad \forall j \\ \mu_{Bk}\, g_{Bk}(\mathbf{x}_A^*) = 0 & \qquad \forall B, k \end{aligned}$$

6.2 Interpretation of the Stationarity Condition

The stationarity condition (i) has a precise physical reading:

$$\underbrace{\frac{\partial U_A}{\partial x_{Aj}}}_{\substack{\text{Marginal energy} \\ \text{gain from resource } j}} = \underbrace{\lambda_j}_{\substack{\text{Shadow price of} \\ \text{resource scarcity}}} + \underbrace{\sum_{B,k} \mu_{Bk} \frac{\partial g_{Bk}}{\partial x_{Aj}}}_{\substack{\text{Aggregate cost of} \\ \text{other agents' constraints}}}$$

At the optimum, the marginal benefit of consuming one more unit of resource j exactly equals the marginal cost, decomposed into two components: the scarcity cost (the resource is finite) and the boundary constraints cost (other agents' boundaries restrict access).

This formalizes the framework's core structure: *every boundary constraint protecting agent B acts as an equal and opposite restriction on agent A's feasible set.*

7 The Shadow Price of a Boundary Constraint

7.1 The Lagrange Multiplier as the Cost of a Boundary Constraint

Theorem 2 (The Shadow Price Theorem)**.** *Let $\mathbf{x}_A^*(\boldsymbol{c})$ denote agent A's optimal strategy as a function of the constraint parameters $\boldsymbol{c} = \{c_{Bk}\}$, where the k-th boundary constraint of agent B is parameterized as $g_{Bk}(\mathbf{x}_A) \leq c_{Bk}$. Then the optimal value function $U_A^*(\boldsymbol{c}) = U_A(\mathbf{x}_A^*(\boldsymbol{c}))$ satisfies:*

$$\frac{\partial U_A^*}{\partial c_{Bk}} = \mu_{Bk}^* \geq 0$$

where μ_{Bk}^* is the optimal Lagrange multiplier associated with agent B's k-th boundary constraint. That is, the multiplier μ_{Bk}^* measures the marginal increase in agent A's optimal energy intake per unit relaxation of agent B's k-th boundary constraint.

Proof. This is a direct application of the shadow price interpretation of Lagrange multipliers [see 25, §5.6.3]. Under Assumption 1 and constraint qualification, the optimal value function is differentiable in the constraint parameters, and the derivative equals the associated multiplier. □

7.2 Physical Interpretation

Corollary 3 (Active vs. Inactive Boundary Constraints)**.** *Every boundary constraint has a non-negative cost. A boundary constraint that is not active (not binding at the optimum) has zero cost. A boundary constraint that is active (binding) has strictly positive cost.*

This follows directly from complementary slackness and dual feasibility:

- If $g_{Bk}(\mathbf{x}_A^*) < 0$ (the constraint is slack, i.e., A is not pressing against B's boundary), then $\mu_{Bk}^* = 0$. The boundary constraint exists but imposes no cost because A's unconstrained optimum already respects it.
- If $g_{Bk}(\mathbf{x}_A^*) = 0$ (the constraint is active, i.e., A is at B's boundary), then $\mu_{Bk}^* > 0$. The boundary constraint actively restricts A's strategy and has a quantifiable energetic cost.

Corollary 4 (Coupled Cost of Active Constraints)**.** *Every active boundary constraint has a strictly positive shadow price $\mu_{Bk}^* > 0$. The same constraint that protects agent B's domain restricts agent A's feasible set, and this restriction is quantified by μ_{Bk}^* in energy units.*

7.3 The "Price" of a Boundary Constraint in the Quadratic Model

For the quadratic energy model with a single resource and the constraint $g_B(x_A) = x_A - (R - \bar{x}_B) \leq 0$ (agent B is guaranteed at least $\bar{x}_B$):

The Lagrangian:

$$\mathcal{L}(x_A, \mu) = \alpha_A x_A - \frac{\beta_A}{2} x_A^2 - \mu\big(x_A - (R - \bar{x}_B)\big)$$

KKT stationarity:

$$\alpha_A - \beta_A x_A - \mu = 0$$

Case 1: Constraint not binding ($x_A^\circ = \alpha_A/\beta_A \leq R - \bar{x}_B$).

Agent A's unconstrained optimum already respects B's boundary constraint. Then $\mu^* = 0$ and $x_A^* = \alpha_A/\beta_A$. The boundary constraint exists but costs A nothing.

Case 2: Constraint binding ($\alpha_A/\beta_A > R - \bar{x}_B$).

Agent A would prefer more than is allowed. The constraint forces $x_A^* = R - \bar{x}_B$ and:

$$\mu^* = \alpha_A - \beta_A(R - \bar{x}_B) > 0$$

The shadow price μ^* is the marginal energy that A forfeits because of B's boundary constraint. This is:

- **Increasing** in α_A (the more A values the resource, the costlier the constraint).
- **Decreasing** in β_A (the faster A's returns diminish, the less it cares about the cap).
- **Increasing** in $\bar{x}_B$ (the more B is guaranteed, the more A must forfeit).
- **Decreasing** in R (the more total resource exists, the less binding the constraint).

These are physically intuitive: boundary constraints are costlier when entities are more capable, more "hungry," and when resources are scarcer.

8 The Impossibility of Unconstrained Coexistence

Theorem 5 (Impossibility of Infinite Freedom)**.** *Let $N \geq 2$ agents share a finite resource endowment $\mathbf{R}$. Suppose each agent has at least one essential resource (a resource j for which $\partial U_i/\partial x_{ij}|_{\mathbf{x}_i=\mathbf{0}} > 0$) in common with at least one other agent, and that aggregate unconstrained demand exceeds supply for at least one resource ($\sum_i x_{ij}^\circ > R_j$ for some j). Then no feasible strategy profile $(\mathbf{x}_1, \ldots, \mathbf{x}_N)$ can simultaneously be unconstrained-optimal for all agents. Formally:*

$$\nexists\, (\mathbf{x}_1, \ldots, \mathbf{x}_N) \in \mathcal{F} \ \text{ such that } \ \mathbf{x}_i = \mathbf{x}_i^\circ \ \ \forall i$$

where $\mathcal{F} = \{(\mathbf{x}_1, \ldots, \mathbf{x}_N) : \sum_i x_{ij} \leq R_j \ \forall j, \ \mathbf{x}_i \geq \mathbf{0} \ \forall i\}$ is the feasible set and $\mathbf{x}_i^\circ$ is agent i's unconstrained optimum.

Proof. Let j^* be a resource dimension that is essential for agents A and B (at minimum). By Assumption 1(c) and strict concavity, both agents have strictly positive unconstrained optima: $x_{Aj^*}^\circ > 0$ and $x_{Bj^*}^\circ > 0$.

Consider the aggregate unconstrained demand:

$$\sum_{i=1}^{N} x_{ij^*}^\circ \geq x_{Aj^*}^\circ + x_{Bj^*}^\circ$$

For the claim to fail, i.e., for all agents to achieve their unconstrained optima simultaneously, we would need:

$$\sum_{i=1}^{N} x^{\circ}_{ij^*} \leq R_{j^*}$$

But as N grows or as agents' capacities increase relative to R_{j^*}, this inequality is violated. More precisely, for any fixed R_{j^*} and any collection of agents each requiring $x^{\circ}_{ij^*} > 0$, there exists a finite N^* such that $\sum_{i=1}^{N^*} x^{\circ}_{ij^*} > R_{j^*}$.

In any biologically relevant system (multiple organisms sharing finite local resources), the collision condition $\sum_i x^{\circ}_{ij} > R_j$ is satisfied for at least one j, making the simultaneous unconstrained-optimal profile infeasible. □

Corollary 6 (Necessity of Active Constraints)**.** *In any multi-agent system with resource collision, at least one boundary constraint must be active ($\mu^*_{Bk} > 0$ for some B, k) at any feasible solution. Boundary constraints are not optional; they are necessary for coexistence.*

Interpretation: Unconstrained coexistence is impossible in any shared finite-resource environment. An agent can achieve its unconstrained optimum only if no other agent competes for any of the same resources.

9 The Multi-Agent Coupled System: Generalized Nash Equilibrium

9.1 Formulation

When all N agents optimize simultaneously, each taking the others' strategies into account, the system is a **Generalized Nash Equilibrium Problem (GNEP)**. Each agent i solves:

$$\max_{\mathbf{x}_i \geq \mathbf{0}} U_i(\mathbf{x}_i) \qquad \text{s.t.} \quad \sum_{l=1}^{N} x_{lj} \leq R_j \;\; \forall j, \quad g_{Bk}(\mathbf{x}_i) \leq 0 \;\; \forall B \neq i, \; k$$

The feasible set for each agent depends on the other agents' strategies (through the shared resource constraints). This is what makes the problem a *generalized* Nash problem, not a standard one.

9.2 Solution Concept: Variational Equilibrium

A **Variational Equilibrium** (VE), also called a Normalized Nash Equilibrium [8], is a strategy profile $(\mathbf{x}^*_1, \ldots, \mathbf{x}^*_N)$ such that:

1. Each agent i's strategy $\mathbf{x}^*_i$ is optimal given the others' strategies, and
2. All agents share the **same** Lagrange multipliers $\boldsymbol{\lambda}^* = (\lambda^*_1, \ldots, \lambda^*_n)$ for the shared resource constraints.

Theorem 7 (Existence and Structure of the Variational Equilibrium). *Under Assumption 1, the GNEP possesses a Variational Equilibrium* $(\mathbf{x}_1^*, \ldots, \mathbf{x}_N^*, \boldsymbol{\lambda}^*)$*. At this equilibrium:*

(a) The shared shadow prices λ_j^* *are identical for all agents: every agent faces the same "price" for each scarce resource.*

(b) The stationarity condition for each agent i *reads:*

$$\left.\frac{\partial U_i}{\partial x_{ij}}\right|_{\mathbf{x}_i^*} = \lambda_j^* + \sum_{B \neq i,k} \mu_{Bk}^{(i)} \left.\frac{\partial g_{Bk}}{\partial x_{ij}}\right|_{\mathbf{x}_i^*} \qquad \forall j$$

(c) No agent can improve its utility by unilaterally deviating from $\mathbf{x}_i^*$*, given the constraints.*

Proof. The shared resource constraints $0 \leq x_{ij} \leq R_j$ give each agent a compact, convex feasible set, each U_i is concave by Assumption 1, and the joint constraint set is convex. These are the hypotheses of [8, Theorem 3], which guarantees that a Normalized Nash Equilibrium exists for every weight vector $r > 0$. The shared-multiplier property is the defining feature of the VE solution concept. □

9.3 Interpretation: Equality Before the Constraints

Corollary 8 (Symmetry of Scarcity). *At the variational equilibrium, all agents experience the same marginal cost* λ_j^* *per unit of each shared resource. No agent is privileged in the structure of the shared constraints.*

The scarcity constraint imposes a uniform shadow-price structure on all agents. Individual differences arise only through differences in utility functions U_i (different agents value resources differently) and through the specific boundary constraints g_{Bk} (different agents have different protected domains).

9.4 Welfare Optimality of the Variational Equilibrium

Theorem 9 (Welfare Optimality). *Under Assumption 1, the variational equilibrium* $(\mathbf{x}_1^*, \ldots, \mathbf{x}_N^*)$ *is the unique maximizer of total social welfare over the feasible set:*

$$\sum_{i=1}^{N} U_i(\mathbf{x}_i^*) > \sum_{i=1}^{N} U_i(\mathbf{x}_i) \qquad \forall\, (\mathbf{x}_1, \ldots, \mathbf{x}_N) \in \mathcal{F} \setminus \{\mathbf{x}^*\}$$

That is, the VE uniquely maximizes $W(\mathbf{x}) = \sum_{i=1}^{N} U_i(\mathbf{x}_i)$ *subject to the shared resource and boundary constraints.*

Proof. The argument is by direct KKT equivalence. Because each U_i depends only on agent i's own action vector $\mathbf{x}_i$ (coupling between agents enters solely through the shared resource and boundary constraints, not through the utilities), the gradient of the social welfare objective decomposes agent-wise:

$$\frac{\partial W}{\partial x_{ij}} = \frac{\partial U_i}{\partial x_{ij}} \qquad \forall i, j.$$

Writing the KKT conditions for the welfare problem $\max_{\mathbf{x} \in \mathcal{F}} \sum_i U_i(\mathbf{x}_i)$ with multipliers λ_j on the shared resource constraints and $\mu_{Bk}^{(i)}$ on the boundary constraints that restrict agent i's action yields, for every i, j:

$$\frac{\partial U_i}{\partial x_{ij}} = \lambda_j + \sum_{B \neq i, k} \mu_{Bk}^{(i)} \frac{\partial g_{Bk}}{\partial x_{ij}},$$

together with primal/dual feasibility and complementary slackness. These are *identical* to the VE stationarity conditions of Theorem Theorem 7(b), with the shared resource multipliers λ_j playing the role of the common shadow prices λ_j^* that defines the variational (normalized) equilibrium.

Under Assumption 1 and Slater's condition (established in the KKT discussion above), these conditions are both necessary and sufficient for optimality of the welfare problem [25, §5.5.3]. Hence the VE allocation satisfies the welfare-maximization KKT system, and conversely any welfare maximizer is a VE.

Uniqueness follows from strict concavity: each U_i is strictly concave by Assumption 1(b), so $W = \sum_i U_i$ is strictly concave on the convex, compact feasible set $\mathcal{F}$, admitting at most one maximizer. Existence is guaranteed by continuity of W on $\mathcal{F}$. Therefore the VE is the unique global maximizer of total social welfare over $\mathcal{F}$. (This uniform-weight coincidence between the normalized Nash equilibrium and the social welfare optimum is the separable-utility case of [8, §5].) □

Corollary 10 (Welfare Dominance)**.** *The variational equilibrium is welfare-superior to every other feasible allocation, including equal-split, greedy (unconstrained-optimal for a single agent), and any arbitrary feasible partition of resources.*

9.5 The "Mutual Constraint" Structure

At the variational equilibrium, the full system of KKT conditions across all agents forms a coupled system:

$$\frac{\partial U_i}{\partial x_{ij}} = \lambda_j^* + \sum_{B \neq i, k} \mu_{Bk}^{(i)} \frac{\partial g_{Bk}}{\partial x_{ij}} \qquad \forall i, j$$

$$\sum_{i=1}^{N} x_{ij}^* \leq R_j, \quad \lambda_j^* \geq 0, \quad \lambda_j^* \left(\sum_i x_{ij}^* - R_j \right) = 0 \qquad \forall j$$

$$g_{Bk}(\mathbf{x}_i^*) \leq 0, \quad \mu_{Bk}^{(i)} \geq 0, \quad \mu_{Bk}^{(i)} g_{Bk}(\mathbf{x}_i^*) = 0 \qquad \forall i, B \neq i, k$$

This coupled system is a self-consistent set of mutual constraints that no agent can improve upon unilaterally, a fixed point of the multi-agent optimization. The system of mutual boundary constraints is not imposed exogenously; it emerges as the necessary structure of any stable multi-agent coexistence under scarcity.

10 Worked Example: Two Agents, One Divisible Resource

10.1 Setup

Let $N = 2$, $n = 1$ (agents A and B, one resource of total quantity R). The quadratic utility functions are:

$$U_A(x_A) = \alpha_A x_A - \frac{\beta_A}{2} x_A^2, \qquad U_B(x_B) = \alpha_B x_B - \frac{\beta_B}{2} x_B^2$$

where $x_B = R - x_A$ (by resource conservation, taking the constraint as binding for simplicity, since this is the interesting case).

10.2 Unconstrained Optima

$$x_A^\circ = \frac{\alpha_A}{\beta_A}, \qquad x_B^\circ = \frac{\alpha_B}{\beta_B}$$

Collision condition: $x_A^\circ + x_B^\circ > R$, i.e., $\frac{\alpha_A}{\beta_A} + \frac{\alpha_B}{\beta_B} > R$.

10.3 With Boundary Constraints: Agent B Has a Guaranteed Minimum $\bar{x}_B$

Agent A maximizes $U_A(x_A)$ subject to $x_A \leq R - \bar{x}_B$.

Lagrangian:

$$\mathcal{L}(x_A, \mu) = \alpha_A x_A - \frac{\beta_A}{2} x_A^2 - \mu\,(x_A - R + \bar{x}_B)$$

KKT conditions:

$$\alpha_A - \beta_A x_A^* - \mu^* = 0$$
$$\mu^* \geq 0, \quad x_A^* \leq R - \bar{x}_B, \quad \mu^*(x_A^* - R + \bar{x}_B) = 0$$

Solution:

If $\frac{\alpha_A}{\beta_A} \leq R - \bar{x}_B$ (constraint not binding):

$$x_A^* = \frac{\alpha_A}{\beta_A}, \qquad \mu^* = 0$$

If $\frac{\alpha_A}{\beta_A} > R - \bar{x}_B$ (constraint binding):

$$x_A^* = R - \bar{x}_B, \qquad \mu^* = \alpha_A - \beta_A(R - \bar{x}_B)$$

10.4 Numerical Illustration

Let $\alpha_A = 10$, $\beta_A = 1$, $\alpha_B = 8$, $\beta_B = 1$, $R = 12$, $\bar{x}_B = 5$.

- Unconstrained optima: $x_A^\circ = 10$, $x_B^\circ = 8$. Collision: $10 + 8 = 18 > 12$.
- With B's boundary constraint: A can take at most $R - \bar{x}_B = 7$.
- Since $x_A^\circ = 10 > 7$, the constraint binds: $x_A^* = 7$, $\mu^* = 10 - 1 \cdot 7 = 3$.
- **The cost of B's boundary constraint to A:** Agent A loses $\mu^* = 3$ units of marginal energy due to B's guaranteed share.
- For these parameters the choice $\bar{x}_B = 5$ recovers the variational equilibrium ($\lambda^* = (10 + 8 - 12)/(1 + 1) = 3$, giving $x_A^* = 7$, $x_B^* = 5$). The welfare comparisons that follow therefore illustrate Theorem 9 rather than a coincidence of parameter choice.
- Agent A's constrained utility: $U_A(7) = 10(7) - \frac{1}{2}(49) = 70 - 24.5 = 45.5$.
- Agent A's unconstrained utility: $U_A(10) = 10(10) - \frac{1}{2}(100) = 100 - 50 = 50$.
- **Total cost to A:** $50 - 45.5 = 4.5$ energy units.

10.5 Social Welfare Comparison

With boundary constraints ($\bar{x}_B = 5$):

- $U_A(7) = 45.5$, $U_B(5) = 8(5) - \frac{1}{2}(25) = 27.5$.
- **Total system utility:** $45.5 + 27.5 = 73.0$.

Without boundary constraints (collision → conflict, modeled as equal split $x_A = x_B = 6$):

- $U_A(6) = 10(6) - \frac{1}{2}(36) = 42$, $U_B(6) = 8(6) - \frac{1}{2}(36) = 30$.
- **Total system utility:** $42 + 30 = 72.0$.

Without boundary constraints (A takes all, unconstrained):

- $U_A(10) = 50$, $U_B(2) = 8(2) - \frac{1}{2}(4) = 14$.
- **Total system utility:** $50 + 14 = 64.0$.

The boundary-constrained allocation achieves the highest total system utility. The unconstrained "greedy" solution is the worst for the system.

> **Remark (Why the gap is small here).** The welfare difference between the VE allocation (73.0) and the naive equal split (72.0) is only ~1.4%. This is a *mathematical artifact of near-symmetric agents*: when $\beta_A = \beta_B$, the VE allocation $x_i^* = (\alpha_i - \lambda^*)/\beta$ differs from equal split only through the α_i differences. With $\alpha_A = 10$ and $\alpha_B = 8$ (a 20% gap), the heterogeneity is small and equal split is nearly optimal. In general, the welfare gain from optimally differentiated allocation scales with agent heterogeneity. The following asymmetric example demonstrates this.

10.6 Asymmetric Example: Heterogeneous Metabolic Efficiency

In biological systems, agents rarely have identical metabolic profiles. Consider two agents with the **same marginal yield** (α) but **different diminishing-returns rates** (β): for example, Agent A is an efficient large-scale processor (forests, apex predators, industrial economies) while Agent B saturates quickly (specialized microorganisms, artisanal producers).

Parameters: $\alpha_A = 10$, $\beta_A = 0.5$ (slow diminishing returns); $\alpha_B = 10$, $\beta_B = 2$ (fast diminishing returns); $R = 12$.

Unconstrained optima: $x_A^\circ = 20$, $x_B^\circ = 5$. Collision: $25 > 12$.

Variational Equilibrium:

$$\lambda^* = \frac{x_A^\circ + x_B^\circ - R}{1/\beta_A + 1/\beta_B} = \frac{20 + 5 - 12}{2 + 0.5} = \frac{13}{2.5} = 5.2$$

$$x_A^* = \frac{10 - 5.2}{0.5} = 9.6, \qquad x_B^* = \frac{10 - 5.2}{2} = 2.4$$

The VE allocates **four times more** to the efficient processor, not because A is "more important," but because A extracts more marginal value from each additional unit.

Social welfare comparison:

Allocation	x_A	x_B	U_A	U_B	**Total SW**
VE (optimal boundary constraints)	9.6	2.4	72.96	18.24	**91.20**
Greedy (A takes all)	12	0	84.00	0.00	84.00
Equal split	6	6	51.00	24.00	75.00

The VE now outperforms equal split by **21.6%** and the greedy allocation by **8.6%**. The equal-split regime is particularly wasteful: it forces 6 units onto Agent B, whose fast-saturating metabolism means each unit beyond ˜2.4 costs more to process than it yields. Biologically, this is analogous to force-feeding an organism past its metabolic optimum: raw resource input does not equate to usable energy output.

Key insight: The framework does not merely prove that *some* constraint system beats anarchy. It proves that the *optimal* constraint system (the VE) allocates resources according to each agent's capacity to convert them into useful work, and this allocation can dramatically outperform equal-split allocation precisely when agents are diverse.

10.7 Symmetric Example: Same-Species Competition and the Scarcity Gradient

A ubiquitous biological scenario is intraspecific competition: agents of the **same species** sharing the **same environment** ($\alpha_A = \alpha_B = \alpha$, $\beta_A = \beta_B = \beta$). Here the framework shows a different but equally important result.

Parameters: $\alpha = 10$, $\beta = 1$; resource R varies.

Unconstrained optima: $x_A^\circ = x_B^\circ = 10$. **Collision** occurs whenever $R < 20$.

VE for symmetric agents: By symmetry of the KKT system, the variational equilibrium assigns each agent exactly half:

$$x_A^* = x_B^* = \frac{R}{2}, \qquad \lambda^* = \alpha - \beta \cdot \frac{R}{2}$$

Equal allocation is not assumed; it is derived. The equal split emerges as a theorem from the symmetry of identical utility functions under shared scarcity constraints. No exogenous distributional rule is required.

The cost of defection: If one agent defects (seizes its unconstrained optimum $x_A = 10$, leaving the other with $R - 10$), the total system welfare drops. The magnitude of this loss scales with scarcity:

R	Scarcity	VE: x_i^*	VE: U_i	VE: SW	Defect: (x_A, x_B)	Defect: SW	**VE Advantage**
20	None	10.0	50.0	100.0	(10, 10)	100.0	0%
16	Mild	8.0	48.0	96.0	(10, 6)	92.0	**4.3%**
14	Moderate	7.0	45.5	91.0	(10, 4)	82.0	**11.0%**
12	Scarce	6.0	42.0	84.0	(10, 2)	68.0	**23.5%**
10	Severe	5.0	37.5	75.0	(10, 0)	50.0	**50.0%**

Three results emerge from this table:

1. **Under abundance ($R \geq 20$), boundary constraints have zero cost**: both agents achieve their unconstrained optima, all multipliers are zero, and the constraint system is slack. Boundary constraints exist in principle but impose no restriction. This is the formal counterpart of the observation that resource conflicts rarely arise when resources are plentiful.

2. **As scarcity increases, the cost of defection grows super-linearly.** Moving from mild to severe scarcity, the VE advantage jumps from 4% to 50%. This is because the defector's gain (moving from $R/2$ to 10) is bounded by the concavity of U, while the victim's loss (moving from $R/2$ to $R - 10$) accelerates as it approaches zero, the region where each lost unit is maximally valuable.

3. **Under severe scarcity ($R = 10$), defection eliminates the other agent entirely** ($x_B = 0 \implies U_B = 0$). The defector captures $U_A = 50$ (a gain of 12.5 over the VE), but agent B loses its entire share of 37.5, yielding a net welfare destruction of 25 units. This is the signature of Thermodynamic Friction [1]: unrestricted competition in scarce environments does not redistribute value, it *dissipates* it.

Corollary 11 (Boundary Constraint Value Scales with Scarcity)**.** *The value of a boundary constraint system is monotonically increasing in scarcity. Under abundance, boundary constraints are costless and unnecessary. Under scarcity, boundary constraints become the dominant determinant of system welfare.*

The scarcity-monotonicity result provides a mathematical basis for the prediction that multi-agent systems will intensify coordination under resource pressure: the welfare cost of failed coordination grows with scarcity.

11 The Transition to Conflict: Removing the Boundary Constraints

11.1 Statement

Proposition 12 (Unconstrained Regime Implies Conflict)**.** *If all boundary constraints are removed (g_{Bk} dropped for all B, k), and the system is in collision ($\sum_i x^\circ_{ij} > R_j$ for some j), then:*

(a) No stable allocation exists without an enforcement mechanism.

(b) Each agent's attempt to realize its unconstrained optimum generates a negative externality on other agents.

(c) The resulting contested resource must be resolved by direct energy expenditure (conflict), creating the interaction cost function quantified in the thermodynamic friction derivation.

Proof. Proof of (a). Without constraints, each agent i attempts $\mathbf{x}^\circ_i$. The aggregate $\sum_i \mathbf{x}^\circ_i > \mathbf{R}$ is physically infeasible. Since no constraint directs agents to reduce consumption, and each agent's unilateral best response is $\mathbf{x}^\circ_i$ regardless of the others' (by strict concavity, the optimal response function without constraints is constant), there is no self-correcting mechanism. The unconstrained Nash equilibrium is infeasible: it requires $\sum_i \mathbf{x}^\circ_i > \mathbf{R}$. Any feasible allocation requires at least one agent to accept less than its optimum, which it has no incentive to do without a constraint. □

Proof of (b). When A claims x°_{Aj}, the remaining resource for all other agents is $R_j - x^\circ_{Aj}$. In the collision regime, $R_j - x^\circ_{Aj} < \sum_{i \neq A} x^\circ_{ij}$, meaning A's strategy forces at least one other agent below its optimum. This is the definition of a negative externality. □

Remark. Part (c) connects directly to the thermodynamic friction derivation [1]. When the constraint is absent, the physical reality of resource exclusion (two agents cannot consume the same molecule) is enforced not by coordination but by direct energy expenditure: adversarial contest over the disputed resource. The constraint $g_{Bk}(\mathbf{x}_A) \leq 0$ replaces costly physical enforcement with costless (zero-friction) boundary recognition.

12 Summary of Results

#	Result	Statement	Significance
T1	Shadow Price Theorem	$\partial U_A^* / \partial c_{Bk} = \mu_{Bk}^* \geq 0$	Every boundary constraint has a quantifiable energetic cost to the constrained agent
C1.1	Active vs. Inactive boundary constraints	Binding constraints have $\mu > 0$; slack constraints have $\mu = 0$	Boundary constraints cost nothing when not contested; cost emerges only at points of conflict
C1.2	Coupled Cost of Active Constraints	Every active boundary constraint has $\mu_{Bk}^* > 0$, simultaneously protecting B and restricting A	The same multiplier measures protection and restriction
L1	Inevitability of Collision	Collision guaranteed for $N \geq N_j^*$ agents sharing finite resources	Scarcity makes constraint necessity quantitative, not merely qualitative
T2	Impossibility of Infinite Freedom	No feasible profile achieves all agents' unconstrained optima simultaneously	Boundary constraints are necessary, not optional
C2.1	Necessity of Active Constraints	At least one $\mu_{Bk}^* > 0$ at any feasible solution	Some agent is always constrained
T3	Variational Equilibrium	GNEP has a solution with shared shadow prices	Uniform shadow-price structure emerges from the shared constraint structure
C3.1	Symmetry of Scarcity	All agents face the same λ_j^* per shared resource	Uniform cost structure; individual differences arise only through U_i and g_{Bk}
T4	Welfare Optimality	VE uniquely maximizes total social welfare over the feasible set	The VE is the best feasible allocation, not merely a stable one
C4.1	Welfare Dominance	VE is welfare-superior to every other feasible allocation	Greedy, equal-split, and all other alternatives are strictly dominated

#	Result	Statement	Significance
C4.2	Boundary Constraint Value Scales with Scarcity	Value of boundary constraints monotonically increasing in scarcity	Coordination intensifies during crises as a consequence of the model
P1	Unconstrained → Conflict	Removing constraints in collision regime produces no stable allocation	The unconstrained regime is unstable; links to thermodynamic friction

13 Notation Index

Symbol	Meaning
N	Number of agents
n	Number of resource dimensions
$\mathbf{x}_i \in \mathbb{R}^n_{\geq 0}$	Agent i's strategy (resource-allocation) vector
R_j	Total available quantity of resource j
$U_i(\mathbf{x}_i)$	Agent i's utility (boundary-maintenance objective)
α_{ij}, β_{ij}	Marginal yield and diminishing-returns parameters (quadratic model)
$\mathbf{x}_i^\circ$	Agent i's unconstrained optimum
$g_{Bk}(\mathbf{x}_A)$	Constraint function encoding agent B's k-th boundary constraint
λ_j	Lagrange multiplier for resource j scarcity
μ_{Bk}	Lagrange multiplier (shadow price) for agent B's k-th boundary constraint
$\hat{x}_{-Aj}$	Aggregate consumption of resource j by all agents other than A
$\mathcal{L}_A$	Lagrangian for agent A's constrained optimization
$\mathcal{F}$	Feasible strategy set under resource and boundary constraints

Symbol disambiguation. The symbol R_j denotes total resource quantity. In subsequent papers in this series, r denotes coupling distance [26], contest decisiveness [1], and time-preference rate [4]. Context and subscripts distinguish these uses.

14 Discussion

14.1 Physical Interpretation

The results form a complete logical chain. Scarcity creates collision (Lemma 1); collision requires boundary constraints (Theorem 5); the constrained system admits a variational equilibrium with uniform shadow prices (Theorem 7); and constraint value scales monotonically with scarcity (Corollary 11).

Each result is denominated in energy units and computationally verified. The shadow prices of Theorem 2 provide an auditable, real-time signal of constraint tension, a computable measure of how costly a given constraint is to the constrained agent.

The central contribution is the demonstration that boundary constraints in multi-agent dissipative systems are not design choices or social conventions but necessities: the GNEP has no feasible solution without them.

14.2 Testable Predictions

The model generates falsifiable predictions with explicit failure conditions:

Prediction	Source	Falsified If...
Uniform shadow prices across agents for each shared resource at the VE	Theorem 7	Resource markets with physical delivery constraints consistently exhibit agent-specific (non-uniform) shadow prices under conditions satisfying Assumption 1
Welfare gain from constraints is monotonically increasing in scarcity	Corollary 11	Cooperation rates or welfare measures decrease as resource scarcity increases, controlling for other variables
Unconstrained regime is strictly dominated for all scarcity levels $R < \sum_i x_i^\circ$	Theorem 5, worked examples (§10)	Experimental resource-sharing games with energy-denominated payoffs find that unconstrained allocation outperforms the constrained VE allocation

The strongest falsification of the model's core result would be a class of physically realistic multi-agent systems, with concave utility functions and finite shared resources satisfying Assumption 1, in which a feasible allocation exists that is simultaneously unconstrained-optimal for all agents. The model's architecture makes this structurally impossible under resource collision, but a counterexample would require revision of Theorem 5.

14.3 Limitations

The model inherits standard idealizations. The utility functions satisfy regularity conditions (Assumption 1) that may not hold for all physical systems; in particular, strict concavity and coercivity exclude agents with increasing returns to scale or threshold effects. The two-agent worked examples illustrate the general N-agent theorems but do not capture heterogeneous or large-N populations. The static resource endowment assumption excludes renewable, seasonal, or stochastic resource environments. The restriction to negative boundary constraints ($g_{Bk}(\mathbf{0}) < 0$) means the model addresses non-interference but not provision obligations (see §14.5).

The constraint necessity result is a feasibility argument: it shows that no feasible allocation exists without constraints, but does not specify the mechanism by which constraints are established or enforced. *Thermodynamic Friction* [1] quantifies the energy cost of constraint violation, and *Cooperative Equilibrium* [4] derives game-theoretic enforcement via repeated interaction.

Scope. The model applies to any system of agents satisfying Assumption 1 that share finite resources, without restriction on species, substrate, or scale. But the results address the *structural rules* of multi-agent allocation: constraint necessity, shadow pricing, and equilibrium existence. Questions of distributive justice beyond Pareto efficiency, the political process by which specific constraint levels are negotiated, and the enforcement architecture that maintains constraints in practice lie outside the optimization problem this paper defines.

14.4 Applications

The model applies wherever multiple self-maintaining systems share finite resources:

1. **Ecological carrying capacity.** Species competing for overlapping resource niches face the collision condition of Lemma 1. The boundary constraints correspond to territorial boundaries, niche partitioning, and behavioral dominance hierarchies, mechanisms that ecologists observe empirically. The model proves that some constraint system is necessary for stable coexistence under resource collision; these ecological mechanisms are empirical instances of such constraints. The scarcity-scaling corollary predicts that niche-partitioning intensity should increase with resource scarcity.

2. **Common-pool resource management.** Ostrom [11] documented institutional constraint systems (harvest limits, access rules, monitoring) that communities develop to manage shared fisheries, irrigation systems, and forests. These are instances of mutually enforced boundary constraints. The scarcity-scaling corollary (Corollary 11) predicts that such institutions intensify as the resource base declines.

3. **Economic markets.** The variational equilibrium's uniform shadow prices are a mathematical analog of competitive market prices. Within the model, the shared Lagrange multipliers play the role conventionally assigned to market-clearing prices, providing one possible microfoundation that does not invoke an auctioneer mechanism. The worked examples (§10) show that the VE allocation outperforms both greedy and equal-split regimes, with the welfare advantage increasing with agent heterogeneity.

14.5 Future Work

Positive boundary constraints and endogenous production. The model is restricted to non-interference conditions ($g_{Bk}(\mathbf{0}) < 0$). Provision constraints, where $g_{Bk}(\mathbf{0}) > 0$ so that inaction itself violates the constraint, arise as stability conditions on an already-established variational equilibrium. Formally, this requires replacing $g_{Bk}(\mathbf{0}) < 0$ with direct Slater's condition (non-empty feasible-set interior) and deriving when equilibrium maintenance requires active resource transfer toward vulnerable agents. The historical pattern in which non-interference constraints tend to precede provision constraints in legal and institutional development is consistent with the ordering the model's constraint hierarchy predicts.

Extending the model to a production network in which agents' outputs feed other agents' inputs would convert positive boundary constraints into a derived result: whenever agent B's output is essential input for agent A, the provision threshold $\tau_B > 0$ appears as an active constraint at any sustainable equilibrium. The shadow price structure would split into a scarcity component λ_j^* and a flow-dependency component π_{ij}^*, the latter measuring the marginal downstream value of production. Such an extension would also allow the model to engage directly with Georgescu-Roegen [13]'s critique that neoclassical production functions ignore thermodynamic constraints.

Dynamic and stochastic environments. The present results assume static resource endowments. Extensions to time-varying resources $\mathbf{R}(t)$, stochastic scarcity, and Bayesian updating on opponent types [27] would test the robustness of the variational equilibrium under distributional uncertainty. The shadow price structure should extend to stochastic settings via robust optimization or chance-constrained programming.

Heterogeneous agent populations. The two-agent worked examples illustrate the general N-agent theorems but do not capture the full complexity of large heterogeneous populations. Computational studies with diverse agent types, varying metabolic efficiency, resource needs, and constraint structures, would test the practical scalability of the variational equilibrium.

15 Conclusion

We have shown that boundary constraints are not design choices but necessary for multi-agent dissipative optimization. When multiple self-maintaining open thermodynamic systems share a finite-resource environment, no feasible joint allocation exists in which all agents simultaneously achieve their unconstrained optima (Theorem 5). Each boundary constraint carries a computable shadow price measuring its energetic cost to the constrained agent (Theorem 2), and the constrained system admits a variational equilibrium in which all agents face the same shadow price for each shared resource (Theorem 7). This equilibrium is not merely stable; it uniquely maximizes total social welfare over the feasible set (Theorem 9). The welfare value of the constraint system scales monotonically with scarcity (Corollary 11): under abundance, constraints are costless; under scarcity, they become the dominant determinant of system welfare. All results are denominated in energy units, obtained entirely within the framework of constrained optimization (KKT conditions, GNEP, and variational equilibrium theory), and computationally verified.

These results establish the mathematical foundation for the subsequent papers in this series. *Thermodynamic Friction* [1] quantifies the energy cost of constraint violation. Strategic entropy injection [2] formalizes the cost of information corruption. Accumulated negentropy [3] establishes the irreplaceability of complex dissipative structures. Cooperative equilibrium [4] proves that cooperation is the unique efficient Nash Equilibrium under energy-denominated payoffs. Value dynamics [26] characterizes stable coexistence as a dynamical attractor. The constraint necessity and shadow price results derived here, the proof that boundary constraints are unavoidable and that each carries a computable cost, provide the structural foundation on which all subsequent results depend.

References

[1] Keith Lostracco. Thermodynamics of cooperation: Thermodynamic friction, 2026. URL https://doi.org/10.5281/zenodo.19635850. doi:10.5281/zenodo.19635850.

[2] Keith Lostracco. Thermodynamics of cooperation: Strategic entropy injection, 2026. URL https://doi.org/10.5281/zenodo.19635814. doi:10.5281/zenodo.19635814.

[3] Keith Lostracco. Thermodynamics of cooperation: Accumulated negentropy, 2026. URL https://doi.org/10.5281/zenodo.19635836. doi:10.5281/zenodo.19635836.

[4] Keith Lostracco. Thermodynamics of cooperation: Cooperative equilibrium, 2026. URL https://doi.org/10.5281/zenodo.19635856. doi:10.5281/zenodo.19635856.

[5] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[6] John F. Nash. Equilibrium points in N-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. doi:10.1073/pnas.36.1.48.

[7] John F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. doi:10.2307/1969529.

[8] J. Ben Rosen. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, 33(3):520–534, 1965. doi:10.2307/1911749.

[9] Francisco Facchinei and Christian Kanzow. Generalized Nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211, 2010. doi:10.1007/s10479-009-0653-x.

[10] Garrett Hardin. The tragedy of the commons. *Science*, 162(3859):1243–1248, 1968. doi:10.1126/science.162.3859.1243.

[11] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.

[12] Alfred J. Lotka. Contribution to the energetics of evolution. *Proceedings of the National Academy of Sciences*, 8(6):147–151, 1922. doi:10.1073/pnas.8.6.147.

[13] Nicholas Georgescu-Roegen. *The Entropy Law and the Economic Process*. Harvard University Press, 1971.

[14] Nicholas Georgescu-Roegen. Comments on the papers by Daly and Stiglitz. In V. Kerry Smith, editor, *Scarcity and Growth Reconsidered*, pages 95–105. Johns Hopkins University Press, Baltimore, 1979.

[15] Erwin Schrödinger. *What Is Life? The Physical Aspect of the Living Cell*. Cambridge University Press, 1944.

[16] Ilya Prigogine. Time, structure, and fluctuations. *Science*, 201(4358):777–785, 1978. doi:10.1126/science.201.4358.777.

[17] Karl Friston. Life as we know it. *Journal of the Royal Society Interface*, 10(86):20130475, 2013. doi:10.1098/rsif.2013.0475.

[18] Jeremy L. England. Statistical physics of self-replication. *The Journal of Chemical Physics*, 139(12):121923, 2013. doi:10.1063/1.4818538.

[19] Humberto R. Maturana and Francisco J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. D. Reidel, 1980. Original work published 1972.

[20] Rudolf Clausius. Ueber verschiedene für die Anwendung bequeme Formen der Hauptgleichungen der mechanischen Wärmetheorie. *Annalen der Physik und Chemie*, 201(7):353–400, 1865. doi:10.1002/andp.18652010702.

[21] Ludwig Boltzmann. Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht. *Wiener Berichte*, 76:373–435, 1877.

[22] Max Planck. *Vorlesungen über Thermodynamik*. De Gruyter, Berlin, 1897. doi:10.1515/9783112341964. URL `https://doi.org/10.1515/9783112341964`.

[23] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

[24] Roger Penrose. *The Road to Reality: A Complete Guide to the Laws of the Universe*. Alfred A. Knopf, 2005.

[25] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[26] Keith Lostracco. Thermodynamics of cooperation: Value dynamics, 2026. URL `https://doi.org/10.5281/zenodo.19635865`. doi:10.5281/zenodo.19635865.

[27] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.