

# Descent into Reality

## From Learning Systems to Field Dynamics

Flyxion

Independent Researcher

April 22, 2026

### Abstract

We develop a unified account of realization across domains by reframing the central problem of intelligence, physics, and economics as one of selection within an admissible space. Modern learning systems reveal that, in regimes of excess capacity, many exact solutions coexist; correctness alone no longer determines outcome. Instead, systems require a secondary principle that selects among admissible configurations.

We formalize this principle through a bias functional that measures residual tension. This functional admits multiple coordinate representations: as curvature in geometric systems, as entropy gradients in thermodynamic systems, and as future work burden in economic systems. These are not independent quantities, but equivalent orderings over admissible configurations.

We show that realization emerges through the interaction of four components: admissibility constraints, decimation of incoherent modes, minimization of residual tension, and monotonic reduction of uncertainty. Continuous dynamics are expressed as gradient flow on the functional, while discrete reasoning systems provide an operational realization through recursive exploration, pruning, and selection.

The resulting framework unifies learning theory, recursive reasoning, and field dynamics under a single principle: among all admissible configurations, those that minimize residual tension and reduce future reconstruction cost are realized. Reality is thus not constructed through accumulation, but through the elimination of incompatible configurations until only coherent structure remains.

# 1 Introduction: From Correctness to Selection

Classical accounts of reasoning, whether in logic, computation, or scientific inference, are organized around the notion of correctness. A system is given a problem, and the task is to identify the correct answer. This paradigm presupposes that correctness is both well-defined and sufficient: once a correct solution is identified, no further structure is required to determine realization.

Recent developments in learning theory disrupt this assumption. In overparameterized regimes, models frequently admit a vast multiplicity of solutions that interpolate the data exactly. The condition of correctness becomes degenerate; many distinct configurations satisfy the same constraints. Yet, in practice, only a small subset of these solutions are realized by learning dynamics. This discrepancy reveals a missing component in the classical picture: correctness alone does not determine outcome. A secondary principle is required to select among admissible configurations.

This phenomenon is not confined to machine learning. In physical systems, many microstates correspond to the same macroscopic constraints, yet only certain configurations emerge as stable realizations. In economic systems, multiple allocations may satisfy local constraints, yet only some persist under dynamic pressures. Across these domains, one observes a common structure: a space of admissible configurations, a degeneracy within that space, and a selection mechanism that resolves this degeneracy.

The central claim of this work is that these selection mechanisms are not domain-specific, but instances of a single underlying principle. We propose that realization is governed by the minimization of residual tension within an admissible space. This tension is not a single scalar quantity fixed across domains, but a functional that admits multiple coordinate representations. In geometric systems, it appears as curvature; in thermodynamic systems, as entropy gradients; in economic systems, as the burden of future work required to maintain coherence.

To formalize this idea, we introduce a bias functional defined over configurations subject to admissibility constraints. The dynamics of realization are then described as a flow that reduces this functional while preserving admissibility. In continuous settings, this corresponds to gradient flow on the functional. In discrete reasoning systems, it corresponds to iterative exploration and pruning of candidate configurations, guided by uncertainty and feasibility.

A key feature of this framework is that it shifts the focus from construction to elimination. Rather than building a correct configuration from elementary steps, the system begins with a space of possibilities and progressively eliminates those that are incompatible with global coherence. What remains is not selected because it was explicitly constructed, but because

it survives the removal of all alternatives.

This perspective allows us to unify several seemingly disparate developments. Excess capacity learning demonstrates the insufficiency of correctness and the necessity of selection [3, 4, 5]. Integral decimation provides a mechanism for eliminating incoherent modes in dynamical systems [2]. Recursive reasoning frameworks implement a discrete analogue of this process through iterative refinement and uncertainty reduction [1, 31, 32]. When viewed through the lens of constraint closure, these become different realizations of the same underlying process.

The remainder of the paper develops this unification. We begin by formalizing the notion of admissibility and degeneracy in learning systems. We then examine the distinction between local admissibility and global realizability. We introduce the bias functional and examine its coordinate representations, leading to a statement of the selection principle governing realization. Finally, we show how both continuous field dynamics and discrete reasoning systems instantiate this principle, and we explore its implications for economic systems where future work becomes the dominant coordinate of residual tension.

The guiding idea throughout is simple but far-reaching: correctness defines what is possible, but selection determines what is real.

## 2 Excess Capacity and Degeneracy of Solutions

We begin by formalizing the breakdown of the classical correctness paradigm in learning systems. Let  $\mathcal{H}$  denote a hypothesis space and let  $\mathcal{D}$  be a dataset. A configuration  $h \in \mathcal{H}$  is said to be *admissible* if it satisfies the imposed constraints, expressed by a general constraint functional  $C$ . We distinguish two levels of admissibility. Let

$$\mathcal{M}_{\text{local}} = \{h \in \mathcal{H} \mid h|_{U_i} \text{ satisfies constraints on each } U_i\},$$

and

$$\mathcal{M}_{\text{global}} = \{h \in \mathcal{M}_{\text{local}} \mid h_i|_{U_i \cap U_j} = h_j|_{U_i \cap U_j}\}.$$

We define the admissible set as

$$\mathcal{M}_{\text{valid}} := \mathcal{M}_{\text{global}} = \{h \in \mathcal{H} \mid C(h) = 0\}.$$

Configurations in  $\mathcal{M}_{\text{local}} \setminus \mathcal{M}_{\text{global}}$  are locally correct but globally obstructed; they are eliminated by the selection process.

In classical regimes,  $\mathcal{M}_{\text{valid}}$  is either empty or contains a small number of isolated solutions. Under such conditions, identifying a correct configuration effectively determines re-

alization. However, in overparameterized systems,  $\mathcal{M}_{\text{valid}}$  becomes large, often forming a high-dimensional manifold. The condition of correctness ceases to distinguish between configurations; it merely defines a constraint surface within which many solutions coexist.

This shift can be understood as a transition from a *selection-limited* regime to a *constraint-limited* regime. In the former, the primary difficulty lies in finding any admissible solution. In the latter, admissibility is abundant, and the difficulty lies in selecting among many equivalent solutions. Empirical results in modern machine learning confirm that training dynamics consistently converge to particular regions of  $\mathcal{M}_{\text{valid}}$ , rather than exploring it uniformly [3, 6]. This indicates the presence of an implicit bias that governs selection.

We formalize this by introducing a functional

$$\mathcal{B}(h),$$

which assigns to each admissible configuration a measure of residual tension. The realized configuration is then given by

$$h^* = \arg \min_{h \in \mathcal{M}_{\text{valid}}} \mathcal{B}(h).$$

The existence of such a functional is not an additional assumption but an empirical necessity: without it, there is no principled explanation for why one admissible configuration is realized over another. The role of  $\mathcal{B}$  is to induce an ordering on  $\mathcal{M}_{\text{valid}}$ , thereby resolving degeneracy.

At this stage,  $\mathcal{B}$  is introduced abstractly. Its specific form depends on the domain and the dynamics of the system. In learning systems, it may correspond to norms or regularization terms that favor simpler or smoother solutions. In physical systems, it may correspond to energy or entropy-related quantities. In economic systems, it may correspond to measures of effort, cost, or instability.

What is essential is not the particular representation, but the role  $\mathcal{B}$  plays: it encodes the system’s preference structure over admissible configurations. The central claim of this work is that, despite differing representations, these functionals are manifestations of a single underlying object when viewed at the appropriate level of abstraction.

The transition from correctness to selection thus introduces two fundamental components. First, a constraint set  $\mathcal{M}_{\text{valid}}$  defined by admissibility conditions. Second, a bias functional  $\mathcal{B}$  that resolves degeneracy within this set. Together, they define a variational problem whose solution determines realization.

This perspective reframes learning as a process of navigating a constrained manifold under an implicit objective that is not explicitly encoded in the loss function. The observed be-

havior of learning systems is then understood as the outcome of minimizing  $\mathcal{B}$  subject to admissibility, rather than merely achieving correctness.

In the following section, we make this structure explicit by examining the distinction between local admissibility and global realizability, a distinction that proves essential for understanding when and why selection succeeds.

### 3 Local Admissibility and Global Obstruction

The admissible set  $\mathcal{M}_{\text{valid}}$  has thus far been treated as a collection of configurations satisfying imposed constraints. However, this definition obscures an important distinction: admissibility is often verified locally, while realization requires global consistency.

Let  $\Omega$  be a domain covered by open sets  $\{U_i\}$ . A configuration  $X$  may satisfy all constraints when restricted to each  $U_i$ , yet fail to assemble into a globally consistent object on  $\Omega$ . This failure is not due to violation of local constraints, but to incompatibility across overlaps  $U_i \cap U_j$ .

This phenomenon is naturally expressed in the language of sheaf theory. Let  $\mathcal{F}$  be a presheaf assigning to each  $U_i$  the set of locally admissible configurations. A global configuration corresponds to a section of  $\mathcal{F}$  over  $\Omega$ . The existence of such a section requires that local data satisfy compatibility conditions on overlaps.

**Definition 3.1** (Obstruction). An obstruction is a failure of locally admissible configurations  $\{X_i\}$  on  $\{U_i\}$  to satisfy the cocycle condition on overlaps — that is,  $X_i|_{U_i \cap U_j} = X_j|_{U_i \cap U_j}$  — preventing the existence of a global section.

In this framework, admissibility is a local property, while realization requires the existence of a global section. The admissible set  $\mathcal{M}_{\text{valid}}$  should therefore be understood not as an arbitrary collection, but as the set of configurations that are globally realizable.

The presence of obstruction implies that local correctness is insufficient. A configuration may satisfy all constraints in isolation while remaining globally inconsistent. The role of the selection process is therefore not only to minimize residual tension, but to eliminate configurations that cannot be extended to global coherence.

This perspective clarifies the role of decimation and uncertainty reduction. Both processes act to remove locally admissible but globally incompatible structures. Modes that fail to glue across scales are eliminated, leaving only those configurations that admit a consistent global section.

We emphasize that the bias functional  $\mathcal{B}$  implicitly encodes this requirement. Configurations with unresolved incompatibilities across overlaps exhibit high residual tension, whether

measured geometrically, entropically, or in terms of future work. Minimizing  $\mathcal{B}$  therefore enforces not only local admissibility, but global realizability.

This observation connects the present framework to classical results in sheaf theory and microlocal analysis, where obstructions to global sections are measured by cohomological invariants [13, 14]. In the present setting, these obstructions manifest dynamically as configurations that cannot persist under refinement.

We conclude that realization requires both admissibility and the absence of obstruction. The selection process eliminates configurations that fail to satisfy global compatibility, ensuring that realized configurations correspond to true global sections. With this refined understanding of admissibility, we turn to the structure of the bias functional that governs selection among globally realizable configurations.

## 4 Realization as a Selection Operator

With the understanding that admissibility implies global realizability and not merely local constraint satisfaction, we now formalize the transition from the admissible set to a structured notion of realization. The bias functional  $\mathcal{B}$  introduced in Section 2 allows us to move beyond an undifferentiated admissible set to identify those configurations that the system's dynamics actually select.

Let  $\mathcal{M}_{\text{valid}}$  denote the admissible set, understood now as the set of globally realizable configurations. We define the realization operator

$$\mathcal{R} : \mathcal{M}_{\text{valid}} \rightarrow \mathcal{M}_{\text{valid}}$$

by

$$\mathcal{R}(h) = \arg \min_{h' \in \mathcal{M}_{\text{valid}}} \mathcal{B}(h').$$

In general,  $\mathcal{R}$  may be multivalued if multiple configurations achieve the same minimal value of  $\mathcal{B}$ . In such cases, the operator selects an equivalence class of minimizers. For simplicity, we will assume either uniqueness or that a consistent tie-breaking mechanism exists, so that  $\mathcal{R}$  may be treated as single-valued.

This operator captures the essential distinction between admissibility and realization. A configuration  $h \in \mathcal{M}_{\text{valid}}$  is *possible*, but it is only *realized* if it is selected by  $\mathcal{R}$ . The presence of  $\mathcal{R}$  therefore induces a reduction from the admissible set to a subset of realized configurations, resolving degeneracy.

We now state a fundamental property of this operator.

**Proposition 4.1** (Idempotence of Realization). *For any  $h \in \mathcal{M}_{\text{valid}}$ , the realization operator*

satisfies

$$\mathcal{R}(\mathcal{R}(h)) = \mathcal{R}(h).$$

*Proof.* By definition,  $\mathcal{R}(h)$  is a minimizer of  $\mathcal{B}$  over  $\mathcal{M}_{\text{valid}}$ . Applying  $\mathcal{R}$  again does not change this fact, since the minimizer of  $\mathcal{B}$  remains the same. Therefore,  $\mathcal{R}(\mathcal{R}(h)) = \mathcal{R}(h)$ .  $\square$

Idempotence expresses a form of stability: once a configuration has been realized, further application of the selection process does not alter it. This aligns with the intuition that realization corresponds to a fixed point of the system's dynamics.

**Definition 4.2** (Realized Configuration). A configuration  $h^* \in \mathcal{M}_{\text{valid}}$  is said to be realized if

$$\mathcal{R}(h^*) = h^*.$$

Under this definition, realized configurations are precisely the fixed points of  $\mathcal{R}$ . The selection problem is thus equivalent to identifying fixed points of the realization operator.

This perspective becomes particularly powerful when combined with dynamics. Rather than computing  $\mathcal{R}$  directly, systems often evolve according to a process that iteratively reduces  $\mathcal{B}$ . In such cases, realization emerges as the limit of a dynamical process:

$$h_{t+1} = \Phi(h_t),$$

where  $\Phi$  is an update rule that decreases  $\mathcal{B}$  while preserving admissibility. Under suitable conditions, this process converges to a fixed point of  $\mathcal{R}$ .

The key requirement is that  $\Phi$  be *bias-reducing* and *constraint-preserving*:

$$\mathcal{B}(\Phi(h)) \leq \mathcal{B}(h), \quad \text{and} \quad \Phi(h) \in \mathcal{M}_{\text{valid}} \text{ whenever } h \in \mathcal{M}_{\text{valid}}.$$

In this formulation, the realization operator is not an explicit computation but the asymptotic outcome of a process that eliminates configurations with higher residual tension. The system does not need to evaluate all elements of  $\mathcal{M}_{\text{valid}}$ ; it only needs to follow a trajectory that monotonically reduces  $\mathcal{B}$ .

This shift from global optimization to local dynamics is crucial for connecting the abstract selection principle to concrete systems. In learning algorithms, gradient-based methods provide such dynamics. In physical systems, dissipative processes drive the system toward lower-energy or lower-tension states. In recursive reasoning systems, iterative refinement and pruning play an analogous role.

We will make this connection explicit in subsequent sections by introducing a continuous

formulation of the bias functional and showing how gradient flow implements the realization process.

## 5 The Bias Functional and Its Coordinate Representations

The realization operator introduced in the previous section reduces the admissible set to its stable representatives. Its action is entirely determined by the bias functional  $\mathcal{B}$ , which assigns to each admissible configuration a measure of residual tension. We now make this functional explicit and examine its structure.

We consider configurations  $X$  defined over a domain  $\Omega$ , where  $X$  may represent a scalar field, a vector field, or a composite object depending on the domain of application. The admissible set is defined by constraints

$$X \in \mathcal{M}_{\text{valid}} \subset \mathcal{X},$$

where  $\mathcal{X}$  is the ambient configuration space.

We introduce a bias functional of the form

$$\mathcal{B}(X) = \int_{\Omega} (\alpha \|\nabla^2 X\|^2 + \beta \rho_{\text{unresolved}}(X) + \gamma w_{\text{future}}(X)) d\Omega,$$

where  $\alpha, \beta, \gamma \geq 0$  are weighting coefficients. These should not be interpreted as free parameters, but as conversion factors between coordinate representations of the same underlying ordering. In regimes where one representation dominates, the others may be omitted without changing the minimizing configurations.

The first term,  $\|\nabla^2 X\|^2$ , measures curvature or geometric irregularity. Configurations with large second derivatives exhibit sharp variation and are typically unstable under perturbation. Minimizing this term enforces smoothness and coherence in the geometric sense.

The second term,  $\rho_{\text{unresolved}}(X)$ , measures unresolved entropy or ambiguity. It quantifies the extent to which the configuration fails to uniquely determine underlying microstates or explanations. High values indicate degeneracy or lack of informational closure. Minimizing this term reduces uncertainty and promotes determinacy.

The third term,  $w_{\text{future}}(X)$ , measures the work required to maintain or reconstruct the configuration over time. This can be interpreted as a cost functional associated with future corrections, repairs, or stabilizing interventions. Configurations with high future work are dynamically unstable or economically unsustainable.

At first glance, these terms appear to describe different quantities: geometric smoothness, informational closure, and temporal cost. However, their role within the functional is iden-

tical. Each assigns a penalty to configurations that deviate from stable, coherent structure. The specific form of the penalty depends on the domain, but the ordering it induces on admissible configurations serves the same purpose.

We now state the central organizing principle of this work.

*Principle 5.1 (Bias Equivalence Principle).* The geometric, thermodynamic, and work-based representations of residual tension induce equivalent orderings on the admissible set  $\mathcal{M}_{\text{valid}}$  up to monotone transformation. In particular, they select the same minimizing configurations under admissibility constraints. We do not assume these representations are identical pointwise, only that they are equivalent under the ordering induced on admissible configurations.

The equivalence asserted here is not pointwise identity, nor does it hold over arbitrary configurations. Rather, it holds on the admissible manifold under dynamics that converge to stable fixed points. Each representation measures sensitivity to perturbations within its domain: curvature under geometric variation, entropy under informational refinement, and future work under temporal evolution. When restricted to configurations that admit stable continuation, these measures induce the same ordering up to monotone transformation. The equivalence is therefore a claim about selection invariance, not metaphysical identity. That is, if one configuration is preferred to another under one representation, it is also preferred under the others. These quantities are equivalent in the only sense required: they induce the same ordering over admissible configurations and therefore select the same minimizers.

The justification for this principle rests on a shared invariant. Each term measures distance from a stable fixed point under admissible dynamics. Curvature measures instability under geometric perturbations; unresolved entropy measures instability under informational refinement; future work measures instability under temporal evolution. In each case, higher values correspond to configurations that cannot persist without correction.

What varies across domains is not the law of selection, but the units in which its cost is measured.

This perspective allows us to treat  $\mathcal{B}$  as a single abstract functional with multiple realizations, rather than as a sum of unrelated terms. In specific applications, one representation may dominate or be more natural to compute, but the underlying selection principle remains the same.

Having established the structure of the bias functional, we now turn to its dynamical implementation. In continuous systems, realization corresponds to gradient flow on  $\mathcal{B}$  subject to admissibility constraints.

## 6 Continuous Dynamics: Gradient Flow and Constraint Preservation

Having defined the bias functional  $\mathcal{B}$  and its role in selection, we now examine how realization is implemented dynamically in continuous systems. Rather than computing the minimizer directly, we consider an evolution process that progressively reduces  $\mathcal{B}$  while preserving admissibility.

Let  $X(t)$  denote a time-dependent configuration evolving within the admissible set  $\mathcal{M}_{\text{valid}}$ . We define the dynamics by a gradient flow on  $\mathcal{B}$ :

$$\frac{\partial X}{\partial t} = -\frac{\delta \mathcal{B}}{\delta X},$$

where  $\frac{\delta \mathcal{B}}{\delta X}$  denotes the variational derivative of  $\mathcal{B}$  with respect to  $X$ .

This evolution has two key properties. First, it is *bias-reducing*:

$$\frac{d}{dt} \mathcal{B}(X(t)) = - \int_{\Omega} \left\| \frac{\delta \mathcal{B}}{\delta X} \right\|^2 d\Omega \leq 0.$$

Thus,  $\mathcal{B}$  is non-increasing along trajectories of the flow. Second, under appropriate conditions, the flow converges to a stationary point satisfying

$$\frac{\delta \mathcal{B}}{\delta X} = 0,$$

which corresponds to a minimizer of  $\mathcal{B}$  and therefore a realized configuration.

To make this explicit, we compute the variational derivative of  $\mathcal{B}$ . For the curvature term, we obtain a fourth-order (biharmonic) operator:

$$\frac{\delta}{\delta X} \int_{\Omega} \|\nabla^2 X\|^2 d\Omega = \nabla^4 X,$$

up to boundary conditions. Well-posedness of this evolution depends on boundary conditions; periodic or Navier boundary conditions ensure existence of solutions. The entropy and work terms contribute additional forces that depend on the specific forms of  $\rho_{\text{unresolved}}$  and  $w_{\text{future}}$ . Denoting these contributions collectively by functional derivatives  $F_{\text{entropy}}(X)$  and  $F_{\text{work}}(X)$ , we obtain the evolution equation

$$\frac{\partial X}{\partial t} = - \left( \alpha \nabla^4 X + \beta F_{\text{entropy}}(X) + \gamma F_{\text{work}}(X) \right).$$

This equation describes a dissipative process that smooths geometric irregularities, resolves

informational ambiguity, and reduces future work burden. Each term acts as a force driving the system toward a configuration that is stable under its respective notion of perturbation.

A crucial requirement is that the dynamics preserve admissibility. In many systems, the admissible set is defined by constraints of the form  $C(X) = 0$ . To enforce these during evolution, one may project the dynamics onto the tangent space of  $\mathcal{M}_{\text{valid}}$ :

$$\frac{\partial X}{\partial t} = -\Pi_{\mathcal{M}_{\text{valid}}} \left( \frac{\delta \mathcal{B}}{\delta X} \right),$$

where  $\Pi_{\mathcal{M}_{\text{valid}}}$  denotes projection onto admissible directions.

Under this formulation, the system evolves within the constraint manifold while descending along  $\mathcal{B}$ . Realization emerges as the limit of this process, corresponding to a configuration that is both admissible and stationary under the flow.

The structure of this equation reveals a deep connection between selection and dissipation. The system does not explore the admissible space arbitrarily; it follows trajectories that eliminate high-tension configurations. Modes that contribute strongly to  $\mathcal{B}$  decay rapidly, while those compatible with low residual tension persist. This observation connects naturally to decimation methods in dynamical systems, which we examine in the next section.

## 7 Decimation and the Elimination of Incoherent Modes

The gradient flow formulation reveals that realization proceeds through the systematic suppression of configurations that contribute disproportionately to the bias functional. We now examine this process from a complementary perspective: decimation [2]. Rather than viewing dynamics as continuous descent, decimation treats realization as the progressive elimination of incoherent modes.

Let  $X$  admit a decomposition into modes,

$$X = \sum_k X_k,$$

where the index  $k$  may represent frequency, scale, or another basis appropriate to the domain. Decimation proceeds by identifying and removing those modes whose contribution to  $\mathcal{B}$  exceeds a threshold. Formally, one may define a truncation operator  $\mathcal{D}$  such that

$$\mathcal{D}(X) = \sum_{k \in \mathcal{K}_{\text{stable}}} X_k,$$

where  $\mathcal{K}_{\text{stable}}$  indexes modes whose contribution to  $\mathcal{B}$  is sufficiently small. The sequence

$$X^{(n+1)} = \mathcal{D}(X^{(n)})$$

converges, under suitable conditions, to a configuration that is stable under further decimation.

**Proposition 7.1** (Decimation as Realization). *Under repeated application of a decimation operator that removes modes with high contribution to  $\mathcal{B}$ , the resulting configuration converges to a fixed point of the realization operator  $\mathcal{R}$ .*

*Proof.* Each application of  $\mathcal{D}$  reduces the contribution of high-tension modes, thereby decreasing  $\mathcal{B}$  or leaving it unchanged. The sequence  $\{X^{(n)}\}$  is therefore non-increasing in  $\mathcal{B}$  and bounded below. Under mild compactness conditions, it converges to a limit configuration  $X^*$ . By construction,  $X^*$  contains no modes that can be removed without increasing  $\mathcal{B}$ , and is therefore a minimizer of  $\mathcal{B}$  within the admissible set. Hence  $\mathcal{R}(X^*) = X^*$ .  $\square$

This result shows that decimation and gradient flow are not competing descriptions but different realizations of the same underlying process. Gradient flow continuously attenuates incoherent modes, while decimation removes them discretely. Both implement selection by reducing residual tension.

An important consequence of this perspective is that realization does not require global knowledge of the admissible set. It is sufficient to identify and eliminate locally incoherent components. By focusing on the removal of high-tension modes, the system converges to a coherent configuration without enumerating all possibilities.

## 8 Discrete Realization: CLIO and Uncertainty-Guided Selection

The preceding sections describe realization as a continuous process of bias reduction and mode elimination. We now consider the discrete analogue of this process as implemented in recursive reasoning systems, in particular the Cognitive Loop via In-Situ Optimization (CLIO) framework [1]. In such systems, configurations are not fields over a domain but candidate trajectories or structured hypotheses, and dynamics proceed through iterative refinement rather than continuous evolution.

Let  $\mathcal{T}$  denote a space of candidate trajectories or reasoning paths. Each  $\tau \in \mathcal{T}$  represents a partial or complete construction consistent with local constraints. We define an admissible set

$$\mathcal{T}_{\text{valid}} = \{ \tau \in \mathcal{T} \mid \tau \text{ satisfies all imposed constraints} \}.$$

We introduce an uncertainty functional  $U(\tau)$ , which measures the degree of ambiguity or instability associated with a trajectory. High values of  $U$  indicate that the trajectory admits multiple incompatible continuations or requires substantial correction under further inference. Low values indicate that the trajectory is coherent and constrains its own extension.

The discrete realization process proceeds through three operations:

$$\text{exploration} \rightarrow \text{evaluation} \rightarrow \text{pruning}.$$

Exploration generates candidate trajectories, evaluation assigns uncertainty scores, and pruning eliminates those with high uncertainty.

**Theorem 8.1** (Uncertainty-Guided Realization). *Under iterative application of exploration, evaluation, and pruning, the uncertainty functional  $U(\tau_n)$  is non-increasing, and the sequence  $\{\tau_n\}$  converges to a trajectory  $\tau^*$  that is stable under further refinement.*

*Proof Sketch.* At each iteration, pruning removes trajectories with higher uncertainty, ensuring that  $U(\tau_{n+1}) \leq U(\tau_n)$ . Since  $U$  is bounded below, the sequence converges. The limit trajectory  $\tau^*$  admits no further refinement that reduces uncertainty, and is therefore stable under the update operator  $\Phi$ .  $\square$

The correspondence with the continuous framework is direct. The uncertainty functional  $U$  plays the role of the bias functional  $\mathcal{B}$ . Exploration introduces new modes, while pruning eliminates those that contribute to residual tension. Convergence occurs when no further reduction is possible, corresponding to a fixed point of the realization operator.

This perspective clarifies an important empirical observation. Correct reasoning trajectories exhibit a monotonic reduction in uncertainty, while incorrect ones display persistent oscillation [30]. This is not incidental. Under a constraint-closure dynamic, unresolved inconsistency must manifest as persistent uncertainty, while convergence necessarily produces its monotonic decay. Trajectories with unresolved ambiguity require future correction and therefore correspond to configurations with high residual tension, linking the discrete framework to the work-based coordinate of the bias functional.

**Proposition 8.2** (CLIO as Discrete Gradient Flow). *The iterative process of exploration, uncertainty evaluation, and pruning implements a discrete, locally projected descent on the uncertainty functional  $U$ , approximating gradient flow when the neighborhood of candidate trajectories is sufficiently fine.*

This proposition establishes that recursive reasoning systems and field dynamics are not fundamentally different processes, but distinct implementations of the same underlying mechanism. Both reduce residual tension by eliminating incoherent configurations, and both converge to stable realizations defined by constraint closure.

## 9 Cognitive Heuristics as Approximate Selection Mechanisms

The preceding section described discrete realization as an iterative process of exploration, evaluation, and pruning guided by an uncertainty functional. In idealized systems, this process approximates gradient descent on the bias functional  $\mathcal{B}$ . However, real cognitive systems operate under severe resource constraints: limited memory, bounded computation, and restricted access to the full admissible set.

Under such constraints, exact evaluation of  $\mathcal{B}$  is infeasible. Instead, systems employ heuristics—computational shortcuts that approximate the ordering induced by  $\mathcal{B}$  using locally available information. These heuristics do not explicitly minimize residual tension, but they tend to correlate with it in typical environments.

We therefore interpret cognitive heuristics as approximate selection mechanisms: they provide low-cost proxies for the bias functional that guide pruning and realization in the absence of full evaluation. In this sense, heuristics are not arbitrary biases, but adaptive approximations to the same underlying selection principle.

We examine three such heuristics: recency-based selection, semantic overlap, and compression.

## 10 Recency Bias as Temporal Approximation of Residual Tension

One of the most robust heuristics observed in cognitive systems is recency bias: the tendency to prioritize recently encountered information. Within the present framework, recency bias can be understood as a temporal approximation to residual tension. Configurations that have been recently updated are, by construction, more consistent with the current constraint set. Older configurations may encode outdated assumptions or unresolved inconsistencies that have not yet been corrected.

Formally, let  $\tau(t)$  denote a trajectory constructed over time. A recency-weighted uncertainty functional may be written as

$$U_{\text{recency}}(\tau) = \int w(t') \delta(\tau, t') dt',$$

where  $w(t')$  decays with temporal distance and  $\delta(\tau, t')$  measures local inconsistency at time  $t'$ . This weighting biases selection toward configurations that have been recently verified or updated, approximating the cost of future correction using temporal proximity.

This heuristic is imperfect: recency does not guarantee global coherence. However, in environments where constraints evolve gradually, it provides a computationally efficient proxy

for selecting configurations that remain compatible with current conditions. Recency bias can thus be interpreted not as a distortion of rational inference, but as a resource-constrained approximation to minimizing residual tension under temporally evolving constraints.

## 11 Semantic Overlap as a Proxy for Constraint Compatibility

Another widely observed heuristic is the use of semantic overlap: the tendency to favor configurations that share features, structure, or representation with known or previously validated configurations. Within the present framework, semantic overlap functions as a proxy for compatibility with existing constraints.

Let  $\tau_1, \tau_2 \in \mathcal{T}$  be trajectories, and let  $\text{sim}(\tau_1, \tau_2)$  measure their structural or semantic similarity. A heuristic selection rule may favor trajectories that maximize overlap with a set of previously accepted configurations  $\mathcal{T}^*$ :

$$\tau^* = \underset{\tau \in \mathcal{T}_{\text{valid}}}{\text{arg max}} \sum_{\tau' \in \mathcal{T}^*} \text{sim}(\tau, \tau').$$

This rule approximates the bias functional by assuming that configurations similar to previously realized ones are more likely to be globally coherent. From the perspective of constraint closure, semantic overlap estimates the likelihood that a candidate trajectory will glue consistently with existing structure, reducing the probability of selecting configurations that will later require correction.

As with recency bias, this heuristic is not universally reliable. It may suppress novel but valid configurations, leading to local minima. However, it provides an efficient mechanism for approximating the ordering induced by  $\mathcal{B}$  without requiring explicit evaluation.

## 12 Compression and Minimum Description Length as Structural Selection

A third class of heuristic arises from compression: the tendency to prefer representations that admit shorter descriptions. This principle appears in cognitive systems, statistical modeling, and information theory under the name Minimum Description Length (MDL) [9].

Within the present framework, compression can be understood as a proxy for residual tension. Let  $\ell(X)$  denote the description length of a configuration  $X$  under a given encoding scheme. The MDL principle selects configurations that minimize  $\ell(X)$ :

$$X^* = \underset{X \in \mathcal{M}_{\text{valid}}}{\text{arg min}} \ell(X).$$

At first glance, description length appears unrelated to curvature, entropy, or future work. However, all three can be interpreted as measures of compressibility. Configurations with high curvature require detailed local specification and therefore longer descriptions. Smooth configurations admit compact representations. Similarly, configurations with high unresolved entropy require additional information to disambiguate, increasing description length. Configurations that minimize future work avoid encoding unstable dependencies that must later be corrected, reducing the need for additional specification.

Thus, compression provides a unified proxy for residual tension. Minimizing  $\ell(X)$  implicitly penalizes geometric irregularity, informational ambiguity, and dynamic instability. Description length corresponds to the negative log-likelihood under an optimal coding scheme, and minimizing it therefore aligns with the entropy-based coordinate of the bias functional. This reinforces the Bias Equivalence Principle: the three coordinate representations of  $\mathcal{B}$  are independently approached by a natural and independently motivated information-theoretic criterion.

## 13 Failure Modes of Heuristic Selection

The preceding sections interpret cognitive heuristics as resource-constrained approximations to the bias functional  $\mathcal{B}$ . This interpretation implies that heuristic performance depends on alignment between the proxy used by the heuristic and the true ordering induced by  $\mathcal{B}$ . When this alignment fails, the system may select configurations that are locally admissible but globally unstable.

### Recency Misalignment

Recency bias assumes that recently updated configurations are more likely to remain consistent under current constraints. This assumption fails when constraints change discontinuously or when recent updates incorporate transient or noisy information. In such cases, recency prioritizes configurations that are not globally coherent, increasing residual tension despite local consistency.

### Semantic Entrenchment

Semantic overlap favors configurations similar to previously realized ones. While this reduces the risk of obstruction, it can suppress novel configurations that are globally valid but structurally distinct. This leads to entrenchment in local minima, where the system repeatedly selects configurations that are consistent with past structure but suboptimal under the

full bias functional.

### **Compression Artifacts**

Compression-based selection depends on the encoding scheme used to measure description length. If the encoding fails to capture the true generative structure of the domain, configurations that are highly compressible under the chosen scheme may nevertheless exhibit high residual tension. This leads to the selection of configurations that are superficially simple but dynamically unstable.

### **Coupled Failure in Complex Systems**

In many systems, heuristics operate simultaneously. Misalignment in one heuristic can amplify errors in others. For example, recency may prioritize recently compressed representations that reflect encoding artifacts, while semantic overlap reinforces these selections by favoring consistency with prior states. The resulting configuration may be highly stable under the heuristics while remaining far from the true minimizer of  $\mathcal{B}$ .

These failure modes share a common structure: the heuristic proxy diverges from the underlying bias functional. When this occurs, the selection process no longer eliminates configurations with high residual tension, allowing unstable structures to persist. From the perspective of constraint closure, such configurations remain in a metastable state, sustained by the heuristic approximation rather than by genuine minimization of  $\mathcal{B}$ . Over time, unresolved tension accumulates, leading to abrupt correction, reconfiguration, or collapse.

We therefore conclude that heuristics are effective to the extent that they preserve the ordering induced by the bias functional. Their failure modes provide diagnostic signals of misalignment, indicating where the approximation breaks down and where full evaluation of  $\mathcal{B}$  would yield different selections.

## **14 Economic Systems: Future Work as Residual Tension**

We now consider the application of the selection framework to economic systems. Configurations correspond to allocations of resources, production structures, or institutional arrangements. As before, we define an admissible set

$$\mathcal{E}_{\text{valid}} = \{ E \mid E \text{ satisfies all imposed constraints } \},$$

where constraints may include budget balance, resource availability, contractual obligations, or local feasibility conditions.

In economic systems, the most natural coordinate of residual tension is future work. We define  $w_{\text{future}}(E)$  as the *excess* future work required to maintain the configuration relative to the baseline cost imposed by its scale and structure. Equivalently,  $w_{\text{future}}$  may be interpreted as a rate of corrective effort per unit time, isolating instability from unavoidable maintenance. This includes costs associated with correcting imbalances, resolving inconsistencies, and responding to perturbations beyond what the configuration’s scale would necessitate.

A concrete example is provided by attention-based advertising systems. Such systems construct allocations in which user attention is routed through recommendation loops that maximize short-term engagement. These configurations are admissible in the sense that they satisfy all local constraints: user interaction occurs, revenue is generated, and system metrics are optimized. However, they often encode high future work. Sustaining engagement requires continuous recalibration, increasing informational load, and escalating intervention to counteract saturation and user fatigue. The system must expend growing effort to maintain coherence, indicating high residual tension. In contrast, configurations that align user incentives with stable value generation reduce the need for ongoing correction, thereby minimizing  $w_{\text{future}}$ . The selection principle predicts that, absent external support, high-tension attention loops are dynamically unstable and will be replaced or collapse under accumulated corrective burden [35, 34].

We therefore define the economic realization problem as

$$E^* = \arg \min_{E \in \mathcal{E}_{\text{valid}}} w_{\text{future}}(E).$$

This formulation aligns directly with the general framework. The functional  $w_{\text{future}}$  plays the role of the bias functional  $\mathcal{B}$ , and realization corresponds to minimizing residual tension within the admissible set.

This perspective provides a reinterpretation of economic efficiency. Rather than defining efficiency in terms of static optimality conditions, we define it in terms of dynamic stability: a configuration is efficient if it minimizes the work required to maintain itself over time. This definition naturally incorporates considerations of robustness, resilience, and sustainability [23, 24].

The connection to the previous sections becomes evident when we interpret future work as a measure of unresolved tension. Just as curvature measures instability under geometric perturbations and entropy measures instability under informational refinement, future work measures instability under temporal evolution. High values indicate configurations that cannot persist without intervention.

This framework also clarifies the role of speculation and proxy-driven systems more generally. Such systems may produce configurations that are locally admissible but encode high future work. Allocations that depend on continued inflows of attention or capital without corresponding structural stability exhibit low immediate cost but high future correction burden. The selection principle predicts that such configurations will eventually be eliminated, not because they violate immediate constraints, but because they fail to minimize residual tension.

We conclude that economic systems, like learning and physical systems, operate under a common selection principle. Admissibility defines the space of possible configurations, while the minimization of residual tension—here expressed as future work—determines which configurations are realized.

## 15 Constraint Closure and the Selection of Reality

We are now in a position to state the unified principle underlying the preceding developments. Across learning systems, physical dynamics, recursive reasoning, and economic structures, we observe the same pattern: a space of admissible configurations, a degeneracy within that space, and a process that selects among them by eliminating those that fail to achieve coherence.

This process can be summarized as *constraint closure*. A configuration is not realized merely because it satisfies local constraints, but because it remains consistent under all admissible extensions, refinements, and perturbations. Realization therefore corresponds to the closure of constraints across scales.

Formally, let  $\mathcal{M}_{\text{valid}}$  denote the admissible set and let  $\mathcal{B}$  be the bias functional. We define the realized set

$$\mathcal{M}_{\text{real}} = \{ X \in \mathcal{M}_{\text{valid}} \mid X = \mathcal{R}(X) \},$$

where  $\mathcal{R}$  is the realization operator induced by  $\mathcal{B}$ .

**Theorem 15.1** (Constraint-Closure Selection). *Let  $\mathcal{M}_{\text{valid}}$  be a set of admissible configurations and let  $\mathcal{B}$  be a bias functional that induces an ordering on  $\mathcal{M}_{\text{valid}}$ . Under dynamics that (i) preserve admissibility and (ii) monotonically reduce  $\mathcal{B}$ , the system converges to a configuration  $X^* \in \mathcal{M}_{\text{real}}$  that minimizes residual tension. This configuration is stable under further refinement, perturbation, and extension.*

*Proof Sketch.* The dynamics define a sequence or trajectory within  $\mathcal{M}_{\text{valid}}$  along which  $\mathcal{B}$  is non-increasing and bounded below. Under standard compactness or coercivity assumptions, this sequence converges to a limit  $X^*$ . At the limit, no admissible perturbation can reduce  $\mathcal{B}$  further, implying that  $X^*$  is a minimizer. Stability under further refinement follows

from the fact that any extension that preserves admissibility would increase  $\mathcal{B}$ , contradicting minimality. Therefore,  $X^*$  is a fixed point of the realization operator and lies in  $\mathcal{M}_{\text{real}}$ .  $\square$

This theorem unifies the various mechanisms examined in earlier sections. Gradient flow implements continuous descent on  $\mathcal{B}$ ; decimation removes modes that contribute to residual tension; recursive reasoning systems eliminate trajectories with high uncertainty; economic dynamics suppress configurations with excessive future work. Each is an instance of the same selection process.

The key conceptual shift is from construction to elimination. Traditional accounts describe systems as building solutions from elementary components. In contrast, the present framework describes systems as beginning with a space of possibilities and progressively eliminating those that fail to satisfy global coherence. What remains is not constructed in a stepwise manner, but revealed through the removal of incompatible alternatives.

This perspective resolves the apparent paradox of excess capacity. When many admissible configurations exist, the system does not require additional constraints to identify a unique solution. Instead, it requires a mechanism to eliminate those configurations that are unstable under the dynamics of constraint closure. The uniqueness of the realized configuration arises from the elimination process, not from the specification of additional constraints.

The equivalence of the bias functional's coordinate representations plays a crucial role in this unification. Whether residual tension is measured as curvature, entropy, or future work, the selection process yields the same realized configurations. This invariance reflects the fact that all three measures quantify distance from stable fixed points under their respective dynamics.

We therefore arrive at a single principle governing realization across domains:

Among all configurations that satisfy the imposed constraints, those that minimize residual tension and remain stable under all admissible extensions are realized.

This principle has several implications. First, it suggests that the apparent diversity of selection mechanisms across domains is superficial; they are different expressions of the same underlying process. Second, it provides a framework for analyzing systems in which correctness is insufficient to determine outcome, including modern learning systems and complex economic structures. Third, it offers a criterion for stability that is independent of domain-specific details: a configuration is stable if it minimizes residual tension under admissibility.

Finally, this framework reframes the notion of reality itself. Reality is not the result of constructing a particular configuration from first principles, but the outcome of eliminating all

configurations that fail to achieve constraint closure. It is the residue of a process of decimation, the fixed point of a dynamics that removes inconsistency until only coherent structure remains.

Reality is not built by accumulating what is possible, but by eliminating what cannot coexist until only the coherent remains.

# Appendices

## A Variational Structure of the Bias Functional

We briefly justify the form of the bias functional and its role in selection.

Let  $\mathcal{M}_{\text{valid}}$  be a constraint-defined manifold embedded in a configuration space  $\mathcal{X}$ . We assume  $\mathcal{B} : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$  is coercive on  $\mathcal{M}_{\text{valid}}$ , i.e.,

$$\|X\| \rightarrow \infty \Rightarrow \mathcal{B}(X) \rightarrow \infty,$$

ensuring existence of minimizers.

The functional

$$\mathcal{B}(X) = \int_{\Omega} (\alpha \|\nabla^2 X\|^2 + \beta \rho_{\text{unresolved}}(X) + \gamma w_{\text{future}}(X)) d\Omega$$

is not assumed to be uniquely defined, but to lie within an equivalence class under monotone transformations:

$$\mathcal{B} \sim f(\mathcal{B}), \quad f \text{ strictly increasing.}$$

This equivalence preserves the ordering on  $\mathcal{M}_{\text{valid}}$  and therefore preserves minimizers. All results in the main text depend only on this ordering.

## B Discrete–Continuous Correspondence

We sketch the correspondence between discrete CLIO dynamics and continuous gradient flow.

Let  $\{\tau_n\}$  be a sequence generated by  $\tau_{n+1} = \Phi(\tau_n)$ , with uncertainty functional  $U(\tau) \geq 0$  satisfying  $U(\tau_{n+1}) \leq U(\tau_n)$ .

Define a time interpolation  $t = n\Delta t$ , and assume smoothness in the limit  $\Delta t \rightarrow 0$ . Then

$$\frac{\tau_{n+1} - \tau_n}{\Delta t} \rightarrow -\nabla U(\tau),$$

yielding a gradient flow:

$$\frac{d\tau}{dt} = -\nabla U(\tau).$$

Thus discrete pruning dynamics converge to continuous descent on a functional equivalent to  $\mathcal{B}$ .

## C On Equivalence of Coordinate Representations

We clarify the status of the Bias Equivalence Principle.

Let  $\mathcal{B}_1, \mathcal{B}_2$  be two representations of residual tension. We say they are *selection-equivalent* if

$$\mathcal{B}_1(X) < \mathcal{B}_1(Y) \Rightarrow \mathcal{B}_2(X) < \mathcal{B}_2(Y)$$

for all  $X, Y \in \mathcal{M}_{\text{valid}}$ .

If  $\mathcal{M}_{\text{valid}}$  is connected and  $\mathcal{B}_1, \mathcal{B}_2$  are continuous, this implies the existence of a strictly increasing function  $f$  such that  $\mathcal{B}_2 = f \circ \mathcal{B}_1$  on the image of  $\mathcal{B}_1$ .

The principle does not assert identity of representations, but equivalence of induced orderings, which is sufficient for all selection results.

## References

- [1] Newman Cheng, Gordon Broadbent, and William Chappell. Cognitive Loop via In-Situ Optimization: Self-Adaptive Reasoning for Science. *arXiv preprint arXiv:2508.02789*, 2025. Available at: <https://arxiv.org/abs/2508.02789>.
- [2] Ryan T. Grimm, Alexander J. Staat, and Joel D. Eaves. The Integral Decimation Method for Quantum Dynamics and Statistical Mechanics. *Quantum*, 10:2064, 2026. DOI: <https://doi.org/10.22331/q-2026-04-13-2064>.
- [3] M. Dubova and Steven J. Sloman. Excess Capacity Learning. *Behavioral and Brain Sciences*, 2026. DOI: <https://doi.org/10.1017/S0140525X2610510X>.
- [4] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.
- [5] Preetum Nakkiran et al. Deep double descent: Where bigger models and more data hurt. *International Conference on Learning Representations (ICLR)*, 2020.
- [6] Trevor Hastie, Andrea Montanari, Saharon Rosset, and Ryan J. Tibshirani. Surprises in high-dimensional ridgeless least squares interpolation. *Annals of Statistics*, 50(2):949–986, 2022.
- [7] Vladimir Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- [8] Shun-ichi Amari. Natural gradient works efficiently in learning. *Neural Computation*, 10(2):251–276, 1998.
- [9] Michael I. Jordan, Zoubin Ghahramani, Tommi Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.
- [10] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.
- [11] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2nd edition, 1998.
- [12] John C. Baez and Mike Stay. Physics, topology, logic and computation: A Rosetta Stone. *New Structures for Physics*, Springer, 2011.
- [13] Glen E. Bredon. *Sheaf Theory*. Springer, 2nd edition, 1997.
- [14] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds*. Springer, 1990.

- [15] Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009.
- [16] Lawrence C. Evans. *Partial Differential Equations*. American Mathematical Society, 2nd edition, 2010.
- [17] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191, 1961.
- [18] Charles H. Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12):905–940, 1982.
- [19] Seth Lloyd. Ultimate physical limits to computation. *Nature*, 406:1047–1054, 2000.
- [20] Edwin T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957.
- [21] Stuart A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [22] Howard T. Odum. Self-organization, transformity, and information. *Science*, 242(4882):1132–1139, 1988.
- [23] Nicholas Georgescu-Roegen. *The Entropy Law and the Economic Process*. Harvard University Press, 1971.
- [24] Herman E. Daly. *Beyond Growth: The Economics of Sustainable Development*. Beacon Press, 1996.
- [25] Robert U. Ayres. Eco-thermodynamics: Economics and the second law. *Ecological Economics*, 26(2):189–209, 1998.
- [26] Erik Verlinde. On the origin of gravity and the laws of Newton. *Journal of High Energy Physics*, 2011(4):29, 2011.
- [27] Ted Jacobson. Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters*, 75(7):1260–1263, 1995.
- [28] Daniel Carney, Guy Gur-Ari, and Prashant Saraswat. Emergent gravity and the dark universe. *Classical and Quantum Gravity*, 36(24), 2019.
- [29] David Silver et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- [30] Hunter Lightman et al. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*, 2023.
- [31] Shunyu Yao et al. Tree of Thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023.

- [32] Maciej Besta et al. Graph of Thoughts: Solving elaborate problems with large language models. *arXiv preprint arXiv:2308.09687*, 2024.
- [33] Bernardo A. Huberman and Lada A. Adamic. The nature of markets in the World Wide Web. *Quarterly Journal of Electronic Commerce*, 1(1):5–12, 2001.
- [34] Hal R. Varian. Artificial intelligence, economics, and industrial organization. *NBER Working Paper*, 2019.
- [35] Shoshana Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.
- [36] Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.
- [37] W. Ross Ashby. *An Introduction to Cybernetics*. Chapman & Hall, 1956.