Finite Categorical Entropy Obstruction Calculus I: Stochastic Factorizations and Bridge Reconstruction

Duy Hoang Nguyen
ORCID: 0009-0009-9228-7988

June 2, 2026

Abstract

We develop a finite state-relative *categorical entropy reconstruction-obstruction calculus*. The ambient setting is the nonempty finite stochastic category $\mathbf{FinStoch}_{\neq\varnothing}$, whose objects are finite nonempty sets and whose morphisms are finite Markov kernels. Composition is convolution. The purpose is not to reaxiomatize probability theory, to introduce a new stochastic category, or to claim new entropy identities. The contribution is more precise: we organize standard finite conditional-entropy identities as exact numerical reconstruction certificates for explicitly specified hidden-variable reconstruction problems attached to pointed stochastic morphisms, factorizations, trajectories, and bridges. In this first finite paper, obstruction means a support-relative numerical zero-set criterion for exact reconstruction, not a cohomological, derived, or metric obstruction theory.

For a stochastic morphism $K : X \rightsquigarrow Y$ and an input distribution μ, we define the kernel entropy

$$\mathcal{E}_\mu(K) = H(Y \mid X).$$

We prove that $\mathcal{E}_\mu(K) = 0$ if and only if K admits a visible deterministic realization on the support of μ, and we obtain genuine global deterministic realization in $\mathbf{FinSet}_{\neq\varnothing}$ when μ has full support. For composable stochastic morphisms

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z,$$

we define the entropy defect

$$\Delta_\mu(K, L) = \mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L) - \mathcal{E}_\mu(L \star K),$$

and prove

$$\Delta_\mu(K, L) = H(Y \mid X, Z).$$

Thus the defect is precisely the conditional entropy of the hidden intermediate state. In particular, endpoint-reconstructible stochastic factorization need not be deterministic or invertible.

For finite stochastic categorical trajectories

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n,$$

we define accumulated local conditional uncertainty $\mathcal{S}^{(n)}_{\mu_0}$, endpoint composite entropy $\mathcal{C}^{(n)}_{\mu_0}$, and hidden path defect $\mathcal{D}^{(n)}_{\mu_0}$. The main result is the path entropy balance law

$$\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.$$

We also prove the backward chain decomposition

$$\mathcal{D}^{(n)}_{\mu_0} = \sum_{k=1}^{n-1} \Delta_k^{\mathrm{comp}}.$$

Here Δ_k^{comp} is the two-step defect of the compressed pointed factorization

$$(X_0, \mu_0) \overset{K_{0:k}}{\rightsquigarrow} (X_k, \mu_k) \overset{K_k}{\rightsquigarrow} X_{k+1}, \qquad \mu_k = \mu_0 K_{0:k}.$$

Finally, we introduce categorical bridge entropy

$$\Phi_{t,\mu_0}^{(n)} = H(X_t \mid X_0, X_n),$$

a state-relative invariant of the marked two-block factorization detecting exact failure of intermediate-state reconstruction and quantifying the residual bridge uncertainty. We show by example that bridge entropy is not determined by the bare endpoint composite. For the binary symmetric channel, we compute this invariant explicitly and prove a fixed-parameter long-trajectory reconstruction discontinuity:

$$\Phi_\infty(\theta) = \begin{cases} 0, & \theta = 0, \\ \log 2, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}$$

This is a fixed-parameter long-trajectory reconstruction discontinuity, not a finite-length singularity. It is invisible to the endpoint-only normalized diagnostic $n^{-1}H(X_n \mid X_0)$ on fixed finite state spaces; this is not a criticism of the classical path entropy rate $n^{-1}H(X_1, \ldots, X_n \mid X_0)$.

Keywords. Finite categorical entropy reconstruction calculus; finite numerical reconstruction certificates; state-relative reconstruction; nonempty finite stochastic category; Markov kernels; stochastic morphisms; visible deterministic realization; entropy defect; path entropy balance; bridge entropy; Markov bridges; Doob bridge formula; binary symmetric channel; stochastic categorical dynamics.

2020 Mathematics Subject Classification. Primary 94A17, 60J10; Secondary 18A25, 60J20, 68Q87.

Contents

Notation and Conventions

All finite sets used as objects in this paper are assumed nonempty. We write

$$\mathbf{FinStoch}_{\neq\varnothing}$$

for the category whose objects are finite nonempty sets and whose morphisms are row-stochastic Markov kernels. Equivalently, $\mathbf{FinStoch}_{\neq\varnothing}$ is the nonempty-object finite stochastic category used throughout this paper. The nonempty convention is not cosmetic: it ensures that every object admits probability laws, that stochastic rows have nonempty codomain, and that the extension arguments below can choose a point in the target set. Whenever the ordinary notation **FinStoch** is mentioned for comparison, the categorical work of this paper takes place in $\mathbf{FinStoch}_{\neq\varnothing}$.

For a finite nonempty set X, the notation

$$\mathsf{Prob}(X)$$

denotes the simplex of probability distributions on X. We deliberately avoid the notation $\mathcal{P}(X)$, which is often reserved for the power set; if the power set is ever needed, it will be written as 2^X.

Deterministic functions are written $f : X \to Y$. Stochastic morphisms are written

$$K : X \rightsquigarrow Y.$$

Labeled stochastic trajectories are written

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n.$$

The ordinary arrow $\to$ is reserved for deterministic maps, limit statements, or informal prose where no stochastic morphism is being specified. Composition is written

$$L \star K : X \rightsquigarrow Z.$$

All logarithms are base 2, so entropy is measured in bits. We use the standard convention

$$0 \log 0 := 0.$$

For $\mu \in \mathsf{Prob}(X)$, the transported state along $K : X \rightsquigarrow Y$ is

$$(\mu K)(y) = \sum_{x \in X} \mu(x) K(x, y).$$

A stochastic morphism $K : X \rightsquigarrow Y$ is called *Dirac* if every row $K(x, \cdot)$ is a Dirac probability distribution.

For long trajectories, the dependence on the initial law μ_0 is written explicitly at definitions:

$$\mathcal{S}^{(n)}_{\mu_0}, \qquad \mathcal{C}^{(n)}_{\mu_0}, \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}.$$

When μ_0 is fixed, this subscript may be suppressed.

1 Introduction

1.1 Motivation: Entropy as Obstruction

Classical category theory is built from deterministic morphisms: an arrow $f : X \to Y$ assigns to each input $x \in X$ a definite output $f(x) \in Y$. Stochastic dynamics replaces this deterministic assignment by a transition law. Instead of a single output, an input x determines a probability distribution $K(x, \cdot)$ on possible outputs. The resulting stochastic morphism

$$K : X \rightsquigarrow Y$$

can be represented by a row-stochastic matrix when X and Y are finite.

The guiding question of this paper is:

$$\boxed{\text{When does stochasticity obstruct deterministic categorical reconstruction?}}$$

The answer developed here is that entropy should be read not as a metric distance from stochastic data to deterministic data, but as a support-relative reconstruction certificate through its vanishing criteria and as a numerical score of residual uncertainty. A single stochastic morphism may fail to be realizable by a deterministic morphism on the visible support. A two-step stochastic factorization may fail to be endpoint-reconstructible while retaining its intermediate state as a marked hidden object. A long stochastic trajectory may hide its intermediate path after only the endpoints are observed. Finally, even a single internal time slice of a trajectory may fail to be reconstructible from the endpoints.

The central principle of the finite theory is therefore:

$$\boxed{\begin{array}{c}\text{conditional entropy detects exact state-relative reconstruction failure by vanishing}\\ \text{and quantifies residual hidden-variable uncertainty.}\end{array}}$$

This principle produces a layered reconstruction-obstruction calculus consisting of the following state-relative quantities:

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}.$$

Their roles are summarized in Table 1.

Level	Quantity	Reconstruction failure detected
Single morphism	$\mathcal{E}_\mu(K)$	visible deterministic realization
Two-step factorization	$\Delta_\mu(K, L)$	endpoint reconstruction of the marked intermediate state
Full trajectory	$\mathcal{D}^{(n)}_{\mu_0}$	hidden path reconstruction
Bridge state	$\Phi^{(n)}_{t,\mu_0}$	intermediate-state reconstruction

Table 1: The four finite CEOT layers. Each entropy quantity is an exact state-relative obstruction by its vanishing criterion and an uncertainty score by its numerical value.

The finite setting is chosen deliberately. It avoids measure-theoretic complications and makes all conditional entropies, bridge kernels, and reconstruction criteria completely explicit. This first paper is therefore not an attempt to develop the most general measurable version of the theory. Its purpose is to isolate the finite obstruction mechanism in its cleanest form.

1.2 Relation to $\mathbf{FinStoch}_{\neq\varnothing}$, Markov Categories, and Information Theory

The classical finite stochastic category **FinStoch**, and in this paper its nonempty-object version $\mathbf{FinStoch}_{\neq\varnothing}$, is a standard setting for finite stochastic matrices; finite stochastic matrices and finite Markov chains are standard objects in probability theory [3, 4, 5]. It is also closely related to the modern framework of Markov categories and categorical probability, where stochastic maps are studied through diagrammatic and axiomatic categorical structures [12, 13, 14]. The information-theoretic ingredients used in this paper are likewise classical: conditional entropy, entropy chain rules, Markov factorization, and finite-state bridge kernels all belong to standard probability and information theory [1, 2, 6].

Accordingly, this paper does *not* claim to introduce a new category of stochastic morphisms, nor does it claim that the entropy identities themselves are new as information-theoretic identities. Rather, the contribution is the systematic organization of these quantities as state-relative numerical obstructions, invariant under relabeling isomorphism of the corresponding pointed stochastic data, attached to stochastic morphisms, factorizations, trajectories, and bridges.

In short:

$$\boxed{\text{classical stochastic category} + \text{state-relative numerical reconstruction invariants} = \text{finite CEOT.}}$$

1.3 Categorical Content of the Finite Theory

The word categorical is used in a concrete finite sense. The ambient category is $\mathbf{FinStoch}_{\neq\varnothing}$, and the data considered in this paper are pointed finite stochastic diagrams in this category: single arrows, composable pairs, finite strings of arrows, and marked two-block factorizations of a composite. The entropy quantities are not functions merely of endpoint distributions. They are functions of isomorphism classes of such pointed stochastic diagrams.

For example, a pointed two-step stochastic factorization is the diagram

$$(X, \mu, K, Y, L, Z), \qquad X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z.$$

An isomorphism from this datum to

$$(X', \mu', K', Y', L', Z')$$

is a triple of bijections

$$\alpha : X \to X', \qquad \beta : Y \to Y', \qquad \gamma : Z \to Z',$$

such that

$$\mu'(\alpha x) = \mu(x), \qquad K'(\alpha x, \beta y) = K(x, y), \qquad L'(\beta y, \gamma z) = L(y, z).$$

The two-step defect

$$\Delta_\mu(K, L) = H(Y \mid X, Z)$$

is invariant under this isomorphism.

Thus finite CEOT is categorical in the finite diagrammatic sense: it assigns numerical reconstruction invariants to pointed stochastic diagrams in $\mathbf{FinStoch}_{\neq\varnothing}$. This paper does not develop a general Markov-category axiomatics, a functorial obstruction theory for arbitrary source categories, or a universal categorical replacement of stochastic arrows by deterministic arrows.

1.4 Relation to Markov Categories and Finite Bayesian Conditioning

The categorical background of this paper is the concrete finite Markov category $\mathbf{FinStoch}_{\neq\varnothing}$. The Markov-category structure provides the diagrammatic language of states, stochastic morphisms, composition, deterministic maps, copying/deleting operations implicit in finite random variables, and finite stochastic diagrams. The entropy quantities themselves, however, are not defined in an arbitrary Markov category in this paper. They are defined from the finite joint laws induced by concrete stochastic kernels.

The bridge kernel

$$B_t^{a,b}(x) = \frac{K_{0:t}(a,x)K_{t:n}(x,b)}{K_{0:n}(a,b)}$$

is a finite Bayes–Doob conditioning formula on the positive endpoint support. It may be viewed as a concrete finite disintegration of the path law over the endpoint pair, but no abstract disintegration theorem for general Markov categories is asserted here.

Thus the role of Markov-category language is organizational rather than axiomatic: it identifies the finite stochastic diagrams to which the CEOT reconstruction certificates are attached. The present paper remains inside finite stochastic matrices and finite conditional entropy.

Used in CEOT1	Not claimed in CEOT1
Concrete finite stochastic morphisms in $\mathbf{FinStoch}_{\neq\varnothing}$	Entropy in arbitrary Markov categories
Composition of Markov kernels	A general categorical entropy functor
Finite joint laws induced by diagrams	A universal disintegration theorem
Finite Bayes–Doob bridge conditioning	A general Bayesian inversion calculus
Isomorphism invariance of pointed diagrams	Invariance under arbitrary coarse-graining or quotienting

Table 2: Exact boundary between the concrete finite Markov-category language used by CEOT1 and the stronger abstract claims not made in this paper.

1.5 Terminology and Separation of Layers

The information-theoretic ingredients are standard Shannon entropy, conditional entropy, chain-rule identities, and finite Markov factorization. We use the terminology "numerical entropy obstruction invariant" internally for a nonnegative conditional-entropy quantity whose zero set classifies exact solvability of a specified finite reconstruction problem. This terminology is not meant to refer to cohomological obstruction classes, derived obstruction classes, or thermodynamic entropy production.

Throughout the paper we separate three layers. First, the mathematical theorems are finite conditional-entropy identities. Second, the zero-entropy criteria identify exact solvability of explicitly stated reconstruction problems. Third, the word obstruction refers to the interpretation of those zero criteria as numerical obstruction invariants attached to pointed stochastic categorical data. No new entropy identity is claimed at the interpretive layer.

The categorical background is the finite stochastic category and its nonempty-object version, together with the broader language of Markov categories and categorical probability [12, 13, 14]. The bridge formulas are finite-state Markov bridge formulas related to Doob conditioning, reciprocal processes, and Schrodinger bridge theory [6, 7, 8, 9, 10, 11]. Thus the paper should be read as combining standard finite information theory with pointed stochastic categorical data, not as replacing either field.

1.6 Precise Scope and Novelty of the Finite Theory

This paper develops a finite state-relative reconstruction-obstruction calculus, not a new probability theory and not a new category of stochastic morphisms. The entropy identities used below are consequences of standard conditional entropy, the chain rule, and Markov factorization. The novelty is not the chain rule itself, and not the finite criterion $H(U \mid V) = 0$ if and only if U is a function of V almost surely. The novelty is the systematic assignment of these classical facts to four categorically natural pointed stochastic data:

single morphisms, two-step factorizations, finite trajectories, marked bridges.

The resulting quantities are support-relative, factorization-sensitive, trajectory-sensitive, and marked-cut-sensitive. In particular, several of them are not determined by endpoint composites alone.

Thus finite CEOT should be read as

$$\boxed{\begin{array}{c}\text{pointed stochastic diagram} + \text{canonical hidden variable}\\ +\text{conditional entropy zero criterion} = \text{finite reconstruction certificate.}\end{array}}$$

No cohomological obstruction classes, derived obstruction theory, metric distance to a deterministic locus, or universal categorical deterministic-replacement theorem are claimed in this finite paper.

Claimed in this paper	Not claimed in this paper
A finite state-relative reconstruction-obstruction calculus for pointed stochastic data	A new category of probability spaces or Markov kernels
Conditional-entropy zero sets classify exact reconstruction problems	New entropy identities beyond standard information theory
Relabeling invariants of finite stochastic factorizations	Cohomological obstruction classes or derived obstruction theory
Bridge entropy as a marked-factorization invariant	A universal categorical deterministic-replacement theorem

Table 3: Scope of the finite CEOT foundation.

The categorical content lies in the fact that the relevant quantities are attached not only to endpoint distributions, but to morphisms, marked intermediate objects, and their factorizations. In particular, bridge entropy depends on the marked two-block factorization of a stochastic composite, not merely on the bare endpoint kernel. This marked-factorization sensitivity is the key reason bridge entropy is treated as a state-relative isomorphism invariant in this paper.

1.7 Terminology Discipline and Safeguards

Four terms used below have restricted finite reconstruction meanings. First, "deterministic realization" means visible deterministic realization relative to the chosen input law, unless full support or global Dirac rows are explicitly assumed. Kernel entropy is state-relative and does not inspect rows outside the visible support of the input law.

Second, "lossless" is not used as a synonym for channel invertibility, thermodynamic reversibility, measure preservation, or deterministic composition. When it appears, it means only endpoint reconstruction of the marked hidden object. For a two-step factorization this is the condition

$$H(Y \mid X, Z) = 0.$$

The preferred term is *endpoint-reconstructible factorization*.

Third, the reconstruction problems are not arbitrary choices of random variables. They arise from finite diagram collapses: one observes the collapsed datum and asks whether the marked datum forgotten by the collapse is recoverable.

Fourth, endpoint-rate invisibility always refers to the endpoint-only quantity

$$n^{-1} H(X_n \mid X_0),$$

not to the classical path entropy rate

$$n^{-1} H(X_1, \ldots, X_n \mid X_0).$$

Finite CEOT does not replace the classical entropy rate; it identifies the part of path entropy hidden by endpoint collapse.

Term	Restricted meaning in this finite paper
obstruction	A nonnegative finite numerical certificate whose zero set is exactly the solvability locus of a specified reconstruction problem.
categorical	Attached to pointed finite stochastic diagrams in $\mathbf{FinStoch}_{\neq\varnothing}$ and invariant under diagram isomorphism.
lossless	Endpoint-reconstructible for the marked hidden variable; not channel invertibility or preservation of all information.
transition	A limiting reconstruction discontinuity; not a finite-length singularity or thermodynamic phase transition.
classification	A finite reconstruction dictionary derived from the zero-entropy lemma; not a classification of categories or Markov kernels.

1.8 State-Relative Obstruction Problems and Relabeling Invariance

The word *obstruction* is used in this paper in a finite numerical sense. The paper does not construct obstruction classes in cohomology, deformation theory, or derived category theory. Instead, each obstruction consists of three items: a finite pointed stochastic datum, an explicitly specified reconstruction problem, and a nonnegative entropy functional whose zero set is exactly the solvability locus of that problem.

Definition 1.1 (Finite numerical obstruction calculus)**.** *A finite numerical obstruction calculus consists of the following data:*

(i) *a class of finite pointed stochastic data* $\mathfrak{D}$*;*

(ii) *for each datum* $D \in \mathfrak{D}$*, a finite reconstruction problem* $\mathcal{R}(D)$*;*

(iii) *a nonnegative numerical invariant* $I(D) \in [0, \infty)$*;*

(iv) *an exact zero criterion*

$$I(D) = 0 \iff \mathcal{R}(D) \text{ is solvable};$$

(v) *invariance of $I(D)$ under relabeling isomorphism of the corresponding pointed data.*

The calculus is called entropy-based *when $I(D)$ is a conditional Shannon entropy of the hidden variable specified by $\mathcal{R}(D)$.*

Remark 1.2 (Meaning of obstruction theory in this paper)**.** *The phrase "obstruction theory" is used only in the restricted finite numerical sense of Definition 1.1. It does not refer to cohomological obstruction classes, deformation-obstruction complexes, derived lifting problems, or a metric distance to a deterministic subspace. The obstruction is exact because vanishing of the numerical invariant is equivalent to solvability of the specified finite reconstruction problem.*

Remark 1.3 (Why the obstruction terminology is not circular)**.** *The obstruction terminology in this paper is not meant to assert the existence of an obstruction class independent of a reconstruction problem. The reconstruction problem is part of the datum. The content of the calculus is that several categorically natural finite stochastic data—single morphisms, two-step factorizations, trajectories, and marked bridges—come with canonical hidden variables whose exact recoverability is detected by a conditional-entropy zero set.*

Thus the framework should not be read as

$$\textit{entropy creates a new obstruction theory from nothing.}$$

It should be read as

$$\begin{gathered}\textit{pointed stochastic datum} + \textit{canonical hidden variable}\\ +\textit{conditional entropy zero set} = \textit{finite reconstruction certificate.}\end{gathered}$$

The nontrivial organizational point is the uniform treatment of these four data types under support-relative stochastic composition and relabeling isomorphism.

Definition 1.4 (Finite CEOT data)**.** *The finite CEOT data used in this paper are:*

(a) *a pointed stochastic morphism*

$$(X, \mu, K, Y), \qquad K : X \rightsquigarrow Y;$$

(b) *a pointed two-step factorization*

$$(X, \mu, K, Y, L, Z), \qquad X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z;$$

(c) *a pointed finite trajectory*

$$(X_0, \mu_0, K_0, X_1, K_1, \dots, K_{n-1}, X_n), \qquad X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n;$$

(d) *a marked bridge datum, namely a pointed finite trajectory of length $n \geq 2$ together with an integer*

$$t \in \{1, \dots, n-1\}.$$

The associated entropy quantities are

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}.$$

Convention 1.5 (Trajectory and bridge lengths)**.** *A finite trajectory may have length $n \geq 1$. A marked bridge datum is only defined for $n \geq 2$, with marked internal time*

$$t \in \{1, \dots, n-1\}.$$

Whenever $\Phi^{(n)}_{t,\mu_0}$ is written, these conditions are assumed. For $n = 1$, the hidden path defect is interpreted as the entropy of the empty hidden path, hence zero, but there is no bridge entropy because no internal state exists.

1.9 Canonical Hidden Variables from Diagram Collapse

Finite CEOT assigns a reconstruction problem only after specifying a diagram collapse. A collapse forgets part of a pointed stochastic diagram while retaining a visible observation. The hidden variable is the finite datum forgotten by that collapse.

Layer	Full stochastic datum	Visible collapse	Hidden variable
Morphism	$X \overset{K}{\rightsquigarrow} Y$	X	Y
Factorization	$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$	(X, Z)	Y
Path	$X_0 \to X_1 \to \cdots \to X_n$	(X_0, X_n)	$(X_1, \ldots, X_{n-1})$
Bridge	$X_0 \to \cdots \to X_t \to \cdots \to X_n$	(X_0, X_n)	X_t

Thus the conditional entropies used in CEOT are not selected arbitrarily: they are the conditional entropies of the canonical hidden variables forgotten by the corresponding diagram collapse.

Definition 1.6 (Finite collapse-reconstruction datum)**.** *A finite collapse-reconstruction datum consists of finite random variables (O, U), where O is the observed variable retained by a diagram collapse and U is the hidden variable forgotten by that collapse. Its reconstruction certificate is*

$$\mathsf{Rec}(U \mid O) := H(U \mid O).$$

The datum is exactly reconstructible if there exists a function r on $\operatorname{supp}(O)$ such that

$$U = r(O) \qquad \textit{almost surely.}$$

Proposition 1.7 (CEOT layers as collapse-reconstruction data)**.** *The four basic CEOT quantities are instances of the same finite collapse-reconstruction certificate:*

$$\mathsf{Rec}(U \mid O) = H(U \mid O).$$

More precisely,

$$\begin{aligned}
\mathcal{E}_\mu(K) &= \mathsf{Rec}(Y \mid X), \\
\Delta_\mu(K, L) &= \mathsf{Rec}(Y \mid X, Z), \\
\mathcal{D}^{(n)}_{\mu_0} &= \mathsf{Rec}(X_{1:n-1} \mid X_0, X_n), \\
\Phi^{(n)}_{t,\mu_0} &= \mathsf{Rec}(X_t \mid X_0, X_n).
\end{aligned}$$

Each hidden variable is the datum forgotten by a corresponding finite diagram collapse.

Proof. The first identity is the definition of kernel entropy under the joint law $\mathbb{P}(x, y) = \mu(x)K(x, y)$. The second is Theorem 3.3. The third is the definition of hidden path defect, and the fourth is the definition of bridge entropy. □

Definition 1.8 (Visible and global deterministic realization)**.** *Let $K : X \rightsquigarrow Y$ and $\mu \in \mathsf{Prob}(X)$.*

(i) *K is* visibly deterministic relative to *μ if $K(x, \cdot)$ is a Dirac law for every $x \in \operatorname{supp}(\mu)$.*

(ii) *K is* globally deterministic *if $K(x, \cdot)$ is a Dirac law for every $x \in X$.*

(iii) *A row $x \in X \setminus \operatorname{supp}(\mu)$ is called* μ-invisible*.*

Definition 1.9 (Visible supports). *For a pointed trajectory*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n$$

with initial law μ_0, the path law is

$$\mathbb{P}(x_0, \dots, x_n) = \mu_0(x_0) \prod_{i=0}^{n-1} K_i(x_i, x_{i+1}).$$

For any coordinate set A, the visible support $\mathrm{Supp}(X_A)$ *is the set of coordinate values with positive probability under this path law. The analogous definition applies to a pointed morphism or a pointed two-step factorization by using the corresponding one-step or two-step joint law.*

Lemma 1.10 (Finite zero-entropy reconstruction lemma). *Let U and V be finite random variables. Then*

$$H(U \mid V) = 0$$

if and only if there exists a function f on the support of V such that

$$U = f(V) \qquad \textit{almost surely.}$$

Equivalently, for every positive-probability value v of V, the conditional law of U given $V = v$ is a Dirac law.

Proof. Write

$$H(U \mid V) = \sum_{v:\, \mathbb{P}(V=v)>0} \mathbb{P}(V = v) H(U \mid V = v).$$

Every term is nonnegative. Hence the sum vanishes if and only if each conditional entropy $H(U \mid V = v)$ vanishes on the support of V. A finite probability law has Shannon entropy zero if and only if it is a Dirac law. Therefore for each such v there is a unique value $f(v)$ with $\mathbb{P}(U = f(v) \mid V = v) = 1$, which is exactly $U = f(V)$ almost surely. The converse is immediate. □

All later zero-set statements are applications of Lemma 1.10 to the hidden variable and observed variable of the relevant diagram-collapse datum. The paper therefore proves the finite zero-entropy criterion once and uses it as a common reconstruction backbone.

Definition 1.11 (The four finite reconstruction problems). *The finite CEOT reconstruction problems are the following.*

(i) ***Visible deterministic realization.*** *For a pointed morphism (X, μ, K, Y), the problem $\mathcal{R}_{\det}(X, \mu, K, Y)$ asks whether there exists a function*

$$f : \mathrm{supp}(\mu) \to Y$$

such that

$$K(x, \cdot) = \delta_{f(x)} \qquad \textit{for every } x \in \mathrm{supp}(\mu).$$

(ii) ***Two-step intermediate reconstruction.*** *For a pointed factorization*

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$$

with input law μ*, let* (X, Y, Z) *have joint law*

$$\mathbb{P}(x, y, z) = \mu(x)K(x, y)L(y, z).$$

The problem $\mathcal{R}_{\mathrm{mid}}(X, \mu, K, L, Z)$ *asks whether there exists a function*

$$r : \mathrm{Supp}(X, Z) \to Y$$

such that

$$Y = r(X, Z) \qquad \textit{almost surely.}$$

(iii) ***Hidden path reconstruction.*** *For a pointed trajectory of length* n*, the problem* $\mathcal{R}_{\mathrm{path}}$ *asks whether there exists a function*

$$R : \mathrm{Supp}(X_0, X_n) \to X_1 \times \cdots \times X_{n-1}$$

such that

$$(X_1, \ldots, X_{n-1}) = R(X_0, X_n) \qquad \textit{almost surely.}$$

(iv) ***Bridge-state reconstruction.*** *For a marked time* $t \in \{1, \ldots, n-1\}$*, the problem* $\mathcal{R}_{\mathrm{bridge},t}$ *asks whether there exists a function*

$$r_t : \mathrm{Supp}(X_0, X_n) \to X_t$$

such that

$$X_t = r_t(X_0, X_n) \qquad \textit{almost surely.}$$

Definition 1.12 (Exact numerical obstruction and uncertainty score)**.** *Let* $\mathcal{R}$ *be one of the four reconstruction problems in Definition 1.11. A nonnegative functional* I *attached to the same pointed stochastic datum is an* exact numerical obstruction *for* $\mathcal{R}$ *if*

$$I = 0 \quad \Longleftrightarrow \quad \mathcal{R} \textit{ is solvable.}$$

It is an uncertainty score *if, when positive, it equals the conditional Shannon entropy of the unreconstructed variable after the available observed variables are fixed. This paper uses entropy in both senses. It does not identify entropy with a metric distance from a stochastic kernel to the deterministic locus; any such statement would require an additional metric on stochastic kernels.*

Remark 1.13 (Exact obstruction is not metric distance)**.** *For example,* $\mathcal{E}_\mu(K) = 0$ *is equivalent to deterministic realization on the visible support. If* $\mathcal{E}_\mu(K) > 0$*, then* $\mathcal{E}_\mu(K)$ *quantifies the residual uncertainty of the output after the input is known. It is not, without further structure, a distance from* K *to the set of deterministic kernels. A possible total-variation distance to deterministic kernels could be introduced separately, but it is not part of the finite CEOT foundation in this paper.*

Definition 1.14 (Relabeling isomorphism of pointed trajectories)**.** *Two pointed finite stochastic trajectories*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n, \qquad X'_0 \overset{K'_0}{\rightsquigarrow} X'_1 \overset{K'_1}{\rightsquigarrow} \cdots \overset{K'_{n-1}}{\rightsquigarrow} X'_n$$

with initial laws μ_0 *and* μ_0' *are* relabeling-isomorphic *if there are bijections*

$$\alpha_i : X_i \to X_i', \qquad 0 \le i \le n,$$

such that

$$\mu_0'(\alpha_0 x) = \mu_0(x)$$

for all $x \in X_0$*, and*

$$K_i'(\alpha_i x, \alpha_{i+1} y) = K_i(x, y)$$

for all $0 \le i < n$*,* $x \in X_i$*, and* $y \in X_{i+1}$*. Pointed stochastic morphisms and pointed two-step factorizations are relabeling-isomorphic by the corresponding special cases* $n = 1$ *and* $n = 2$*.*

Lemma 1.15 (Conditional entropy is invariant under finite relabeling)**.** *Let* (A, B) *and* (A', B') *be finite pairs of random variables. Suppose there are bijections* $\alpha : A \to A'$ *and* $\beta : B \to B'$ *such that*

$$\mathbb{P}(A' = \alpha(a), B' = \beta(b)) = \mathbb{P}(A = a, B = b)$$

for all a, b*. Then*

$$H(A') = H(A), \qquad H(A' \mid B') = H(A \mid B).$$

The same conclusion holds for any finite tuple of variables relabeled by coordinatewise bijections.

Proof. A finite relabeling only permutes the atoms of the joint law. Shannon entropy is a symmetric function of the atom probabilities, and conditional entropy can be written as

$$H(A \mid B) = H(A, B) - H(B).$$

Both joint and marginal atom probabilities are merely permuted by the given bijections. Hence the entropies are unchanged. □

1.10 Dependence on the Input Law

All obstruction quantities in this paper are state-relative. Thus $\mathcal{E}_\mu(K)$ is not an invariant of the bare stochastic morphism K alone, but of the pointed datum (X, μ, K, Y). Similarly, $\Delta_\mu(K, L)$ is attached to the pointed two-step factorization (X, μ, K, Y, L, Z), while $\mathcal{D}^{(n)}_{\mu_0}$ and $\Phi^{(n)}_{t,\mu_0}$ are attached to pointed finite trajectories. The input law determines which rows, endpoint atoms, and hidden states are visible to the obstruction problem.

Proposition 1.16 (Input-law dependence is essential at every CEOT layer)**.** *The finite CEOT quantities are invariants of pointed stochastic data, not of bare kernels or bare factorizations. More precisely, there are fixed finite kernels for which changing only the input law changes the value of each of the following quantities:*

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}.$$

Proof. All logarithms are base 2.

Kernel entropy. Let $X = Y = \{0, 1\}$, and define

$$K(0, \cdot) = \delta_0, \qquad K(1, \cdot) = \left(\frac{1}{2}, \frac{1}{2}\right).$$

For $\mu = \delta_0$, one has $\mathcal{E}_\mu(K) = 0$. For $\nu = \delta_1$, one has $\mathcal{E}_\nu(K) = 1$.

Two-step defect. Let $X = Y = Z = \{0,1\}$. Use the same kernel $K : X \rightsquigarrow Y$, and let $L : Y \rightsquigarrow Z$ be the constant collapse

$$L(0,\cdot) = L(1,\cdot) = \delta_0.$$

For $\mu = \delta_0$, the intermediate variable Y is determined by (X,Z), so

$$\Delta_\mu(K,L) = H(Y \mid X,Z) = 0.$$

For $\nu = \delta_1$, the endpoint Z is constant and $X = 1$ is fixed, while Y is fair. Hence

$$\Delta_\nu(K,L) = H(Y \mid X,Z) = 1.$$

Path defect. Use the same two-step trajectory

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} X_2$$

with $K_0 = K$ and $K_1 = L$. For $n = 2$,

$$\mathcal{D}^{(2)}_{\mu_0} = \Delta_{\mu_0}(K_0,K_1).$$

Thus the preceding calculation gives

$$\mathcal{D}^{(2)}_{\delta_0} = 0, \qquad \mathcal{D}^{(2)}_{\delta_1} = 1.$$

Bridge entropy. Again take the same two-step trajectory and the unique internal marked time $t = 1$. Since

$$\Phi^{(2)}_{1,\mu_0} = H(X_1 \mid X_0,X_2) = \Delta_{\mu_0}(K_0,K_1),$$

we obtain

$$\Phi^{(2)}_{1,\delta_0} = 0, \qquad \Phi^{(2)}_{1,\delta_1} = 1.$$

Therefore changing only the input law can change every finite CEOT obstruction quantity. Hence these quantities are invariants of pointed data, not of bare kernels, bare composites, or unpointed trajectories. □

Proposition 1.17 (Finite CEOT separation principles)**.** *Finite CEOT separates four pieces of information that are not determined by one another in general.*

(i) *Kernel entropy is state-relative: the value of $\mathcal{E}_\mu(K)$ is not, in general, determined by the bare stochastic morphism K without the chosen input law μ.*

(ii) *The two-step defect is factorization-sensitive: it is not, in general, determined by the endpoint composite $L \star K$ alone.*

(iii) *The hidden path defect is trajectory-sensitive: it is not, in general, determined by the endpoint composite $K_{n-1} \star \cdots \star K_0$ alone.*

(iv) *Bridge entropy is marked-cut-sensitive: it is not, in general, determined by the endpoint composite alone, nor by the unmarked trajectory without the chosen internal time t.*

Proof. Item (i) is Proposition 1.16.

For (ii), let $X = Z = \{*\}$, let $Y = \{0, 1\}$, put $\mu(*) = 1$, define $K(*, 0) = K(*, 1) = 1/2$, and let $L(y, *) = 1$. The endpoint composite is the unique deterministic kernel $X \rightsquigarrow Z$, but

$$\Delta_\mu(K, L) = H(Y \mid X, Z) = H(Y) = 1.$$

By contrast, the trivial factorization through $Y' = \{*\}$ has the same endpoint composite and zero defect. Hence the two-step defect is not a function of the endpoint composite alone.

For (iii), regard the preceding two-step factorizations as length-two trajectories. Then

$$\mathcal{D}_\mu^{(2)} = H(X_1 \mid X_0, X_2) = \Delta_\mu(K, L),$$

so the hidden path defect also fails to be determined by the endpoint composite alone.

For (iv), the same length-two example already shows that bridge entropy is not determined by the endpoint composite alone. To see marked-cut sensitivity inside a fixed unmarked trajectory, take

$$X_0 = \{*\}, \qquad X_1 = \{0, 1\}, \qquad X_2 = \{*\}, \qquad X_3 = \{*\},$$

with $K_0(*, 0) = K_0(*, 1) = 1/2$, $K_1(i, *) = 1$, and $K_2(*, *) = 1$. Then the endpoints X_0 and X_3 are deterministic, while

$$H(X_1 \mid X_0, X_3) = 1, \qquad H(X_2 \mid X_0, X_3) = 0.$$

Thus the same unmarked trajectory has different bridge entropies at different marked cuts. □

1.11 Main Results

We now state the main results informally. Precise definitions and proofs appear in the subsequent sections.

Theorem A: zero-entropy visible realization. Let $K : X \rightsquigarrow Y$ be a stochastic morphism and let $\mu \in \mathsf{Prob}(X)$. The kernel entropy $\mathcal{E}_\mu(K)$ vanishes if and only if every row $K(x, \cdot)$ is Dirac for $x \in \operatorname{supp}(\mu)$. Equivalently, K is deterministic on the support of μ. If μ has full support, then

$$\mathcal{E}_\mu(K) = 0 \quad \Longleftrightarrow \quad K = \delta_f$$

for a unique deterministic function $f : X \to Y$. Thus single-morphism entropy is an exact obstruction to visible deterministic realization, with full deterministic realization obtained under full support.

Theorem B: one-step entropy defect. For composable stochastic morphisms

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z,$$

we define

$$\Delta_\mu(K, L) = \mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L) - \mathcal{E}_\mu(L \star K).$$

If $X \to Y \to Z$ is the induced Markov chain, then

$$\boxed{\Delta_\mu(K, L) = H(Y \mid X, Z).}$$

Consequently $\Delta_\mu(K, L) \geq 0$, and $\Delta_\mu(K, L) = 0$ if and only if the intermediate state Y is reconstructible from (X, Z) almost surely. This shows that the two-step entropy defect measures failure of endpoint reconstruction of the marked intermediate state.

A key point is that endpoint-reconstructible stochastic factorization does not require deterministic morphisms. It only requires that the stochastic intermediate state remain recoverable from the endpoints. Thus CEOT separates *determinism* from *endpoint-reconstructible factorization.*

Theorem C: path entropy balance. For a finite stochastic categorical trajectory

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n,$$

with initial law μ_0, define accumulated local conditional uncertainty $\mathcal{S}^{(n)}_{\mu_0}$, endpoint composite entropy $\mathcal{C}^{(n)}_{\mu_0}$, and hidden path defect $\mathcal{D}^{(n)}_{\mu_0}$. The main balance law is

$$\boxed{\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.}$$

Equivalently,

$$\boxed{H(X_n \mid X_0) + H(X_1, \ldots, X_{n-1} \mid X_0, X_n) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k).}$$

This identity says that the accumulated local conditional uncertainty decomposes into visible endpoint uncertainty plus hidden path uncertainty.

Theorem D: backward chain defect decomposition. The hidden path defect admits a backward conditional-chain decomposition into one-step reconstruction defects of accumulated composites:

$$\boxed{\mathcal{D}^{(n)}_{\mu_0} = \sum_{k=1}^{n-1} \Delta_k^{\mathrm{comp}}, \qquad \Delta_k^{\mathrm{comp}} = \mathcal{E}_{\mu_0}(K_{0:k}) + \mathcal{E}_{\mu_k}(K_k) - \mathcal{E}_{\mu_0}(K_{0:k+1}).}$$

Equivalently,

$$\boxed{H(X_1, \ldots, X_{n-1} \mid X_0, X_n) = \sum_{k=1}^{n-1} H(X_k \mid X_0, X_{k+1}).}$$

This formula is not the same as summing adjacent two-step defects. The correct composite one-step defect compares the accumulated composite $K_{0:k}$ with the next morphism K_k. This distinction is essential for long trajectories.

Theorem E: fixed-parameter long-trajectory reconstruction discontinuity. For the binary symmetric channel $K^{(\theta)}$ on $X = \{0, 1\}$, define the fixed-parameter limiting midpoint bridge entropy

$$\Phi_\infty(\theta) = \lim_{n\to\infty} H(X_{\lfloor n/2 \rfloor} \mid X_0, X_n).$$

Then

$$\boxed{\Phi_\infty(\theta) = \begin{cases} 0, & \theta = 0, \\ \log 2, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}}$$

Since logarithms are base 2, the interior value is $\log 2 = 1$ bit. The discontinuity is a fixed-parameter long-trajectory statement, not a finite-n singularity, and not a uniform statement over joint limits $\theta = \theta_n \to 0$ or $\theta_n \to 1$. For every fixed $\theta \in (0, 1)$, however small, midpoint reconstructibility is destroyed in the limit $n \to \infty$, even though the endpoint-only normalized diagnostic is trivial on fixed finite state spaces.

Structural result F: diagram-isomorphism invariance. The four CEOT quantities descend to well-defined numerical functions on isomorphism classes of finite pointed stochastic diagram groupoids. This is stronger than saying that entropy is invariant under renaming: the domain of each invariant is explicitly the corresponding quotient of pointed stochastic morphisms, factorizations, trajectories, or marked bridges by diagram isomorphism.

Structural result G: marked-factorization sensitivity. The two-step defect and bridge entropy do not factor through the bare endpoint composite. There are pointed factorizations with the same endpoint composite but different values of $\Delta_\mu(K, L)$, and already in length two this implies different bridge entropies for the same endpoint composite.

Structural result H: bridge-path residual decomposition. For each marked time t, bridge entropy is a coordinate projection of the full hidden-path obstruction:

$$\mathcal{D}^{(n)}_{\mu_0} = \Phi^{(n)}_{t,\mu_0} + H(X_{I_t} \mid X_0, X_n, X_t), \qquad I_t = \{1, \ldots, n-1\} \setminus \{t\}.$$

Thus equality $\Phi^{(n)}_{t,\mu_0} = \mathcal{D}^{(n)}_{\mu_0}$ holds exactly when the remaining hidden coordinates are reconstructible from (X_0, X_t, X_n).

1.12 Organization of the Paper

The rest of the paper is organized as follows.

Section 2 recalls the nonempty finite stochastic category $\mathbf{FinStoch}_{\neq\varnothing}$, deterministic embeddings, state transport, kernel entropy, and zero-entropy deterministic realization. This section establishes the single-morphism layer of CEOT.

Section 3 introduces one-step entropy defects for composable stochastic morphisms. The main result identifies the defect with $H(Y \mid X, Z)$, giving the first reconstruction-theoretic interpretation of entropy obstruction.

Section 4 develops stochastic categorical trajectories. It proves the path entropy balance law and the backward chain defect decomposition. It also explains why endpoint-only normalized diagnostic is insufficient on fixed finite state spaces.

Section 5 introduces categorical bridge entropy. It proves the bridge reconstruction criterion, derives the finite bridge kernel formula with the correct positive-probability support condition, proves the bridge-path residual decomposition, and establishes that defect and bridge entropy are marked-factorization invariants rather than invariants of the bare endpoint composite.

Section 6 computes the binary symmetric channel example in detail. It derives explicit formulas for endpoint entropy, path defect, and bridge entropy, checks the two-step numerical case, and proves the fixed-parameter long-trajectory reconstruction discontinuity.

Section 7 discusses the finite scope of the theory and outlines future directions, including measurable CEOT, growing state spaces, entropy obstruction invariants for stochastic functors, and applications to hidden-state dynamics.

2 Finite Stochastic Morphisms and Visible Deterministic Realization

This section fixes the finite categorical setting used throughout the paper and develops the first layer of the finite state-relative obstruction calculus. The ambient category is the nonempty finite stochastic category $\mathbf{FinStoch}_{\neq\varnothing}$: objects are finite nonempty sets, and morphisms are Markov kernels. This is the finite stochastic-matrix setting familiar from finite Markov chains and categorical probability [3, 4, 13]. Deterministic maps embed into this category as row-wise Dirac kernels. Kernel entropy then detects precisely the exact failure of such a stochastic morphism to be deterministic on the visible support of the chosen input state, and its value records residual output uncertainty.

The main result of the section is the visible deterministic-realization criterion: for a stochastic morphism $K : X \rightsquigarrow Y$ and an input law $\mu \in \mathsf{Prob}(X)$, the entropy $\mathcal{E}_\mu(K)$ vanishes if and only if every row of K visible from μ is Dirac. In the full-support case this is exactly the statement that K comes from a deterministic function $X \to Y$.

2.1 Finite Probability Simplexes

For a finite nonempty set X, write $\mathsf{Prob}(X)$ for the set of probability distributions on X:

$$\mathsf{Prob}(X) = \left\{ \mu : X \to [0,1] \ : \ \sum_{x\in X} \mu(x) = 1 \right\}.$$

The support of $\mu \in \mathsf{Prob}(X)$ is

$$\operatorname{supp}(\mu) = \{x \in X : \mu(x) > 0\}.$$

Since all sets are finite and nonempty, every probability law has nonempty support.

2.2 Stochastic Morphisms

Definition 2.1 (Finite stochastic morphism)**.** *Let X and Y be finite nonempty sets. A* stochastic morphism

$$K : X \rightsquigarrow Y$$

is a function

$$K : X \times Y \to [0,1]$$

such that

$$\sum_{y\in Y} K(x,y) = 1$$

for every $x \in X$. Equivalently, each row $K(x,\cdot)$ is a probability distribution on Y:

$$K(x,\cdot) \in \mathsf{Prob}(Y).$$

Thus a stochastic morphism is the finite version of a Markov kernel. It assigns to each input state $x \in X$ a distribution of possible outputs in Y. When X and Y are listed as finite sets, K is a row-stochastic matrix whose rows are indexed by X and whose columns are indexed by Y.

2.3 Convolution Composition

Let

$$K : X \rightsquigarrow Y, \qquad L : Y \rightsquigarrow Z$$

be stochastic morphisms. Their composite is defined by convolution:

$$(L \star K)(x,z) = \sum_{y\in Y} K(x,y)L(y,z).$$

The order convention is the categorical one: first apply K, then apply L. Thus

$$L \star K : X \rightsquigarrow Z.$$

Lemma 2.2 (Closure under convolution)**.** *If $K : X \rightsquigarrow Y$ and $L : Y \rightsquigarrow Z$, then $L \star K : X \rightsquigarrow Z$ is a stochastic morphism.*

Proof. For all $x \in X$ and $z \in Z$, each term $K(x, y)L(y, z)$ is nonnegative, so

$$(L \star K)(x, z) \geq 0.$$

Moreover,

$$\sum_{z \in Z}(L \star K)(x, z) = \sum_{z \in Z}\sum_{y \in Y} K(x, y)L(y, z).$$

Since all sums are finite, we may exchange the order of summation:

$$\sum_{z \in Z}(L \star K)(x, z) = \sum_{y \in Y} K(x, y) \sum_{z \in Z} L(y, z).$$

Each row of L sums to one, hence

$$\sum_{z \in Z}(L \star K)(x, z) = \sum_{y \in Y} K(x, y) = 1.$$

Therefore $L \star K$ is row-stochastic. □

2.4 Dirac Identities

For every finite nonempty set X, define

$$\mathrm{id}_X^{\mathsf{P}} : X \rightsquigarrow X$$

by

$$\mathrm{id}_X^{\mathsf{P}}(x, x') = \delta_x(x') = \begin{cases} 1, & x' = x, \\ 0, & x' \neq x. \end{cases}$$

Equivalently,

$$\mathrm{id}_X^{\mathsf{P}}(x, \cdot) = \delta_x.$$

This is the stochastic version of the identity map: it sends each input to the Dirac distribution concentrated at itself.

Proposition 2.3 (The finite stochastic category)**.** *Finite nonempty sets, stochastic morphisms, convolution composition, and Dirac identities form a category, denoted* $\mathbf{FinStoch}_{\neq\varnothing}$.

Proof. We have already shown that convolution of stochastic morphisms is again stochastic. It remains to prove associativity and the identity laws.

Let

$$K : X \rightsquigarrow Y, \qquad L : Y \rightsquigarrow Z, \qquad M : Z \rightsquigarrow W.$$

For $x \in X$ and $w \in W$,

$$(M \star (L \star K))(x, w) = \sum_{z \in Z}(L \star K)(x, z)M(z, w).$$

Expanding $L \star K$ gives

$$(M \star (L \star K))(x, w) = \sum_{z \in Z}\sum_{y \in Y} K(x, y)L(y, z)M(z, w).$$

Since the sums are finite,

$$(M \star (L \star K))(x,w) = \sum_{y \in Y} K(x,y) \sum_{z \in Z} L(y,z) M(z,w).$$

The inner sum is $(M \star L)(y,w)$, so

$$(M \star (L \star K))(x,w) = \sum_{y \in Y} K(x,y)(M \star L)(y,w) = ((M \star L) \star K)(x,w).$$

Thus composition is associative.

For the right identity law,

$$(K \star \mathrm{id}_X^{\mathsf{P}})(x,y) = \sum_{x' \in X} \mathrm{id}_X^{\mathsf{P}}(x,x') K(x',y) = K(x,y).$$

For the left identity law,

$$(\mathrm{id}_Y^{\mathsf{P}} \star K)(x,y) = \sum_{y' \in Y} K(x,y') \, \mathrm{id}_Y^{\mathsf{P}}(y',y) = K(x,y).$$

Therefore the Dirac kernels are identities, and $\mathbf{FinStoch}_{\neq \varnothing}$ is a category. ☐

2.5 Dirac Stochastic Morphisms

Definition 2.4 (Dirac stochastic morphism)**.** *A stochastic morphism $K : X \rightsquigarrow Y$ is called* Dirac *if every row $K(x,\cdot)$ is a Dirac probability distribution.*

Equivalently, K is Dirac if for each $x \in X$ there exists a unique element $f(x) \in Y$ such that

$$K(x,\cdot) = \delta_{f(x)}.$$

This observation identifies row-wise Dirac stochastic morphisms with ordinary deterministic functions.

Lemma 2.5 (Dirac morphisms are deterministic maps)**.** *A stochastic morphism $K : X \rightsquigarrow Y$ is Dirac if and only if there exists a unique function $f : X \to Y$ such that*

$$K = \delta_f,$$

where

$$\delta_f(x,y) = \begin{cases} 1, & y = f(x), \\ 0, & y \neq f(x). \end{cases}$$

Proof. Suppose K is Dirac. For each $x \in X$, the row $K(x,\cdot)$ is a Dirac distribution on Y. Since Y is finite, there is a unique point $f(x) \in Y$ at which this distribution has mass one. This defines a unique function $f : X \to Y$. By construction,

$$K(x,y) = 1 \iff y = f(x),$$

so $K = \delta_f$.

Conversely, if $K = \delta_f$ for some function $f : X \to Y$, then each row $K(x,\cdot) = \delta_{f(x)}$ is Dirac. Hence K is a Dirac stochastic morphism. ☐

2.6 The Deterministic Embedding

Every deterministic function $f : X \to Y$ induces the stochastic morphism $\delta_f : X \rightsquigarrow Y$ defined above. This construction is compatible with composition.

Lemma 2.6 (Compatibility with composition)**.** *If*

$$f : X \to Y, \qquad g : Y \to Z$$

are deterministic functions, then

$$\delta_g \star \delta_f = \delta_{g \circ f}.$$

Proof. For $x \in X$ and $z \in Z$,

$$(\delta_g \star \delta_f)(x, z) = \sum_{y \in Y} \delta_f(x, y)\delta_g(y, z).$$

Only the term $y = f(x)$ contributes, so

$$(\delta_g \star \delta_f)(x, z) = \delta_g(f(x), z).$$

This equals 1 exactly when $z = g(f(x))$, and equals 0 otherwise. Therefore

$$\delta_g \star \delta_f = \delta_{g \circ f}.$$

□

Proposition 2.7 (Faithful deterministic embedding)**.** *The assignment*

$$X \mapsto X, \qquad f \mapsto \delta_f$$

defines a faithful functor

$$\Delta : \mathbf{FinSet}_{\neq \varnothing} \hookrightarrow \mathbf{FinStoch}_{\neq \varnothing}.$$

Its image consists exactly of Dirac stochastic morphisms. The functor is not full whenever there exist non-Dirac stochastic morphisms between the same finite sets.

Proof. The previous lemma proves compatibility with composition, and the identity function $\mathrm{id}_X : X \to X$ is sent to the Dirac identity $\mathrm{id}_X^{\mathsf{P}}$. Hence Δ is a functor.

If $f, g : X \to Y$ and $\delta_f = \delta_g$, then for each $x \in X$ the Dirac distributions $\delta_{f(x)}$ and $\delta_{g(x)}$ are equal. Hence $f(x) = g(x)$ for all x, so $f = g$. Thus Δ is faithful.

By Lemma 2.5, the image of Δ is exactly the class of Dirac stochastic morphisms. The functor is not full in general because $\mathbf{FinStoch}_{\neq \varnothing}(X, Y)$ contains stochastic morphisms whose rows are not Dirac whenever Y has at least two elements. □

This proposition is the categorical baseline for deterministic realization. It identifies ordinary deterministic functions with Dirac stochastic morphisms. No universal deterministic-replacement construction is asserted here. In the rest of the paper, a stochastic morphism admits deterministic realization precisely when it agrees with one of these deterministic arrows, possibly only on the support of a chosen input state.

2.7 State Transport

Entropy of a kernel is state-dependent. A stochastic morphism $K : X \rightsquigarrow Y$ does not carry a canonical input distribution; one must choose $\mu \in \mathsf{Prob}(X)$. The kernel then transports μ to a distribution on Y.

Definition 2.8 (State transport)**.** *Let $\mu \in \mathsf{Prob}(X)$ and let $K : X \rightsquigarrow Y$. Define $\mu K \in \mathsf{Prob}(Y)$ by*

$$(\mu K)(y) = \sum_{x \in X} \mu(x) K(x, y).$$

Lemma 2.9 (State transport)**.** *If $\mu \in \mathsf{Prob}(X)$ and $K : X \rightsquigarrow Y$, then $\mu K \in \mathsf{Prob}(Y)$. Moreover, if*

$$X \sim \mu, \qquad Y \mid X = x \sim K(x, \cdot),$$

then

$$Y \sim \mu K.$$

Proof. Nonnegativity is immediate because $\mu(x) \geq 0$ and $K(x, y) \geq 0$. Also,

$$\sum_{y \in Y} (\mu K)(y) = \sum_{y \in Y} \sum_{x \in X} \mu(x) K(x, y) = \sum_{x \in X} \mu(x) \sum_{y \in Y} K(x, y) = \sum_{x \in X} \mu(x) = 1.$$

Thus $\mu K \in \mathsf{Prob}(Y)$.

If $X \sim \mu$ and $Y \mid X = x \sim K(x, \cdot)$, then by the law of total probability,

$$\mathbb{P}(Y = y) = \sum_{x \in X} \mathbb{P}(X = x) \mathbb{P}(Y = y \mid X = x) = \sum_{x \in X} \mu(x) K(x, y) = (\mu K)(y).$$

Hence $Y \sim \mu K$. $\square$

State transport is essential for compositional entropy formulas. If

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$$

and the input state on X is μ, then the natural input state for L is not μ, but μK. This is why later formulas contain $\mathcal{E}_{\mu K}(L)$ rather than a state-free entropy of L.

2.8 Kernel Entropy

Let $K : X \rightsquigarrow Y$ be a stochastic morphism and let $\mu \in \mathsf{Prob}(X)$. The pair (μ, K) defines a joint law on $X \times Y$ by

$$\mathbb{P}(X = x, Y = y) = \mu(x) K(x, y).$$

The conditional entropy of Y given X is then the average row entropy of K, weighted by μ.

Definition 2.10 (Kernel entropy)**.** *The* kernel entropy *of $K : X \rightsquigarrow Y$ with respect to $\mu \in \mathsf{Prob}(X)$ is*

$$\mathcal{E}_\mu(K) = -\sum_{x \in X} \mu(x) \sum_{y \in Y} K(x, y) \log K(x, y).$$

Equivalently,

$$\mathcal{E}_\mu(K) = \sum_{x \in X} \mu(x) H(K(x, \cdot)).$$

By construction,

$$\boxed{\mathcal{E}_\mu(K) = H(Y \mid X).}$$

Thus $\mathcal{E}_\mu(K)$ measures the conditional uncertainty produced by the stochastic morphism after the input state is known.

Proposition 2.11 (Ground states)**.** *For every finite nonempty set X and every $\mu \in \mathsf{Prob}(X)$,*

$$\mathcal{E}_\mu(\mathrm{id}_X^{\mathsf{P}}) = 0.$$

More generally, if $f : X \to Y$ is deterministic, then

$$\mathcal{E}_\mu(\delta_f) = 0.$$

Proof. Every row of $\mathrm{id}_X^{\mathsf{P}}$ is a Dirac distribution, hence has entropy zero. Therefore the weighted average of row entropies is zero. The same argument applies to δ_f, since $\delta_f(x, \cdot) = \delta_{f(x)}$ for every $x \in X$. □

This proposition is the first obstruction-theoretic signal: deterministic morphisms have zero kernel entropy. The converse is true on the support of the input law, as shown next.

2.9 Zero-Entropy Rows

Lemma 2.12 (Zero-entropy rows)**.** *Let $K : X \rightsquigarrow Y$ and $\mu \in \mathsf{Prob}(X)$. Then*

$$\mathcal{E}_\mu(K) = 0$$

if and only if $K(x, \cdot)$ is a Dirac distribution for every $x \in \mathrm{supp}(\mu)$.

Proof. By definition,

$$\mathcal{E}_\mu(K) = \sum_{x \in X} \mu(x) H(K(x, \cdot)).$$

Each term is nonnegative. Hence the sum is zero if and only if

$$\mu(x) H(K(x, \cdot)) = 0$$

for every $x \in X$. For $x \in \mathrm{supp}(\mu)$, one has $\mu(x) > 0$, so this is equivalent to

$$H(K(x, \cdot)) = 0.$$

A finite probability distribution has Shannon entropy zero if and only if it is Dirac. Therefore $\mathcal{E}_\mu(K) = 0$ if and only if every row visible from μ is Dirac. □

Rows outside $\mathrm{supp}(\mu)$ do not contribute to $\mathcal{E}_\mu(K)$. Thus zero entropy with respect to a non-full-support state detects deterministic behavior only where the input law actually places mass.

2.10 Visible Deterministic Realization

For an input law $\mu \in \mathsf{Prob}(X)$, write

$$S_\mu := \operatorname{supp}(\mu)$$

for the visible part of the source. The visible restriction of a stochastic morphism $K : X \rightsquigarrow Y$ is the stochastic morphism

$$K|_{S_\mu} : S_\mu \rightsquigarrow Y, \qquad K|_{S_\mu}(x, \cdot) := K(x, \cdot).$$

Theorem 2.13 (Visible deterministic realization)**.** *Let $K : X \rightsquigarrow Y$ and $\mu \in \mathsf{Prob}(X)$. Then*

$$\mathcal{E}_\mu(K) = 0$$

if and only if there exists a deterministic function

$$f_\mu : S_\mu \to Y$$

such that

$$K|_{S_\mu} = \delta_{f_\mu}.$$

Equivalently,

$$\boxed{\mathcal{E}_\mu(K) = 0 \iff K \textit{ is deterministic on the visible support } S_\mu.}$$

Proof. By Lemma 2.12, $\mathcal{E}_\mu(K) = 0$ if and only if every row $K(x, \cdot)$, for $x \in S_\mu$, is Dirac. For each such x, let $f_\mu(x)$ be the unique point of Y at which $K(x, \cdot)$ has mass one. This defines a function

$$f_\mu : S_\mu \to Y$$

and gives

$$K|_{S_\mu}(x, \cdot) = \delta_{f_\mu(x)}$$

for all $x \in S_\mu$. Hence $K|_{S_\mu} = \delta_{f_\mu}$.

Conversely, if such a function f_μ exists, then every row on the visible support is Dirac and hence has entropy zero. Rows outside S_μ have zero μ-weight. Therefore $\mathcal{E}_\mu(K) = 0$. □

Lemma 2.14 (Extension from the visible support)**.** *Every function*

$$f_\mu : S_\mu \to Y$$

extends to at least one function

$$\bar{f} : X \to Y.$$

If $S_\mu = X$, the extension is unique; if $X \setminus S_\mu \neq \varnothing$ and Y has more than one element, the extension is not unique.

Proof. Since objects of $\mathbf{FinStoch}_{\neq\varnothing}$ are nonempty, choose a point $y_0 \in Y$. Define

$$\bar{f}(x) = \begin{cases} f_\mu(x), & x \in S_\mu, \\ y_0, & x \notin S_\mu. \end{cases}$$

This gives an extension. If $X \setminus S_\mu \neq \varnothing$, different choices on invisible points give different extensions whenever Y has more than one element. If $S_\mu = X$, there are no invisible points, so the extension is unique. □

Proposition 2.15 (Global deterministic realization criterion)**.** *Let $K : X \rightsquigarrow Y$. There exists a function*

$$f : X \to Y$$

such that

$$K = \delta_f$$

if and only if every row $K(x, \cdot)$, including rows invisible to any specified input law, is Dirac.

Proof. If $K = \delta_f$, then each row is the Dirac mass at $f(x)$. Conversely, if every row is Dirac, define $f(x)$ to be the unique atom of the Dirac row $K(x, \cdot)$. Then $K(x, \cdot) = \delta_{f(x)}$ for all $x \in X$, so $K = \delta_f$. □

Remark 2.16 (Invisible rows)**.** *Outside S_μ, the entropy $\mathcal{E}_\mu(K)$ does not detect the rows of K. Hence visible deterministic realization is the correct general zero-entropy statement. Full deterministic realization in $\mathbf{FinSet}_{\neq\varnothing}$ requires all rows of K, including invisible rows, to be Dirac.*

Proposition 2.17 (Invisible rows are not detected by kernel entropy)**.** *Let $K : X \rightsquigarrow Y$ and $\mu \in \mathsf{Prob}(X)$. The value of*

$$\mathcal{E}_\mu(K)$$

depends only on the rows $K(x, \cdot)$ with $x \in \operatorname{supp}(\mu)$. In particular, changing $K(x, \cdot)$ arbitrarily on $X \setminus \operatorname{supp}(\mu)$ does not change $\mathcal{E}_\mu(K)$. Consequently,

$$\mathcal{E}_\mu(K) = 0$$

does not imply global deterministic realization unless μ has full support or the rows outside $\operatorname{supp}(\mu)$ *are separately assumed to be Dirac.*

Proof. By definition,

$$\mathcal{E}_\mu(K) = \sum_{x \in X} \mu(x) H(K(x, \cdot)).$$

All terms with $\mu(x) = 0$ vanish. Hence the expression depends only on rows with $x \in \operatorname{supp}(\mu)$. The remaining claims follow immediately. □

Corollary 2.18 (Full-support deterministic realization)**.** *If $\mu \in \mathsf{Prob}(X)$ has full support, then*

$$\mathcal{E}_\mu(K) = 0$$

if and only if there exists a unique function

$$f : X \to Y$$

such that

$$K = \delta_f.$$

Equivalently,

$$\boxed{\mathcal{E}_\mu(K) = 0 \iff K = \delta_f \textit{ for a unique deterministic map } f : X \to Y.}$$

Proof. If μ has full support, then $\operatorname{supp}(\mu) = X$. Theorem 2.13 therefore gives a function $f : X \to Y$ such that $K(x, \cdot) = \delta_{f(x)}$ for all $x \in X$. Hence $K = \delta_f$. Uniqueness follows from Lemma 2.5. □

Remark 2.19 (No universal deterministic-replacement claim)**.** *The phrase deterministic realization is used deliberately. The result above does not construct a universal replacement functor and does not assert a categorical universal property. It says only that zero kernel entropy is equivalent to the relevant rows of the stochastic morphism being Dirac. Under full support, this identifies the whole stochastic morphism with a unique deterministic map through the Dirac embedding.*

Theorem 2.13 and Corollary 2.18 give the single-morphism layer of CEOT. Kernel entropy is the exact obstruction to visible deterministic realization; full deterministic realization is recovered under full support, or more generally when all rows are Dirac as in Proposition 2.15. The next section passes from individual stochastic morphisms to two-step stochastic factorizations, where the relevant obstruction is not whether each arrow is deterministic, but whether the intermediate state can be reconstructed from the endpoints.

3 One-Step Entropy Defects

Section 2 identified the first obstruction layer: a stochastic morphism has zero kernel entropy exactly when it is deterministic on the support of the input law. We now pass from a single stochastic morphism to a two-step stochastic factorization

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z.$$

At this level the central question is no longer whether each arrow is deterministic. A genuinely stochastic arrow may still be part of an endpoint-reconstructible factorization if the intermediate state can be reconstructed from the two endpoints. The entropy defect introduced below measures precisely the failure of such reconstruction.

Throughout this section let

$$K : X \rightsquigarrow Y, \qquad L : Y \rightsquigarrow Z,$$

be composable stochastic morphisms, and let $\mu \in \mathsf{Prob}(X)$ be an input law. They determine a finite Markov chain

$$X \longrightarrow Y \longrightarrow Z$$

with joint law

$$\mathbb{P}(X = x, Y = y, Z = z) = \mu(x)K(x, y)L(y, z).$$

The transported state on Y is μK, and the endpoint kernel from X to Z is the composite

$$L \star K : X \rightsquigarrow Z, \qquad (L \star K)(x, z) = \sum_{y \in Y} K(x, y)L(y, z).$$

Thus

$$\mathcal{E}_\mu(K) = H(Y \mid X), \qquad \mathcal{E}_{\mu K}(L) = H(Z \mid Y), \qquad \mathcal{E}_\mu(L \star K) = H(Z \mid X).$$

Definition 3.1 (Entropy defect of a two-step factorization)**.** *The entropy defect of the composable pair (K, L), relative to the input law μ, is*

$$\boxed{\Delta_\mu(K, L) := \mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L) - \mathcal{E}_\mu(L \star K).}$$

Equivalently, under the induced Markov chain $X \to Y \to Z$,

$$\Delta_\mu(K, L) = H(Y \mid X) + H(Z \mid Y) - H(Z \mid X).$$

Remark 3.2 (State dependence)**.** *The defect $\Delta_\mu(K,L)$ is not an invariant of the pair (K,L) alone. It is an invariant of the pair together with the input state μ. This state dependence is unavoidable: both kernel entropy and reconstruction entropy are evaluated only on the part of the stochastic factorization seen by the input law.*

By definition, the defect satisfies the formal balance identity

$$\boxed{\mathcal{E}_\mu(L \star K) + \Delta_\mu(K,L) = \mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L).}$$

This identity should not be read as the main theorem. Its content becomes meaningful only after identifying the defect with a genuine conditional entropy. The theorem below supplies that interpretation.

Theorem 3.3 (Defect as hidden intermediate entropy)**.** *For the finite Markov chain*

$$X \longrightarrow Y \longrightarrow Z$$

induced by μ, K, L, one has

$$\boxed{\Delta_\mu(K,L) = H(Y \mid X, Z).}$$

Proof. Apply the chain rule to $H(Y,Z \mid X)$ in two ways. First,

$$H(Y,Z \mid X) = H(Y \mid X) + H(Z \mid X,Y).$$

Since $X \to Y \to Z$ is Markov, conditioning on Y screens off X from Z, so

$$H(Z \mid X,Y) = H(Z \mid Y).$$

Therefore

$$H(Y,Z \mid X) = H(Y \mid X) + H(Z \mid Y).$$

Second, the chain rule also gives

$$H(Y,Z \mid X) = H(Z \mid X) + H(Y \mid X,Z).$$

Comparing the two expansions yields

$$H(Y \mid X) + H(Z \mid Y) - H(Z \mid X) = H(Y \mid X,Z).$$

Using Definition 3.1, this is exactly

$$\Delta_\mu(K,L) = H(Y \mid X,Z).$$

□

Remark 3.4 (Identity, zero criterion, and interpretation)**.** *The theorem contains the mathematical identity*

$$\Delta_\mu(K,L) = H(Y \mid X,Z).$$

The finite zero-entropy criterion gives the exact solvability condition

$$\Delta_\mu(K,L) = 0 \quad \Longleftrightarrow \quad Y = r(X,Z) \text{ a.s. for some function } r.$$

The obstruction-theoretic reading is attached to this zero criterion: $\Delta_\mu(K,L)$ measures the residual uncertainty of the intermediate state after the endpoints are known.

Corollary 3.5 (Nonnegativity and subadditivity)**.** *For every composable pair $K : X \rightsquigarrow Y$, $L : Y \rightsquigarrow Z$, and every input law $\mu \in \mathsf{Prob}(X)$,*

$$\boxed{\Delta_\mu(K, L) \geq 0.}$$

Consequently,

$$\boxed{\mathcal{E}_\mu(L \star K) \leq \mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L).}$$

Proof. By Theorem 3.3,

$$\Delta_\mu(K, L) = H(Y \mid X, Z),$$

and conditional entropy is nonnegative in the finite setting. Rearranging the definition of $\Delta_\mu(K, L)$ gives the stated subadditivity inequality. ☐

Corollary 3.6 (Endpoint-reconstructible factorization criterion)**.** *The following are equivalent:*

(i) $\Delta_\mu(K, L) = 0$.

(ii) $H(Y \mid X, Z) = 0$.

(iii) *There exists a function*

$$r : \mathrm{Supp}(X, Z) \to Y$$

such that

$$Y = r(X, Z)$$

almost surely under the joint law induced by μ, K, L.

Equivalently, the marked intermediate state Y is reconstructible from the endpoint pair (X, Z) almost surely. The values of r outside the support of the joint endpoint law of (X, Z) are irrelevant and may be chosen arbitrarily.

Proof. The equivalence of (i) and (ii) follows from Theorem 3.3. The equivalence with (iii) is Lemma 1.10 applied to the hidden variable Y and the observed variable (X, Z). Extending the resulting function arbitrarily outside $\mathrm{Supp}(X, Z)$ gives the stated form. ☐

Definition 3.7 (Endpoint-reconstructible and endpoint-hidden two-step factorizations)**.** *The factorization*

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$$

is called endpoint-reconstructible relative to μ *if*

$$\Delta_\mu(K, L) = 0.$$

It is called endpoint-hidden relative to μ *if*

$$\Delta_\mu(K, L) > 0.$$

Equivalently, endpoint-reconstructibility means that the marked intermediate state Y is reconstructible from the endpoint pair (X, Z) almost surely; endpoint-hidden means that no such almost-sure reconstruction is possible.

Remark 3.8 (Restricted meaning of losslessness)**.** *When the word "lossless" is used in this paper, it has only the restricted reconstruction-theoretic meaning of endpoint-reconstructibility:*

$$H(Y \mid X, Z) = 0.$$

It does not mean that the stochastic morphisms K or L are invertible, deterministic, measure-preserving, thermodynamically reversible, or information-lossless as standalone channels. The property belongs to the pointed factorization with its marked intermediate object, not to the bare composite $L \star K$ alone.

Remark 3.9 (Endpoint hiddenness is not thermodynamic dissipation)**.** *The term* endpoint-hidden *has only the reconstruction-theoretic meaning*

$$\Delta_\mu(K, L) = H(Y \mid X, Z) > 0.$$

It does not refer to thermodynamic entropy production, violation of detailed balance, or irreversibility of a Markov semigroup. Throughout this paper, "loss" means loss of exact endpoint reconstruction of the hidden intermediate object.

Proposition 3.10 (Endpoint reconstruction versus invertibility)**.** *Endpoint-reconstructibility of a two-step factorization is strictly weaker than deterministic invertibility of the constituent arrows. There exist finite factorizations for which*

$$\Delta_\mu(K, L) = 0$$

while K is not deterministic. Conversely, the statement $\Delta_\mu(K, L) = 0$ asserts only

$$Y = r(X, Z) \quad a.s.$$

and does not imply invertibility of K, invertibility of L, or invertibility of the endpoint composite $L \star K$.

Proof. The example below gives a factorization with a genuinely stochastic first arrow and zero defect. The converse statement follows directly from Corollary 3.6, whose conclusion is only the almost-sure reconstruction of the marked intermediate state from the endpoint pair. □

Example 3.11 (Stochastic but endpoint-reconstructible factorization)**.** *Let*

$$X = \{*\}, \qquad Y = \{0, 1\}, \qquad Z = \{0, 1\},$$

and let $\mu(*) = 1$. Define $K : X \rightsquigarrow Y$ by sending $*$ to a fair bit:

$$K(*, 0) = K(*, 1) = \frac{1}{2}.$$

Let $L : Y \rightsquigarrow Z$ be the deterministic identity copy,

$$L(y, z) = \begin{cases} 1, & z = y, \\ 0, & z \neq y. \end{cases}$$

Then K is genuinely stochastic, since

$$\mathcal{E}_\mu(K) = H(Y \mid X) = 1 \text{ bits}.$$

However, $Z = Y$ almost surely. Hence the intermediate state is recovered from the endpoint pair by

$$r(*, z) = z,$$

and therefore

$$H(Y \mid X, Z) = 0.$$

By Theorem 3.3,

$$\Delta_\mu(K, L) = 0.$$

Equivalently, one can compute directly:

$$\mathcal{E}_\mu(K) = 1, \qquad \mathcal{E}_{\mu K}(L) = 0, \qquad \mathcal{E}_\mu(L \star K) = 1,$$

so

$$\Delta_\mu(K, L) = 1 + 0 - 1 = 0.$$

Thus a stochastic factorization may be endpoint-reconstructible even when one of its morphisms is random.

Example 3.12 (Erasure creates defect)**.** *Let again*

$$X = \{*\}, \qquad Y = \{0, 1\}, \qquad Z = \{*\},$$

with $\mu(*) = 1$. Let $K : X \rightsquigarrow Y$ send $*$ to a fair bit, and let $L : Y \rightsquigarrow Z$ collapse both points of Y deterministically to $*$. Then

$$\mathcal{E}_\mu(K) = 1, \qquad \mathcal{E}_{\mu K}(L) = 0, \qquad \mathcal{E}_\mu(L \star K) = 0.$$

Therefore

$$\Delta_\mu(K, L) = 1.$$

Equivalently,

$$\Delta_\mu(K, L) = H(Y \mid X, Z) = H(Y) = 1 \text{ bits}.$$

The composite endpoint morphism is deterministic, but the two-step factorization contains one hidden bit of intermediate information. This is the simplest form of entropy obstruction to endpoint reconstruction of a marked intermediate state.

The examples show the distinction between stochasticity and defect. Randomness introduced by a morphism is not automatically endpoint-hidden: it becomes a defect only when the chosen factorization hides the corresponding intermediate state from the endpoints. Section 4 extends this two-step obstruction to arbitrary finite stochastic trajectories.

4 Path Entropy Balance

The previous section attached an entropy defect to a two-step stochastic factorization

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z.$$

In this section we extend the same obstruction principle to finite stochastic trajectories. The main point is that the local conditional uncertainty accumulated along the trajectory splits into two parts: uncertainty visible at the endpoint composite and uncertainty hidden in the intermediate path. This gives the path entropy balance law

$$\mathcal{C}_{\mu_0}^{(n)} + \mathcal{D}_{\mu_0}^{(n)} = \mathcal{S}_{\mu_0}^{(n)}.$$

The section also proves a backward chain decomposition expressing the hidden path defect as a sum of composite one-step defects.

4.1 Stochastic Categorical Trajectories

Definition 4.1 (Finite stochastic categorical trajectory). *Let*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n$$

be a sequence of finite nonempty sets and stochastic morphisms. Given an initial state

$$\mu_0 \in \mathsf{Prob}(X_0),$$

we define the induced path law on

$$X_0 \times X_1 \times \cdots \times X_n$$

by

$$\mathbb{P}(x_0, \ldots, x_n) = \mu_0(x_0) \prod_{k=0}^{n-1} K_k(x_k, x_{k+1}).$$

Equivalently, the random variables

$$X_0, X_1, \ldots, X_n$$

form a finite Markov chain with initial law μ_0 and transition kernels $K_0, \ldots, K_{n-1}$.

For $0 \le a < b \le n$, define the composite kernel

$$K_{a:b} = K_{b-1} \star K_{b-2} \star \cdots \star K_a.$$

Thus

$$K_{a:b} : X_a \rightsquigarrow X_b.$$

We also use the convention

$$K_{a:a} = \mathrm{id}^{\mathsf{P}}_{X_a}.$$

In particular,

$$K_{0:n} : X_0 \rightsquigarrow X_n$$

is the endpoint composite of the whole trajectory.

The transported states are

$$\mu_k = \mu_0 K_{0:k}, \qquad 0 \le k \le n.$$

Thus $\mu_0 = \mu_0 K_{0:0}$, and

$$\mu_{k+1} = \mu_k K_k.$$

Under the path law, μ_k is exactly the distribution of the random variable X_k.

Convention 4.2 (Conditional laws on positive atoms). *In all finite path-law arguments, conditional distributions given a finite random variable are evaluated only on atoms of the conditioning variable with positive probability. Conditional entropy sums are therefore always taken over the support of the conditioning variable. Equalities of conditional laws are understood almost surely on these positive atoms.*

4.2 Local Entropy, Endpoint Entropy, and Hidden Path Defect

The local conditional entropy contribution at step k is

$$e_k = \mathcal{E}_{\mu_k}(K_k).$$

Since $X_{k+1} \mid X_k = x$ has law $K_k(x, \cdot)$, this is equivalently

$$e_k = H(X_{k+1} \mid X_k).$$

Definition 4.3 (Accumulated local conditional uncertainty). *The accumulated local conditional uncertainty of the trajectory is*

$$\boxed{\mathcal{S}_{\mu_0}^{(n)} = \sum_{k=0}^{n-1} \mathcal{E}_{\mu_k}(K_k) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k).}$$

Remark 4.4 (No thermodynamic entropy-production claim). *The term local conditional entropy is purely information-theoretic here. The quantity*

$$H(X_{k+1} \mid X_k)$$

is the conditional uncertainty of the next categorical state after the current state is known. It is not asserted to be thermodynamic entropy production, dissipation, heat, or entropy flux. Any such physical interpretation would require additional structure not present in the finite stochastic categorical data of this paper.

Definition 4.5 (Endpoint composite entropy). *The endpoint composite entropy is*

$$\boxed{\mathcal{C}_{\mu_0}^{(n)} = \mathcal{E}_{\mu_0}(K_{0:n}) = H(X_n \mid X_0).}$$

It measures the uncertainty visible after the whole trajectory is collapsed into the single endpoint morphism

$$K_{0:n} : X_0 \rightsquigarrow X_n.$$

Definition 4.6 (Hidden path defect). *The hidden path defect is*

$$\boxed{\mathcal{D}_{\mu_0}^{(n)} = H(X_1, \dots, X_{n-1} \mid X_0, X_n).}$$

It is the entropy of the intermediate trajectory after both endpoints are known.

The interpretation is direct. The quantity $\mathcal{S}_{\mu_0}^{(n)}$ records all local conditional uncertainty accumulated along the chosen factorization. The quantity $\mathcal{C}_{\mu_0}^{(n)}$ records only what remains visible at the final endpoint. The difference is the residual conditional entropy of the hidden intermediate path.

4.3 The Path Entropy Balance Law

Theorem 4.7 (Path entropy balance). *For every finite stochastic categorical trajectory*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n$$

with initial state μ_0, one has

$$\boxed{\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.}$$

Equivalently,

$$\boxed{H(X_n \mid X_0) + H(X_1, \ldots, X_{n-1} \mid X_0, X_n) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k).}$$

Proof. We expand

$$H(X_1, \ldots, X_n \mid X_0)$$

in two different ways.

First, by the chain rule,

$$H(X_1, \ldots, X_n \mid X_0) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_0, \ldots, X_k).$$

The path law is Markov, so

$$H(X_{k+1} \mid X_0, \ldots, X_k) = H(X_{k+1} \mid X_k).$$

Therefore

$$H(X_1, \ldots, X_n \mid X_0) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k) = \mathcal{S}^{(n)}_{\mu_0}.$$

Second, applying the chain rule by first separating the endpoint X_n,

$$H(X_1, \ldots, X_n \mid X_0) = H(X_n \mid X_0) + H(X_1, \ldots, X_{n-1} \mid X_0, X_n).$$

The two terms on the right are, by definition,

$$\mathcal{C}^{(n)}_{\mu_0} \qquad \text{and} \qquad \mathcal{D}^{(n)}_{\mu_0}.$$

Comparing the two expansions gives

$$\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.$$

□

Remark 4.8 (Identity versus reconstruction interpretation)**.** *The path balance law is the entropy identity*

$$\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.$$

The reconstruction statement is the separate zero criterion

$$\mathcal{D}^{(n)}_{\mu_0} = 0 \quad \Longleftrightarrow \quad (X_1, \ldots, X_{n-1}) = R(X_0, X_n) \text{ a.s. for some } R.$$

The phrase hidden path obstruction refers to this exact zero criterion, not to a new entropy identity.

Corollary 4.9 (Endpoint subadditivity)**.** *For every finite stochastic categorical trajectory,*

$$\boxed{\mathcal{C}^{(n)}_{\mu_0} \leq \mathcal{S}^{(n)}_{\mu_0}.}$$

Equivalently,

$$H(X_n \mid X_0) \leq \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k).$$

Proof. This follows immediately from Theorem 4.7 and the nonnegativity of conditional entropy:

$$\mathcal{D}^{(n)}_{\mu_0} = H(X_1, \dots, X_{n-1} \mid X_0, X_n) \geq 0.$$

□

Corollary 4.10 (Endpoint-reconstructible path criterion)**.** *The following are equivalent:*

(i) $\mathcal{D}^{(n)}_{\mu_0} = 0$.

(ii) *There exists a function*

$$r : X_0 \times X_n \to X_1 \times \cdots \times X_{n-1}$$

such that

$$(X_1, \dots, X_{n-1}) = r(X_0, X_n)$$

almost surely.

The values of r outside the support of the endpoint law of (X_0, X_n) are irrelevant and may be chosen arbitrarily.

Proof. For finite random variables, conditional entropy

$$H(U \mid V)$$

vanishes if and only if U is a function of V almost surely. Apply this with

$$U = (X_1, \dots, X_{n-1}), \qquad V = (X_0, X_n).$$

□

Consequently, positive hidden path defect has a precise reconstruction meaning:

$$\mathcal{D}^{(n)}_{\mu_0} > 0$$

if and only if the hidden path is not reconstructible from (X_0, X_n) almost surely.

4.4 Composite One-Step Defects

The path defect can also be decomposed into a sum of one-step defects, but the correct decomposition is not the adjacent-pair sum. The correct terms compare the accumulated composite up to time k with the next morphism. To avoid ambiguity about which law is used at which object, we introduce an explicit notation.

Definition 4.11 (Composite one-step reconstruction defect)**.** *For $k = 1, \dots, n-1$, let*

$$\mu_k := \mu_0 K_{0:k}.$$

The k-th composite one-step reconstruction defect is

$$\boxed{\Delta_k^{\mathrm{comp}} := \mathcal{E}_{\mu_0}(K_{0:k}) + \mathcal{E}_{\mu_k}(K_k) - \mathcal{E}_{\mu_0}(K_{0:k+1}).}$$

Equivalently, Δ_k^{comp} is the two-step defect of the compressed pointed factorization

$$(X_0, \mu_0) \overset{K_{0:k}}{\rightsquigarrow} (X_k, \mu_k) \overset{K_k}{\rightsquigarrow} X_{k+1}.$$

When no ambiguity is possible, we write this compactly as

$$\Delta_k^{\mathrm{comp}} = \Delta_{\mu_0 \to \mu_k}(K_{0:k}, K_k),$$

where the notation $\mu_0 \to \mu_k$ *records explicitly that the first leg is evaluated at the input law* μ_0 *and the second leg at the transported law* $\mu_k = \mu_0 K_{0:k}$.

Lemma 4.12 (Composite one-step defect)**.** *For* $k = 1, \dots, n-1$,

$$\boxed{\Delta_k^{\mathrm{comp}} = H(X_k \mid X_0, X_{k+1}).}$$

Equivalently,

$$\Delta_{\mu_0 \to \mu_k}(K_{0:k}, K_k) = H(X_k \mid X_0, X_{k+1}).$$

Proof. Consider the compressed two-step Markov chain

$$X_0 \longrightarrow X_k \longrightarrow X_{k+1}.$$

The first transition is the composite kernel

$$K_{0:k} : X_0 \rightsquigarrow X_k,$$

and the second transition is

$$K_k : X_k \rightsquigarrow X_{k+1}.$$

The input law on X_0 is μ_0, and the transported law on X_k is $\mu_k = \mu_0 K_{0:k}$. Applying Theorem 3.3 to this compressed two-step factorization gives

$$H(X_k \mid X_0, X_{k+1}) = \mathcal{E}_{\mu_0}(K_{0:k}) + \mathcal{E}_{\mu_k}(K_k) - \mathcal{E}_{\mu_0}(K_{0:k+1}).$$

By Definition 4.11, the right-hand side is Δ_k^{comp}. □

Lemma 4.13 (Backward screening identity)**.** *For the finite Markov chain*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n,$$

one has

$$\boxed{H(X_1, \dots, X_{n-1} \mid X_0, X_n) = \sum_{k=1}^{n-1} H(X_k \mid X_0, X_{k+1}).}$$

Proof. The backward chain rule gives

$$H(X_1, \dots, X_{n-1} \mid X_0, X_n) = \sum_{k=1}^{n-1} H(X_k \mid X_0, X_{k+1}, \dots, X_n).$$

It remains to prove, for each k, that the future tail can be removed from the conditioning variable.

Fix k and write

$$x_{k+2:n} := (x_{k+2}, \dots, x_n).$$

Consider a positive atom

$$(x_0, x_{k+1}, x_{k+2:n})$$

of the conditioning variable $(X_0, X_{k+1}, \dots, X_n)$. By the Markov factorization of the path law,

$$\begin{aligned}&\mathbb{P}(x_k, x_{k+1}, x_{k+2:n} \mid x_0)\\ &\qquad = K_{0:k}(x_0, x_k)\, K_k(x_k, x_{k+1}) \prod_{j=k+1}^{n-1} K_j(x_j, x_{j+1}).\end{aligned}$$

Therefore, on positive atoms,

$$\begin{aligned}&\mathbb{P}(x_k \mid x_0, x_{k+1}, x_{k+2:n})\\ &\qquad = \frac{K_{0:k}(x_0, x_k) K_k(x_k, x_{k+1})}{\sum_{u \in X_k} K_{0:k}(x_0, u) K_k(u, x_{k+1})}.\end{aligned}$$

The future-tail factor

$$\prod_{j=k+1}^{n-1} K_j(x_j, x_{j+1})$$

cancels from numerator and denominator. The remaining expression is exactly

$$\mathbb{P}(x_k \mid x_0, x_{k+1}).$$

Thus

$$X_k \perp (X_{k+2}, \dots, X_n) \mid (X_0, X_{k+1})$$

on the positive conditioning atoms, and hence

$$H(X_k \mid X_0, X_{k+1}, \dots, X_n) = H(X_k \mid X_0, X_{k+1}).$$

Substituting this equality into the backward chain-rule expansion proves the identity. □

Theorem 4.14 (Backward chain decomposition of the path defect)**.** *For every finite stochastic categorical trajectory,*

$$\boxed{\mathcal{D}_{\mu_0}^{(n)} = \sum_{k=1}^{n-1} \Delta_k^{\text{comp}}.}$$

Equivalently,

$$\boxed{\mathcal{D}_{\mu_0}^{(n)} = \sum_{k=1}^{n-1} \Delta_{\mu_0 \to \mu_k}(K_{0:k}, K_k), \qquad \mu_k = \mu_0 K_{0:k}.}$$

In conditional-entropy form,

$$\boxed{H(X_1, \dots, X_{n-1} \mid X_0, X_n) = \sum_{k=1}^{n-1} H(X_k \mid X_0, X_{k+1}).}$$

Proof. By Lemma 4.12,

$$\Delta_k^{\text{comp}} = H(X_k \mid X_0, X_{k+1})$$

for every $k = 1, \dots, n-1$. Summing over k and applying Lemma 4.13 gives

$$\sum_{k=1}^{n-1} \Delta_k^{\text{comp}} = \sum_{k=1}^{n-1} H(X_k \mid X_0, X_{k+1}) = H(X_1, \dots, X_{n-1} \mid X_0, X_n).$$

The last term is precisely $\mathcal{D}_{\mu_0}^{(n)}$. □

Remark 4.15 (Backward chain decomposition, not scalar telescoping)**.** *The identity in Theorem 4.14 is a backward conditional-chain decomposition. It is not a telescoping cancellation of scalar terms in the elementary algebraic sense. The word "composite" refers to the fact that each summand compares the accumulated composite $K_{0:k}$ with the next morphism K_k; the proof itself uses the backward chain rule and Markov screening.*

4.5 Endpoint Rate, Path Rate, and Defect Rate

The balance law also explains a limitation of endpoint-only normalized diagnostics on fixed finite state spaces. The endpoint entropy $\mathcal{C}_{\mu_0}^{(n)} = H(X_n \mid X_0)$ remains bounded by the logarithm of the state-space size, so its per-step normalized rate vanishes. This statement concerns only an endpoint-only diagnostic; it is not a criticism of the classical path entropy rate.

Definition 4.16 (Endpoint, path, and defect rates)**.** *For a finite trajectory define three normalized quantities:*

$$\mathsf{ER}_{\mu_0}^{(n)} := \frac{1}{n} H(X_n \mid X_0) = \frac{1}{n}\mathcal{C}_{\mu_0}^{(n)},$$

$$\mathsf{PR}_{\mu_0}^{(n)} := \frac{1}{n} H(X_1, \dots, X_n \mid X_0) = \frac{1}{n}\mathcal{S}_{\mu_0}^{(n)},$$

and

$$\mathsf{DR}_{\mu_0}^{(n)} := \frac{1}{n} H(X_1, \dots, X_{n-1} \mid X_0, X_n) = \frac{1}{n}\mathcal{D}_{\mu_0}^{(n)}.$$

By the path entropy balance law,

$$\mathsf{PR}_{\mu_0}^{(n)} = \mathsf{ER}_{\mu_0}^{(n)} + \mathsf{DR}_{\mu_0}^{(n)}.$$

Throughout this paper, claims about endpoint-rate invisibility refer only to $\mathsf{ER}_{\mu_0}^{(n)}$, not to the unnormalized endpoint entropy and not to the classical path entropy rate $\mathsf{PR}_{\mu_0}^{(n)}$.

Proposition 4.17 (Vanishing endpoint-only rate on fixed finite state spaces)**.** *Assume the trajectory is time-homogeneous on a fixed finite state space:*

$$X_k = X, \qquad K_k = K$$

for all k. Then

$$K_{0:n} = K^n.$$

Since X is finite,

$$\mathcal{C}_{\mu_0}^{(n)} = H(X_n \mid X_0) \le \log |X|.$$

Consequently,

$$\boxed{\lim_{n\to\infty} \mathsf{ER}_{\mu_0}^{(n)} = \lim_{n\to\infty} \frac{\mathcal{C}_{\mu_0}^{(n)}}{n} = 0.}$$

Proof. The bound

$$\mathcal{C}_{\mu_0}^{(n)} = H(X_n \mid X_0) \le \log |X|$$

is immediate because X_n takes values in the fixed finite set X. Dividing by n and letting $n \to \infty$ gives the endpoint-rate limit. $\square$

Proposition 4.18 (Endpoint rate is not the classical entropy rate)**.** *On a fixed finite state space,*

$$\mathsf{ER}^{(n)}_{\mu_0} = \frac{1}{n} H(X_n \mid X_0) \to 0.$$

This statement is only a limitation of endpoint-only normalized diagnostics. It is not a statement about the classical path entropy rate

$$\mathsf{PR}^{(n)}_{\mu_0} = \frac{1}{n} H(X_1, \ldots, X_n \mid X_0).$$

The CEOT defect rate satisfies

$$\mathsf{DR}^{(n)}_{\mu_0} = \mathsf{PR}^{(n)}_{\mu_0} - \mathsf{ER}^{(n)}_{\mu_0}.$$

Hence, whenever $\mathsf{PR}^{(n)}_{\mu_0}$ *has a nonzero limit, the same limit is carried by the defect rate because* $\mathsf{ER}^{(n)}_{\mu_0} \to 0$*.*

Proof. The convergence of $\mathsf{ER}^{(n)}_{\mu_0}$ is Proposition 4.17. The identity for $\mathsf{DR}^{(n)}_{\mu_0}$ is the normalized path entropy balance law from the preceding definition. The final assertion follows by taking limits. □

Proposition 4.19 (Path-defect rate under finite irreducibility)**.** *Let* $K : X \rightsquigarrow X$ *be a finite irreducible Markov kernel with stationary law* π*. For the time-homogeneous trajectory with initial law* μ_0*,*

$$X_0 \overset{K}{\rightsquigarrow} X_1 \overset{K}{\rightsquigarrow} \cdots \overset{K}{\rightsquigarrow} X_n,$$

one has

$$\boxed{\lim_{n\to\infty} \frac{1}{n} \mathcal{D}^{(n)}_{\mu_0} = \sum_{x \in X} \pi(x) H(K(x, \cdot)).}$$

No aperiodicity assumption is required for this Cesaro-type rate statement.

Proof. Write

$$h(x) := H(K(x, \cdot)).$$

The accumulated local conditional uncertainty is

$$\mathcal{S}^{(n)}_{\mu_0} = \sum_{k=0}^{n-1} \sum_{x \in X} \mu_k(x) h(x), \qquad \mu_k = \mu_0 K^k.$$

For a finite irreducible Markov chain, the Cesaro averages of the laws converge to the stationary law:

$$\frac{1}{n} \sum_{k=0}^{n-1} \mu_k \longrightarrow \pi.$$

Therefore

$$\frac{1}{n} \mathcal{S}^{(n)}_{\mu_0} \longrightarrow \sum_{x \in X} \pi(x) h(x).$$

By the path entropy balance law,

$$\mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0} - \mathcal{C}^{(n)}_{\mu_0}.$$

Since Proposition 4.17 gives $\mathcal{C}^{(n)}_{\mu_0}/n \to 0$, the asserted limit follows. □

Proposition 4.20 (No new asymptotic entropy rate is claimed)**.** *In the finite irreducible homogeneous case of Proposition 4.19, the limiting defect rate is exactly the classical stationary entropy rate*

$$h_\pi(K) = \sum_{x \in X} \pi(x) H(K(x, \cdot)).$$

Thus CEOT does not introduce a new asymptotic entropy rate in this setting. It gives the same limiting number the endpoint-collapse interpretation

$$\mathcal{D}_{\mu_0}^{(n)} = H(X_1, \ldots, X_{n-1} \mid X_0, X_n),$$

namely the amount of trajectory information hidden after the full path is collapsed to its two endpoints.

Remark 4.21 (Relation to the classical entropy rate)**.** *The equality between the defect-rate limit and the classical entropy-rate limit is intentional. The value of the limit is not the novelty. The CEOT content is that the same rate is identified as the per-step uncertainty of the intermediate categorical trajectory hidden by endpoint collapse.*

Remark 4.22 (Role of aperiodicity)**.** *Aperiodicity is needed only if one wants ordinary marginal convergence $\mu_k \to \pi$. The path-defect rate above is an average over time, so finite irreducibility and Cesaro convergence suffice. If the chain is irreducible and aperiodic, the same limit also follows from ordinary marginal convergence.*

Proposition 4.17 shows that the endpoint-only normalized diagnostic alone cannot classify long stochastic categorical trajectories on a fixed finite state space. Proposition 4.18 separates this diagnostic from the classical path entropy rate. Proposition 4.19 then shows that the path defect rate can still converge to the stationary one-step conditional entropy under finite irreducibility. This motivates the bridge entropy developed in Section 5, which detects reconstruction phenomena invisible to the endpoint-only normalized diagnostic.

5 Categorical Bridge Entropy

The path defect $\mathcal{D}_{\mu_0}^{(n)}$ measures the full residual uncertainty of the hidden trajectory $(X_1, \ldots, X_{n-1})$ after the endpoints (X_0, X_n) are known. This section extracts a more localized obstruction: the uncertainty of a single intermediate state X_t under endpoint conditioning. The resulting quantity is the *categorical bridge entropy.*

Bridge entropy is useful for two reasons. First, it detects reconstruction failure at a specified internal time. Second, unlike endpoint entropy, it is sensitive to the marked intermediate object and to the two-block factorization of the endpoint composite through that object. In particular, two factorizations with the same endpoint composite may have different bridge entropy.

Throughout this section, $n \geq 2$ and every marked bridge time is an integer

$$t \in \{1, \ldots, n-1\}.$$

Thus X_t is an internal state of the finite trajectory, not a continuous-time interpolation parameter.

5.1 Bridge Entropy

Definition 5.1 (Bridge entropy)**.** *Let*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n$$

be a finite stochastic categorical trajectory with initial law μ_0. For $t \in \{1, \dots, n-1\}$, define the time-t bridge entropy *by*

$$\boxed{\Phi^{(n)}_{t,\mu_0} := H(X_t \mid X_0, X_n).}$$

For $n \geq 2$, the midpoint bridge entropy is

$$\boxed{\Phi^{(n)}_{\mu_0} := \Phi^{(n)}_{\lfloor n/2 \rfloor,\mu_0}.}$$

Here $\lfloor n/2 \rfloor \in \{1, \dots, n-1\}$. When μ_0 is fixed, we may write $\Phi^{(n)}_t$ and $\Phi^{(n)}$.

Thus $\Phi^{(n)}_{t,\mu_0}$ detects the exact obstruction to reconstructing the intermediate categorical state X_t from the endpoints by its vanishing, and its numerical value quantifies the residual bridge uncertainty. It is a one-time marginal version of the full path defect $\mathcal{D}^{(n)}_{\mu_0}$.

5.2 Relation to the Path Defect

Proposition 5.2 (Bridge entropy is bounded by path defect)**.** *For every $t \in \{1, \dots, n-1\}$,*

$$\boxed{0 \leq \Phi^{(n)}_{t,\mu_0} \leq \mathcal{D}^{(n)}_{\mu_0}.}$$

In the two-step case $n = 2$, one has

$$\boxed{\Phi^{(2)}_{1,\mu_0} = \mathcal{D}^{(2)}_{\mu_0} = \Delta_{\mu_0}(K_0, K_1).}$$

Proof. The lower bound follows from nonnegativity of conditional entropy. For the upper bound, observe that

$$\Phi^{(n)}_{t,\mu_0} = H(X_t \mid X_0, X_n),$$

whereas

$$\mathcal{D}^{(n)}_{\mu_0} = H(X_1, \dots, X_{n-1} \mid X_0, X_n).$$

Conditioned on (X_0, X_n), the variable X_t is one coordinate of the hidden path. A marginal entropy is bounded above by the corresponding joint entropy, hence

$$H(X_t \mid X_0, X_n) \leq H(X_1, \dots, X_{n-1} \mid X_0, X_n).$$

This proves the inequality.

If $n = 2$ and $t = 1$, the hidden path consists only of X_1. Therefore

$$\Phi^{(2)}_{1,\mu_0} = H(X_1 \mid X_0, X_2) = \mathcal{D}^{(2)}_{\mu_0}.$$

By Theorem 3.3, applied to

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} X_2,$$

this is also $\Delta_{\mu_0}(K_0, K_1)$. $\square$

Proposition 5.3 (Bridge-path defect decomposition at a marked time)**.** *Fix $n \geq 2$ and $t \in \{1, \dots, n-1\}$. Let*

$$I_t := \{1, \dots, n-1\} \setminus \{t\}, \qquad X_{I_t} := (X_i)_{i \in I_t}.$$

Then

$$\boxed{\mathcal{D}^{(n)}_{\mu_0} = \Phi^{(n)}_{t,\mu_0} + H(X_{I_t} \mid X_0, X_n, X_t).}$$

Consequently,

$$0 \leq \Phi^{(n)}_{t,\mu_0} \leq \mathcal{D}^{(n)}_{\mu_0}.$$

Moreover,

$$\Phi^{(n)}_{t,\mu_0} = \mathcal{D}^{(n)}_{\mu_0}$$

if and only if the remaining hidden coordinates X_{I_t} are reconstructible from the triple (X_0, X_t, X_n) almost surely.

Proof. By the chain rule for conditional entropy,

$$\begin{aligned}
\mathcal{D}^{(n)}_{\mu_0} &= H(X_1, \dots, X_{n-1} \mid X_0, X_n) \\
&= H(X_t \mid X_0, X_n) + H(X_{I_t} \mid X_0, X_n, X_t) \\
&= \Phi^{(n)}_{t,\mu_0} + H(X_{I_t} \mid X_0, X_n, X_t).
\end{aligned}$$

The inequality follows from nonnegativity of conditional entropy. Equality holds if and only if

$$H(X_{I_t} \mid X_0, X_n, X_t) = 0,$$

which, by Lemma 1.10, is equivalent to almost-sure reconstructibility of X_{I_t} from (X_0, X_t, X_n). □

Example 5.4 (A zero bridge with positive hidden path defect)**.** *Let*

$$X_0 = X_1 = X_3 = \{*\}, \qquad X_2 = \{0, 1\},$$

with deterministic transitions $X_0 \to X_1$, a fair-bit transition $X_1 \to X_2$, and the deterministic collapse $X_2 \to X_3$. For the marked time $t = 1$,

$$\Phi^{(3)}_1 = H(X_1 \mid X_0, X_3) = 0,$$

but

$$\mathcal{D}^{(3)} = H(X_1, X_2 \mid X_0, X_3) = H(X_2) = 1.$$

Thus exact reconstruction of one marked bridge state does not imply exact reconstruction of the full hidden path.

Corollary 5.5 (Bridge zero does not imply path zero)**.** *For a fixed marked time t, the condition*

$$\Phi^{(n)}_{t,\mu_0} = 0$$

does not imply

$$\mathcal{D}^{(n)}_{\mu_0} = 0$$

in general. It only says that the single marked coordinate X_t is endpoint-reconstructible. The remaining hidden coordinates may still carry positive conditional entropy.

Proof. Example 5.4 gives $\Phi_1^{(3)} = 0$ and $\mathcal{D}^{(3)} = 1$. $\square$

These propositions give one layer relation in the finite CEOT calculus: bridge entropy is a one-time marginal of the full hidden-path obstruction, and the difference is exactly the residual hidden-path entropy after revealing the marked coordinate. The following proposition records the precise meaning of the word "layer" used in this paper. It is a hierarchy of reconstruction tasks and marginalizations, not a total numerical ordering among all obstruction quantities.

Proposition 5.6 (Layer relations in finite CEOT)**.** *The finite CEOT layers satisfy the following relations.*

(i) *For every two-step factorization*
$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z,$$
one has
$$0 \le \Delta_\mu(K, L) \le \mathcal{E}_\mu(K).$$
In particular, $\mathcal{E}_\mu(K) = 0$ *implies* $\Delta_\mu(K, L) = 0$.

(ii) *In the two-step trajectory case,*
$$\Phi^{(2)}_{1,\mu_0} = \mathcal{D}^{(2)}_{\mu_0} = \Delta_{\mu_0}(K_0, K_1).$$

(iii) *For every finite trajectory,*
$$\mathcal{D}^{(n)}_{\mu_0} = \sum_{k=1}^{n-1} \Delta_k^{\mathrm{comp}}.$$
Equivalently,
$$\mathcal{D}^{(n)}_{\mu_0} = \sum_{k=1}^{n-1} \Delta_{\mu_0 \to \mu_k}(K_{0:k}, K_k), \qquad \mu_k = \mu_0 K_{0:k}.$$

(iv) *For every marked time* $t \in \{1, \dots, n-1\}$,
$$0 \le \Phi^{(n)}_{t,\mu_0} \le \mathcal{D}^{(n)}_{\mu_0}.$$

Thus the finite CEOT layers are organized by reconstruction tasks, compression of trajectories, and marginalization of hidden paths, not by a single universal numerical order.

Proof. For (i), Theorem 3.3 gives
$$\Delta_\mu(K, L) = H(Y \mid X, Z),$$
while $\mathcal{E}_\mu(K) = H(Y \mid X)$. Since conditioning cannot increase entropy,
$$H(Y \mid X, Z) \le H(Y \mid X).$$
Nonnegativity is standard. Statement (ii) is the two-step part of Proposition 5.2. Statement (iii) is Theorem 4.14. Statement (iv) is the main inequality in Proposition 5.2. $\square$

5.3 Bridge Reconstruction Criterion

The bridge entropy vanishes exactly when the intermediate state is determined by the endpoints.

Proposition 5.7 (Bridge reconstruction criterion)**.** *For $t \in \{1, \ldots, n-1\}$,*

$$\boxed{\Phi^{(n)}_{t,\mu_0} = 0}$$

if and only if there exists a function

$$r_t : X_0 \times X_n \to X_t$$

such that

$$X_t = r_t(X_0, X_n)$$

almost surely. The values of r_t outside the support of the endpoint law of (X_0, X_n) are irrelevant and may be chosen arbitrarily.

Proof. By Definition 5.1,

$$\Phi^{(n)}_{t,\mu_0} = H(X_t \mid X_0, X_n).$$

For finite random variables, conditional entropy $H(A \mid B)$ is zero if and only if A is a measurable function of B almost surely. Applying this with $A = X_t$ and $B = (X_0, X_n)$ gives the result. □

Remark 5.8 (Bridge theorem versus bridge interpretation)**.** *Bridge entropy is mathematically the conditional entropy*

$$H(X_t \mid X_0, X_n).$$

Its zero set is the exact reconstruction criterion

$$H(X_t \mid X_0, X_n) = 0 \quad \iff \quad X_t = r_t(X_0, X_n) \text{ a.s. for some } r_t.$$

Calling this a bridge obstruction means only that the numerical invariant vanishes exactly when the specified bridge-state reconstruction problem is solvable. Thus bridge entropy is the single-time analogue of the endpoint-reconstructible path criterion in Corollary 4.10: $\mathcal{D}^{(n)}_{\mu_0}$ asks whether the entire hidden trajectory is reconstructible from the endpoints, whereas $\Phi^{(n)}_{t,\mu_0}$ asks the same question only for the state at time t.

5.4 Zero-Entropy Reconstruction Dictionary and Diagram Invariance

The preceding sections now provide the promised formal reconstruction interpretation: each entropy quantity is attached to a specified hidden variable and a specified observed variable. The zero criterion is not a new entropy theorem; it is the finite zero-entropy reconstruction lemma applied uniformly to the four CEOT layers.

Theorem 5.9 (Finite CEOT zero-entropy reconstruction dictionary)**.** *Each CEOT zero criterion is an instance of Lemma 1.10 applied to the hidden and observed variables listed below:*

Layer	*Hidden variable*	*Observed variable*	*Entropy*
morphism	Y	X	$H(Y \mid X)$
factorization	Y	(X, Z)	$H(Y \mid X, Z)$
path	$(X_1, \ldots, X_{n-1})$	(X_0, X_n)	$H(X_{1:n-1} \mid X_0, X_n)$
bridge	X_t	(X_0, X_n)	$H(X_t \mid X_0, X_n)$

Consequently, for the four finite CEOT layers,

$$\begin{aligned}\mathcal{E}_\mu(K) = 0 &\iff \mathcal{R}_{\mathrm{det}}(X, \mu, K, Y) \text{ is solvable,}\\ \Delta_\mu(K, L) = 0 &\iff \mathcal{R}_{\mathrm{mid}}(X, \mu, K, L, Z) \text{ is solvable,}\\ \mathcal{D}^{(n)}_{\mu_0} = 0 &\iff \mathcal{R}_{\mathrm{path}} \text{ is solvable,}\\ \Phi^{(n)}_{t,\mu_0} = 0 &\iff \mathcal{R}_{\mathrm{bridge},t} \text{ is solvable.}\end{aligned}$$

Thus

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}$$

are exact numerical reconstruction obstructions in the sense of Definition 1.12.

Proof. For the morphism layer, Theorem 2.13 identifies $\mathcal{E}_\mu(K)$ with $H(Y \mid X)$ and applies Lemma 1.10 row-wise on the visible support. For the factorization layer, Theorem 3.3 gives $\Delta_\mu(K, L) = H(Y \mid X, Z)$, so Lemma 1.10 gives Corollary 3.6. For the path layer, Corollary 4.10 is the same lemma applied to $U = (X_1, \dots, X_{n-1})$ and $V = (X_0, X_n)$. For the bridge layer, Proposition 5.7 is the same lemma applied to $U = X_t$ and $V = (X_0, X_n)$. These are precisely the four reconstruction problems listed in Definition 1.11. □

Proposition 5.10 (Relabeling invariance of finite CEOT quantities)**.** *Under relabeling isomorphism of the corresponding pointed stochastic data,*

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{S}^{(n)}_{\mu_0}, \qquad \mathcal{C}^{(n)}_{\mu_0}, \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}$$

are unchanged.

Proof. A relabeling isomorphism identifies the relevant finite path law with the primed path law by coordinatewise bijections. By Lemma 1.15, all conditional entropies of corresponding relabeled variables are preserved.

For $\mathcal{E}_\mu(K)$, this gives invariance of $H(Y \mid X)$. For $\mathcal{S}^{(n)}_{\mu_0}$, it gives invariance of each $H(X_{k+1} \mid X_k)$, hence of their sum. For $\mathcal{C}^{(n)}_{\mu_0}$, it gives invariance of $H(X_n \mid X_0)$. For $\mathcal{D}^{(n)}_{\mu_0}$, it gives invariance of

$$H(X_1, \dots, X_{n-1} \mid X_0, X_n).$$

For $\Phi^{(n)}_{t,\mu_0}$, it gives invariance of

$$H(X_t \mid X_0, X_n).$$

Finally, $\Delta_\mu(K, L)$ is invariant directly from its definition as

$$\mathcal{E}_\mu(K) + \mathcal{E}_{\mu K}(L) - \mathcal{E}_\mu(L \star K),$$

since transported laws and composites are relabeled accordingly; equivalently, after Theorem 3.3, this is the invariance of $H(Y \mid X, Z)$. □

Definition 5.11 (Finite pointed stochastic diagram groupoids)**.** *We write*

$$\mathsf{Mor}_*, \qquad \mathsf{Fact}_*, \qquad \mathsf{Traj}_{n,*}, \qquad \mathsf{Bridge}_{n,t,*}$$

for the finite pointed diagram groupoids whose objects are respectively:

(i) *pointed single morphisms (X, μ, K, Y);*

(ii) *pointed two-step factorizations*

$$(X, \mu, K, Y, L, Z);$$

(iii) *pointed n-step trajectories*

$$(X_0, \mu_0, K_0, X_1, \dots, K_{n-1}, X_n);$$

(iv) *pointed marked bridges*

$$(X_0, \mu_0, K_0, \dots, K_{n-1}, X_n; t).$$

Morphisms in these groupoids are coordinatewise bijections that transport the initial law and every stochastic kernel exactly.

Theorem 5.12 (CEOT invariants descend to diagram-isomorphism classes)**.** *The assignments*

$$(X, \mu, K, Y) \longmapsto \mathcal{E}_\mu(K),$$

$$(X, \mu, K, Y, L, Z) \longmapsto \Delta_\mu(K, L),$$

$$(X_0, \mu_0, K_0, \dots, K_{n-1}, X_n) \longmapsto \mathcal{D}^{(n)}_{\mu_0},$$

and

$$(X_0, \mu_0, K_0, \dots, K_{n-1}, X_n; t) \longmapsto \Phi^{(n)}_{t,\mu_0}$$

descend to well-defined maps on the corresponding isomorphism classes:

$$[\mathsf{Mor}_*] \to [0, \infty), \qquad [\mathsf{Fact}_*] \to [0, \infty),$$

$$[\mathsf{Traj}_{n,*}] \to [0, \infty), \qquad [\mathsf{Bridge}_{n,t,*}] \to [0, \infty).$$

Proof. A diagram isomorphism identifies the joint laws of all variables by coordinatewise bijections. Conditional entropy is invariant under bijection of the hidden and observed variables by Lemma 1.15. Therefore the conditional entropies defining the four CEOT quantities are unchanged. This proves that the displayed assignments are constant on isomorphism classes. □

Remark 5.13 (No invariance under coarse-graining is claimed)**.** *The preceding theorem is an isomorphism-invariance statement. CEOT quantities are not claimed to be invariant under arbitrary surjective coarse-grainings, quotient maps, lumpings, or non-invertible stochastic post-processings. Such operations may change the hidden-variable reconstruction problem and hence may change the obstruction value.*

Proposition 5.14 (Finite CEOT is an entropy-based numerical obstruction calculus)**.** *The four layers*

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K, L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}$$

form an entropy-based finite numerical obstruction calculus in the sense of Definition 1.1.

Proof. The pointed data are those of Definition 1.4. The reconstruction problems are $\mathcal{R}_{\mathrm{det}}$, $\mathcal{R}_{\mathrm{mid}}$, $\mathcal{R}_{\mathrm{path}}$, and $\mathcal{R}_{\mathrm{bridge},t}$. The zero criteria are the reconstruction dictionary of Theorem 5.9. Relabeling invariance is Proposition 5.10. Each invariant is a finite conditional Shannon entropy of the hidden variable in the corresponding reconstruction problem. □

5.5 Bridge Kernel Formula

Bridge entropy can be computed explicitly from the factorization of the endpoint composite. For $0 \le a < b \le n$, recall that

$$K_{a:b} = K_{b-1} \star \cdots \star K_a.$$

Thus

$$K_{0:t} : X_0 \rightsquigarrow X_t, \qquad K_{t:n} : X_t \rightsquigarrow X_n, \qquad K_{0:n} = K_{t:n} \star K_{0:t}.$$

Definition 5.15 (Pointwise finite bridge kernel)**.** *Let $t \in \{1, \dots, n-1\}$. For $a \in X_0$ and $b \in X_n$ with*

$$K_{0:n}(a, b) > 0,$$

define the pointwise bridge kernel *from a to b at time t by*

$$\boxed{B_t^{a,b}(x) := \frac{K_{0:t}(a, x)K_{t:n}(x, b)}{K_{0:n}(a, b)}, \qquad x \in X_t.}$$

Lemma 5.16 (Normalization of the pointwise bridge kernel)**.** *For every a, b with $K_{0:n}(a, b) > 0$, the function $B_t^{a,b}$ is a probability distribution on X_t.*

Proof. Nonnegativity is immediate. Moreover,

$$\sum_{x \in X_t} K_{0:t}(a, x)K_{t:n}(x, b) = (K_{t:n} \star K_{0:t})(a, b) = K_{0:n}(a, b).$$

Dividing by $K_{0:n}(a, b) > 0$ gives

$$\sum_{x \in X_t} B_t^{a,b}(x) = 1.$$

□

Proposition 5.17 (Finite bridge kernel under the path law)**.** *Fix the path law induced by the initial distribution μ_0 and by the kernels $K_0, \dots, K_{n-1}$. Let $t \in \{1, \dots, n-1\}$. For $a \in X_0$, $x \in X_t$, and $b \in X_n$, whenever*

$$\mu_0(a)K_{0:n}(a, b) > 0,$$

one has

$$\boxed{\mathbb{P}(X_t = x \mid X_0 = a, X_n = b) = B_t^{a,b}(x) = \frac{K_{0:t}(a, x)K_{t:n}(x, b)}{K_{0:n}(a, b)}.}$$

Proof. The condition $\mu_0(a)K_{0:n}(a, b) > 0$ ensures that the endpoint event $(X_0 = a, X_n = b)$ has positive probability under the path law. In particular, $K_{0:n}(a, b) > 0$, so $B_t^{a,b}$ is defined. By the Markov path factorization,

$$\mathbb{P}(X_t = x, X_n = b \mid X_0 = a) = K_{0:t}(a, x)K_{t:n}(x, b),$$

and

$$\mathbb{P}(X_n = b \mid X_0 = a) = K_{0:n}(a, b).$$

Bayes' formula under the conditional law given $X_0 = a$ gives

$$\mathbb{P}(X_t = x \mid X_0 = a, X_n = b) = \frac{K_{0:t}(a, x)K_{t:n}(x, b)}{K_{0:n}(a, b)} = B_t^{a,b}(x).$$

□

The pointwise kernel $B_t^{a,b}$ is the finite categorical bridge kernel. If the endpoint pair also has positive probability under the initial law μ_0, Proposition 5.17 identifies it with the ordinary conditional law under the path measure. This distinction is important: pairs with $K_{0:n}(a,b) > 0$ but $\mu_0(a) = 0$ have a well-defined pointwise bridge kernel, but they do not define conditioning events of positive probability for the path law induced by μ_0. The formula is the finite-state version of the Markov bridge, equivalently a finite Doob bridge formula, and is closely related to the classical theory of reciprocal processes and Schrödinger bridges [6, 4, 7, 8, 9, 10, 11].

Remark 5.18 (Relation to Markov bridges and reciprocal processes)**.** *The bridge kernel $B_t^{a,b}$ used here is the finite-state endpoint conditioning of a Markov chain. In this paper it is used only as an entropy-reconstruction device: the corresponding obstruction quantity is*

$$H(X_t \mid X_0, X_n).$$

This should be distinguished from the broader Schrödinger bridge problem, where one prescribes endpoint marginal distributions and minimizes relative entropy over path measures, and from the general theory of reciprocal processes. The finite CEOT bridge entropy uses the same endpoint-conditioning geometry, but it does not solve an entropy-minimizing path interpolation problem.

5.6 Weighted Bridge Entropy Formula

The bridge kernel formula gives an explicit expression for $\Phi_{t,\mu_0}^{(n)}$, but the expression must be written only over endpoint atoms of positive probability. The endpoint pair (X_0, X_n) has law

$$\mathbb{P}(X_0 = a, X_n = b) = \mu_0(a)K_{0:n}(a,b).$$

Definition 5.19 (Endpoint support)**.** *For a finite trajectory with initial law μ_0, define the endpoint support by*

$$\mathrm{Supp}_{0,n}^{\mu_0} := \{(a,b) \in X_0 \times X_n : \mu_0(a)K_{0:n}(a,b) > 0\}.$$

Proposition 5.20 (Support-safe weighted bridge entropy formula)**.** *For $t \in \{1, \ldots, n-1\}$, one has*

$$\boxed{\Phi_{t,\mu_0}^{(n)} = \sum_{(a,b)\in \mathrm{Supp}_{0,n}^{\mu_0}} \mu_0(a)K_{0:n}(a,b)\, H(B_t^{a,b}).}$$

Equivalently,

$$\boxed{\Phi_{t,\mu_0}^{(n)} = \sum_{(a,b)\in \mathrm{Supp}_{0,n}^{\mu_0}} \mu_0(a)K_{0:n}(a,b) H\left(x \longmapsto \frac{K_{0:t}(a,x)K_{t:n}(x,b)}{K_{0:n}(a,b)}\right).}$$

Proof. The finite conditional-entropy decomposition is taken only over atoms of the conditioning variable with positive probability. Therefore

$$H(X_t \mid X_0, X_n) = \sum_{(a,b)\in \mathrm{Supp}_{0,n}^{\mu_0}} \mathbb{P}(X_0 = a, X_n = b)\, H(X_t \mid X_0 = a, X_n = b).$$

For every $(a,b) \in \mathrm{Supp}_{0,n}^{\mu_0}$, Proposition 5.17 identifies the conditional law of X_t given $(X_0 = a, X_n = b)$ with $B_t^{a,b}$. Since

$$\mathbb{P}(X_0 = a, X_n = b) = \mu_0(a)K_{0:n}(a,b),$$

substitution gives the formula. ∎

This formula makes two dependencies explicit. First, the conditional bridge kernel depends on the factorization

$$K_{0:n} = K_{t:n} \star K_{0:t}.$$

Second, the averaged bridge entropy also depends on the endpoint law induced by μ_0. No conditional law is evaluated on endpoint pairs outside $\mathrm{Supp}_{0,n}^{\mu_0}$.

5.7 Factorization Sensitivity

Definition 5.21 (Marked two-block factorization)**.** *Fix a trajectory*

$$X_0 \overset{K_0}{\rightsquigarrow} X_1 \overset{K_1}{\rightsquigarrow} \cdots \overset{K_{n-1}}{\rightsquigarrow} X_n$$

and a marked internal time $t \in \{1, \dots, n-1\}$. The associated marked two-block factorization *is*

$$(X_0, \mu_0) \overset{K_{0:t}}{\rightsquigarrow} X_t \overset{K_{t:n}}{\rightsquigarrow} X_n.$$

Bridge entropy is attached to this marked two-block factorization. It is not an invariant of the unmarked endpoint composite alone, and it is not intended to record arbitrary refinements inside the two blocks once the two block composites are fixed.

Bridge entropy has an intermediate level of sensitivity. It is not an invariant of the bare endpoint composite $K_{0:n}$ alone, because different marked two-block factorizations

$$K_{0:n} = K_{t:n} \star K_{0:t}$$

can give different reconstruction uncertainty for X_t. On the other hand, once the marked two-block factorization

$$(X_0, \mu_0) \overset{K_{0:t}}{\rightsquigarrow} X_t \overset{K_{t:n}}{\rightsquigarrow} X_n$$

is fixed, bridge entropy is insensitive to finer refinements inside the left block $0 \to t$ and the right block $t \to n$. Thus

$$\boxed{\text{Bridge entropy sees the marked two-block factorization, not the full fine path.}}$$

This is the precise categorical reason bridge entropy is useful.

Proposition 5.22 (Two-step defect is not determined by the endpoint composite)**.** *There exist two pointed two-step factorizations*

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z, \qquad X \overset{K'}{\rightsquigarrow} Y' \overset{L'}{\rightsquigarrow} Z$$

with the same input law μ and the same endpoint composite

$$L \star K = L' \star K',$$

but with different entropy defects:

$$\Delta_\mu(K, L) \neq \Delta_\mu(K', L').$$

Consequently, $\Delta_\mu(K, L)$ is an invariant of the marked factorization, not of the bare endpoint composite alone.

Proof. Let $X = Z = \{*\}$, let $Y = \{0, 1\}$, and let $\mu(*) = 1$. Define

$$K(*, 0) = K(*, 1) = \frac{1}{2}, \qquad L(0, *) = L(1, *) = 1.$$

Then $L \star K$ is the deterministic one-point kernel. Since X and Z are both deterministic while Y is a fair bit,

$$\Delta_\mu(K, L) = H(Y \mid X, Z) = H(Y) = 1.$$

Now let $Y' = \{*\}$, with K' and L' both deterministic one-point kernels. The endpoint composite $L' \star K'$ is the same deterministic one-point kernel, but

$$\Delta_\mu(K', L') = H(Y' \mid X, Z) = 0.$$

Thus the endpoint composite does not determine the two-step defect. □

Corollary 5.23 (Bridge entropy is not determined by the endpoint composite)**.** *Bridge entropy is not a function of the bare endpoint composite $K_{0:n}$ alone. Already for $n = 2$ and $t = 1$, there exist two marked two-block factorizations with the same endpoint composite and different bridge entropies.*

Proof. Apply Proposition 5.22 to the length-two trajectory

$$X_0 = X, \qquad X_1 = Y, \qquad X_2 = Z.$$

For $n = 2$ and $t = 1$,

$$\Phi^{(2)}_{1,\mu} = H(X_1 \mid X_0, X_2) = H(Y \mid X, Z) = \Delta_\mu(K, L).$$

The two examples in Proposition 5.22 have the same endpoint composite but bridge entropies 1 and 0, respectively. □

Theorem 5.24 (Marked-factorization sensitivity)**.** *The quantities $\Delta_\mu(K, L)$ and $\Phi^{(n)}_{t,\mu_0}$ are marked-factorization invariants. They do not factor through the endpoint-composite map from marked factorizations to bare stochastic kernels.*

Proof. If either quantity factored through the endpoint composite, then any two pointed factorizations with the same endpoint composite would have the same value. Proposition 5.22 disproves this for Δ_μ, and Corollary 5.23 disproves this for $\Phi^{(n)}_{t,\mu_0}$. □

Proposition 5.25 (Marked-cut dependence and block-refinement invariance)**.** *Fix $t \in \{1, \ldots, n-1\}$. The bridge entropy*

$$\Phi^{(n)}_{t,\mu_0} = H(X_t \mid X_0, X_n)$$

is completely determined by the pointed marked two-block factorization

$$(X_0, \mu_0) \overset{K_{0:t}}{\rightsquigarrow} X_t \overset{K_{t:n}}{\rightsquigarrow} X_n.$$

Consequently:

(i) *two trajectories with the same $\mu_0, K_{0:t}, K_{t:n}$ have the same bridge entropy, even if their internal refinements inside the two blocks differ;*

(ii) *two trajectories with the same endpoint composite $K_{0:n}$ may have different bridge entropy if their marked two-block factorizations differ.*

Proof. Part (i) follows from the support-safe weighted bridge formula:

$$\Phi_{t,\mu_0}^{(n)} = \sum_{(a,b)\in \mathrm{Supp}_{0,n}^{\mu_0}} \mu_0(a) K_{0:n}(a,b) H(B_t^{a,b}),$$

where

$$B_t^{a,b}(x) = \frac{K_{0:t}(a,x)K_{t:n}(x,b)}{K_{0:n}(a,b)}, \qquad K_{0:n} = K_{t:n} \star K_{0:t}.$$

The right-hand side uses only $\mu_0, K_{0:t}, K_{t:n}$. Therefore any finer internal decomposition preserving these two composites leaves $\Phi_{t,\mu_0}^{(n)}$ unchanged.

Part (ii) is witnessed by Example 5.27. □

Data level	Does $\Phi_{t,\mu_0}^{(n)}$ depend on it?
Bare endpoint composite $K_{0:n}$ alone	No. The same endpoint composite can have different bridge entropy.
Marked two-block factorization $(\mu_0, K_{0:t}, K_{t:n})$	Yes. This is the exact data level seen by bridge entropy.
Finer refinements inside $0 \to t$ or $t \to n$	No, provided the two block composites are preserved.
Full hidden path	Only through the marked bridge state X_t, not through all intermediate coordinates.

Table 4: Sensitivity level of bridge entropy.

Remark 5.26 (What bridge entropy sees)**.** *Bridge entropy belongs exactly to a marked two-block factorization of an endpoint composite. The full path defect $\mathcal{D}_{\mu_0}^{(n)}$ records the whole hidden trajectory, whereas $\Phi_{t,\mu_0}^{(n)}$ records the reconstruction obstruction for one marked intermediate object X_t.*

The next example gives a minimal witness using the same intermediate object in both factorizations.

Example 5.27 (Same endpoint composite and same bridge object, different bridge entropy)**.** *Let*

$$X_0 = X_2 = \{*\}, \qquad X_1 = \{0,1\},$$

and take the unique initial distribution on X_0.

Factorization A: deterministic marked bridge. *Define*

$$K_0^A(*,\cdot) = \delta_0, \qquad K_1^A(0,\cdot) = K_1^A(1,\cdot) = \delta_*.$$

Then the endpoint composite is the unique deterministic morphism

$$K_1^A \star K_0^A : \{*\} \rightsquigarrow \{*\}.$$

Moreover $X_1 = 0$ almost surely, so

$$\Phi_1^{(2),A} = H(X_1 \mid X_0, X_2) = 0.$$

Factorization B: hidden fair marked bridge. *Define*

$$K_0^B(*,0) = K_0^B(*,1) = \frac{1}{2}, \qquad K_1^B(0,\cdot) = K_1^B(1,\cdot) = \delta_*.$$

The endpoint composite is again the same unique deterministic morphism

$$K_1^B \star K_0^B : \{*\} \rightsquigarrow \{*\}.$$

However, X_1 is a fair bit and the endpoints carry no information about it. Therefore

$$\Phi_1^{(2),B} = H(X_1 \mid X_0, X_2) = H(X_1) = 1 \text{ bits}.$$

Thus two pointed factorizations with the same objects

$$X_0 = \{*\}, \qquad X_1 = \{0,1\}, \qquad X_2 = \{*\},$$

and the same endpoint composite $X_0 \rightsquigarrow X_2$ can have different bridge entropy. The difference is caused by the marked two-block factorization, not by changing the intermediate state space.

5.8 Transition to Explicit Models

The bridge entropy formalism developed above is abstract but computable. The next section evaluates it for the binary symmetric channel. That example shows that bridge entropy can detect a fixed-parameter long-trajectory reconstruction discontinuity that is invisible to the endpoint-only normalized diagnostic on fixed finite state spaces.

6 Binary Symmetric Channel and Boundary Reconstruction

This section computes the obstruction quantities introduced above for the binary symmetric channel. The example has three purposes. First, it gives a concrete verification of the path entropy balance law and the bridge entropy formula. Second, it illustrates how endpoint entropy can remain too coarse on a fixed finite state space. Third, it exhibits a long-trajectory reconstruction discontinuity: deterministic boundary dynamics have zero asymptotic bridge entropy, while every genuinely stochastic interior channel has one full bit of limiting midpoint uncertainty.

6.1 The Binary Symmetric Channel

Let

$$X = \{0,1\}.$$

For $0 \le \theta \le 1$, the binary symmetric channel with crossover parameter θ is the standard two-state noise model from information theory [2]. It is the stochastic morphism

$$K^{(\theta)} : X \rightsquigarrow X$$

defined by

$$K^{(\theta)}(x,y) = \begin{cases} 1-\theta, & y = x, \\ \theta, & y \neq x. \end{cases}$$

Remark 6.1 (Parameter range for the BSC)**.** *In coding-theoretic conventions, the binary symmetric channel is often reduced to $0 \leq \theta \leq 1/2$, since a channel with $\theta > 1/2$ is equivalent to a deterministic bit flip followed by a BSC with crossover probability $1 - \theta < 1/2$. In this paper we deliberately keep the full Markov-kernel parameter interval*

$$0 \leq \theta \leq 1.$$

The reason is that the two boundary points have distinct deterministic categorical dynamics:

$$\theta = 0 \quad \textit{is the identity channel,} \qquad \theta = 1 \quad \textit{is the deterministic flip channel.}$$

Both endpoints have zero bridge reconstruction obstruction, while every interior parameter $0 < \theta < 1$ has full support and gives the long-trajectory bridge limit in Theorem 6.13.

Lemma 6.2 (Flip symmetry of the full parameter interval)**.** *Let $F : \{0, 1\} \to \{0, 1\}$ be the deterministic flip. Then*

$$K^{(\theta)} = \delta_F \star K^{(1-\theta)} = K^{(1-\theta)} \star \delta_F.$$

Consequently, for coding-theoretic capacity questions one may restrict to $\theta \leq 1/2$. For deterministic-boundary reconstruction, however, the full interval $[0, 1]$ is natural.

Proof. Composing with the deterministic flip interchanges the diagonal and off-diagonal transition probabilities. Checking the two rows gives the stated identity. □

Equivalently, if $B \sim \text{Bernoulli}(\theta)$, then one step of the chain is

$$X_{k+1} = X_k \oplus B,$$

where $\oplus$ denotes addition modulo 2. We write

$$\alpha = 1 - 2\theta.$$

Let $\mu_0 \in \mathsf{Prob}(\{0, 1\})$ be arbitrary. The endpoint entropy and bridge entropy formulas below are independent of this initial law after all conditional laws are restricted to positive-probability endpoint atoms.

The binary entropy function is

$$H_\mathrm{b}(p) = -p \log p - (1 - p) \log(1 - p),$$

with logarithms base 2 and the convention $0 \log 0 = 0$.

6.2 Conventions for the BSC Example

The BSC computation below uses three conventions.

First, all rational bridge formulas are stated only on the interior parameter range

$$0 < \theta < 1,$$

where the endpoint denominators are strictly positive. The deterministic endpoints $\theta = 0$ and $\theta = 1$ are evaluated separately.

Second, initial-law independence always means support-safe independence: conditional laws are evaluated only on endpoint atoms of positive probability under the chosen initial law.

Third, the limiting boundary phenomenon is a long-trajectory reconstruction discontinuity. It is not a finite-length singularity and is not claimed to be a thermodynamic phase transition.

Definition 6.3 (Reconstruction discontinuity versus phase transition)**.** *In this finite BSC example, a* long-trajectory reconstruction discontinuity *means a discontinuity of the limiting function*

$$\theta \longmapsto \lim_{n\to\infty} H(X_{\lfloor n/2\rfloor} \mid X_0, X_n),$$

where θ is fixed before the limit $n \to \infty$ is taken. This is not a finite-length singularity, because for every fixed $n \geq 2$ the bridge curve is regular on the open interval $(0,1)$. It is also not a thermodynamic phase transition unless a separate scaling theory, such as growing state spaces or joint parameter-length limits, is introduced.

6.3 Powers of the Channel

The powers of the binary symmetric channel remain binary symmetric. The following elementary formula is used throughout the section.

Lemma 6.4 (Powers of the binary symmetric channel)**.** *For every integer $r \geq 1$,*

$$\left(K^{(\theta)}\right)^r(x,y) = \begin{cases} s_r, & y = x, \\ q_r, & y \neq x, \end{cases}$$

where

$$s_r = \frac{1+\alpha^r}{2}, \qquad q_r = \frac{1-\alpha^r}{2}.$$

Proof. Let

$$p_r = \left(K^{(\theta)}\right)^r(0,1).$$

Starting from state 0, the chain is in state 1 after $r+1$ steps either because it was in state 0 after r steps and flips, or because it was already in state 1 after r steps and does not flip. Hence

$$p_{r+1} = (1-p_r)\theta + p_r(1-\theta) = \theta + (1-2\theta)p_r.$$

Since $p_0 = 0$, solving the linear recursion gives

$$p_r = \frac{1-(1-2\theta)^r}{2} = q_r.$$

Therefore the probability of returning to the starting state is

$$1 - p_r = \frac{1+(1-2\theta)^r}{2} = s_r.$$

By symmetry the same formula holds starting from either state. ☐

Lemma 6.5 (Interior positivity of BSC endpoint denominators)**.** *Let $0 < \theta < 1$, set $\alpha = 1 - 2\theta$, and define*

$$s_r = \frac{1+\alpha^r}{2}, \qquad q_r = \frac{1-\alpha^r}{2}.$$

Then, for every integer $r \geq 1$,

$$0 < s_r < 1, \qquad 0 < q_r < 1.$$

Consequently, the denominators s_n and q_n in the interior BSC bridge formula are strictly positive.

Proof. Since $0 < \theta < 1$, one has $|\alpha| < 1$. Hence

$$-1 < \alpha^r < 1$$

for every $r \geq 1$. Therefore

$$0 < \frac{1+\alpha^r}{2} < 1, \qquad 0 < \frac{1-\alpha^r}{2} < 1.$$

☐

6.4 Endpoint and Path Quantities

Consider the time-homogeneous stochastic trajectory

$$X_0 \overset{K^{(\theta)}}{\rightsquigarrow} X_1 \overset{K^{(\theta)}}{\rightsquigarrow} \cdots \overset{K^{(\theta)}}{\rightsquigarrow} X_n.$$

For every state $x \in \{0,1\}$, the row $K^{(\theta)}(x,\cdot)$ has binary entropy $H_{\mathrm{b}}(\theta)$. Hence for every transported law μ_k,

$$\mathcal{E}_{\mu_k}(K^{(\theta)}) = H_{\mathrm{b}}(\theta).$$

Thus the accumulated local conditional uncertainty is

$$\boxed{\mathcal{S}^{(n)} = nH_{\mathrm{b}}(\theta).} \tag{1}$$

By Lemma 6.4, the endpoint composite $(K^{(\theta)})^n$ flips the initial bit with probability

$$q_n = \frac{1-(1-2\theta)^n}{2}.$$

Hence

$$\boxed{\mathcal{C}^{(n)} = H(X_n \mid X_0) = H_{\mathrm{b}}(q_n) = H_{\mathrm{b}}\left(\frac{1-(1-2\theta)^n}{2}\right).} \tag{2}$$

Proposition 6.6 (Endpoint entropy and endpoint-only normalized diagnostic for the BSC)**.** *For the binary symmetric channel,*

$$\mathcal{C}^{(n)}(\theta) = H_{\mathrm{b}}\left(\frac{1-(1-2\theta)^n}{2}\right).$$

Consequently,

$$\lim_{n\to\infty} \frac{1}{n}\mathcal{C}^{(n)}(\theta) = 0$$

for every $\theta \in [0,1]$. Moreover,

$$\lim_{n\to\infty} \mathcal{C}^{(n)}(\theta) = \begin{cases} 0, & \theta = 0, \\ 1, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}$$

Proof. The displayed formula is Equation (2). Since $0 \le \mathcal{C}^{(n)}(\theta) \le 1$ for a binary endpoint, the endpoint-only normalized diagnostic tends to zero for every parameter. If $\theta = 0$, then $q_n = 0$ for all n, hence $\mathcal{C}^{(n)} = 0$. If $\theta = 1$, then $q_n = (1-(-1)^n)/2$, so $q_n \in \{0,1\}$ and again $\mathcal{C}^{(n)} = 0$. If $0 < \theta < 1$, then $|1-2\theta| < 1$, hence $q_n \to 1/2$, and continuity of H_{b} gives $\mathcal{C}^{(n)} \to H_{\mathrm{b}}(1/2) = 1$. □

Proposition 6.7 (Support-safe initial-law independence for BSC endpoint and bridge entropies)**.** *For the binary symmetric channel on $\{0,1\}$, the endpoint entropy*

$$H(X_n \mid X_0)$$

is independent of the initial law μ_0 for every $0 \le \theta \le 1$. For $0 < \theta < 1$, the bridge entropy

$$H(X_t \mid X_0, X_n)$$

is also independent of μ_0 in the support-safe sense: conditional bridge laws are evaluated only on endpoint atoms

$$(a, b) \in \mathrm{Supp}_{0,n}^{\mu_0}.$$

More precisely, for every positive-probability endpoint atom, the bridge entropy depends only on whether $b = a$ or $b \neq a$, not on the actual value of a. Moreover, for every initial law μ_0,

$$\mathbb{P}(X_n = X_0) = s_n, \qquad \mathbb{P}(X_n \neq X_0) = q_n.$$

Therefore the averaged interior bridge entropy equals

$$s_n H_{\mathrm{b}}\left(\frac{q_t q_{n-t}}{s_n}\right) + q_n H_{\mathrm{b}}\left(\frac{q_t s_{n-t}}{q_n}\right),$$

independently of μ_0. No zero-probability endpoint atom is conditioned on.

Proof. For every starting state $a \in \{0, 1\}$, the n-step transition row of the BSC has the same entropy:

$$H\big((K^{(\theta)})^n(a, \cdot)\big) = H_{\mathrm{b}}(q_n).$$

Therefore

$$H(X_n \mid X_0) = \sum_a \mu_0(a) H\big((K^{(\theta)})^n(a, \cdot)\big) = H_{\mathrm{b}}(q_n),$$

which is independent of μ_0.

Now assume $0 < \theta < 1$. For each fixed starting state a, the n-step BSC row assigns probability s_n to returning to a and probability q_n to ending at $1 - a$. Hence, after averaging over any μ_0,

$$\mathbb{P}(X_n = X_0) = \sum_a \mu_0(a) s_n = s_n, \qquad \mathbb{P}(X_n \neq X_0) = \sum_a \mu_0(a) q_n = q_n.$$

Restrict to endpoint atoms of positive probability. By flip symmetry, the pointwise bridge law for (a, a) is a relabeling of the pointwise bridge law for $(1-a, 1-a)$, hence has the same entropy. Likewise, the bridge law for $(a, 1-a)$ is a relabeling of the bridge law for $(1-a, a)$. Therefore there are only two entropy values to average: one for equal endpoints and one for unequal endpoints. The equal-endpoint value is

$$H_{\mathrm{b}}\left(\frac{q_t q_{n-t}}{s_n}\right),$$

and the unequal-endpoint value is

$$H_{\mathrm{b}}\left(\frac{q_t s_{n-t}}{q_n}\right).$$

Weighting them by s_n and q_n gives the displayed formula. The conditioning events used in this proof are exactly the positive-probability endpoint atoms under the selected initial law. □

The path entropy balance law gives the hidden path defect:

$$\boxed{\mathcal{D}^{(n)} = n H_{\mathrm{b}}(\theta) - H_{\mathrm{b}}\left(\frac{1 - (1 - 2\theta)^n}{2}\right).} \tag{3}$$

Remark 6.8 (Endpoint entropy versus endpoint-only normalized diagnostic)**.** *Proposition 6.6 separates two endpoint quantities:*

$$\mathcal{C}^{(n)}(\theta) = H(X_n \mid X_0), \qquad \mathsf{ER}^{(n)}(\theta) = \frac{1}{n} H(X_n \mid X_0).$$

The unnormalized quantity $\mathcal{C}^{(n)}$ distinguishes deterministic boundary parameters from fixed stochastic interior parameters in the long-time limit:

$$\lim_{n\to\infty} \mathcal{C}^{(n)}(\theta) = \begin{cases} 0, & \theta = 0, \\ 1, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}$$

Thus CEOT does not claim that endpoint entropy is blind to deterministic boundary behavior. The precise claim is only that the endpoint-only normalized diagnostic $\mathsf{ER}^{(n)}$ vanishes for every fixed parameter on a fixed finite state space. Bridge entropy records midpoint reconstruction uncertainty, whereas endpoint-only normalized diagnostic cannot record a nonzero per-step uncertainty rate in a fixed binary endpoint.

6.5 Bridge Entropy Formula

We next compute the bridge entropy

$$\Phi_t^{(n)}(\theta) = H(X_t \mid X_0, X_n)$$

for $t \in \{1, \ldots, n-1\}$. The endpoint cases $\theta = 0$ and $\theta = 1$ are deterministic and are treated separately below. For this reason the displayed bridge formula in this subsection is stated for

$$0 < \theta < 1.$$

Proposition 6.9 (Interior bridge entropy of the binary symmetric channel)**.** *Fix $n \geq 2$, $t \in \{1, \ldots, n-1\}$, and $0 < \theta < 1$. Then, by Lemma 6.5, $s_n > 0$ and $q_n > 0$, and*

$$\boxed{\Phi_t^{(n)}(\theta) = s_n H_{\mathrm{b}}\left(\frac{q_t\, q_{n-t}}{s_n}\right) + q_n H_{\mathrm{b}}\left(\frac{q_t\, s_{n-t}}{q_n}\right).} \tag{4}$$

This displayed rational formula is not evaluated at $\theta = 0$ or $\theta = 1$; the deterministic endpoint cases are handled separately in Lemma 6.12.

Proof. By Lemma 6.4 and Lemma 6.5,

$$\mathbb{P}(X_n = X_0) = s_n > 0, \qquad \mathbb{P}(X_n \neq X_0) = q_n > 0.$$

First condition on the event $X_n = X_0$. By the bridge kernel formula from Proposition 5.17, the probability that the midpoint state differs from the starting state is

$$\mathbb{P}(X_t \neq X_0 \mid X_0, X_n = X_0) = \frac{q_t\, q_{n-t}}{s_n}.$$

Indeed, the path must move from X_0 to the opposite state in t steps and then return to X_0 in $n-t$ steps.

Next condition on the event $X_n \neq X_0$. In this case

$$\mathbb{P}(X_t \neq X_0 \mid X_0, X_n \neq X_0) = \frac{q_t\, s_{n-t}}{q_n}.$$

The conditional law of X_t is binary in either endpoint case. Averaging the two conditional binary entropies with weights s_n and q_n gives (4). ☐

6.6 Worked Numerical Check

The following computation checks the bridge formula against the path entropy balance law in the smallest nontrivial case.

Example 6.10 (Two-step bridge with $\theta = 1/4$)**.** *Let*

$$n = 2, \qquad t = 1, \qquad \theta = \frac{1}{4}.$$

Then

$$\alpha = 1 - 2\theta = \frac{1}{2},$$

and therefore

$$s_1 = \frac{3}{4}, \qquad q_1 = \frac{1}{4},$$

while

$$s_2 = \frac{5}{8}, \qquad q_2 = \frac{3}{8}.$$

If the endpoints agree, then

$$\mathbb{P}(X_1 \neq X_0 \mid X_0, X_2 = X_0) = \frac{q_1^2}{s_2} = \frac{1/16}{5/8} = \frac{1}{10}.$$

Thus

$$H_{\mathrm{b}}(1/10) \approx 0.469.$$

If the endpoints differ, then

$$\mathbb{P}(X_1 \neq X_0 \mid X_0, X_2 \neq X_0) = \frac{q_1 s_1}{q_2} = \frac{(1/4)(3/4)}{3/8} = \frac{1}{2}.$$

Thus

$$H_{\mathrm{b}}(1/2) = 1.$$

The bridge entropy is therefore

$$\Phi_1^{(2)}(1/4) = \frac{5}{8} H_{\mathrm{b}}(1/10) + \frac{3}{8} H_{\mathrm{b}}(1/2) \approx 0.668 \text{ bits}.$$

Since $n = 2$, the bridge entropy is the full hidden path defect:

$$\Phi_1^{(2)} = \mathcal{D}^{(2)}.$$

The balance law gives

$$\mathcal{D}^{(2)} = 2H_{\mathrm{b}}(1/4) - H_{\mathrm{b}}(3/8) \approx 0.668 \text{ bits},$$

which agrees with the bridge computation.

Remark 6.11 (Interior formula is not evaluated at deterministic endpoints)**.** *Formula (4) is an interior formula. It is not evaluated at $\theta = 0$ or $\theta = 1$, where endpoint weights or normalizing constants may vanish. The endpoint values are supplied separately by Lemma 6.12.*

Lemma 6.12 (Deterministic endpoint parameters)**.** *For the binary symmetric channel:*

(i) *If $\theta = 0$, then $X_t = X_0$ for every t, and hence*

$$H(X_t \mid X_0, X_n) = 0.$$

(ii) *If $\theta = 1$, then*

$$X_t = X_0 \oplus t \pmod 2$$

for every t, and hence

$$H(X_t \mid X_0, X_n) = 0.$$

Proof. For $\theta = 0$, every step preserves the bit. For $\theta = 1$, every step flips the bit. In both cases X_t is a deterministic function of X_0 and t, hence is determined by (X_0, X_n). The conditional entropy is therefore zero. □

6.7 Fixed-Parameter Long-Trajectory Reconstruction Discontinuity

The explicit bridge formula reveals a discontinuous limiting reconstruction behavior at the deterministic boundary parameters $\theta = 0$ and $\theta = 1$. Define the midpoint bridge entropy by

$$\Phi^{(n)}(\theta) = H(X_{\lfloor n/2 \rfloor} \mid X_0, X_n).$$

Theorem 6.13 (Fixed-parameter long-trajectory reconstruction discontinuity)**.** *For the binary symmetric channel on $\{0, 1\}$, for every initial law $\mu_0 \in \mathsf{Prob}(\{0, 1\})$, the limit*

$$\Phi_\infty(\theta) = \lim_{n\to\infty} \Phi^{(n)}(\theta)$$

exists for every $\theta \in [0, 1]$. This is a long-trajectory reconstruction discontinuity in the sense of Definition 6.3, and

$$\boxed{\Phi_\infty(\theta) = \begin{cases} 0, & \theta = 0, \\ \log 2, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}}$$

Since logarithms are base 2, the interior value is $\log 2 = 1$ bit.

Proof. By Lemma 6.12, the bridge entropy is zero for $\theta = 0$ and $\theta = 1$, for every trajectory length and every initial law.

It remains to treat $0 < \theta < 1$. In that case the interior bridge formula of Proposition 6.9 applies, and

$$|\alpha| = |1 - 2\theta| < 1.$$

Let

$$t = \lfloor n/2 \rfloor.$$

Then

$$t \to \infty, \qquad n - t \to \infty, \qquad n \to \infty.$$

Hence

$$s_n, q_n, s_t, q_t, s_{n-t}, q_{n-t} \longrightarrow \frac{1}{2}.$$

In particular,

$$\frac{q_t\, q_{n-t}}{s_n} \longrightarrow \frac{(1/2)(1/2)}{1/2} = \frac{1}{2},$$

and

$$\frac{q_t\, s_{n-t}}{q_n} \longrightarrow \frac{(1/2)(1/2)}{1/2} = \frac{1}{2}.$$

By continuity of H_{b},

$$H_{\mathrm{b}}\left(\frac{q_t\, q_{n-t}}{s_n}\right) \to H_{\mathrm{b}}(1/2) = \log 2,$$

and

$$H_{\mathrm{b}}\left(\frac{q_t\, s_{n-t}}{q_n}\right) \to H_{\mathrm{b}}(1/2) = \log 2.$$

Using Proposition 6.9 together with Proposition 6.7,

$$\Phi^{(n)}(\theta) = s_n H_{\mathrm{b}}\left(\frac{q_t\, q_{n-t}}{s_n}\right) + q_n H_{\mathrm{b}}\left(\frac{q_t\, s_{n-t}}{q_n}\right).$$

Taking $n \to \infty$, we obtain

$$\Phi^{(n)}(\theta) \to \frac{1}{2}\log 2 + \frac{1}{2}\log 2 = \log 2.$$

Thus $\Phi_\infty(\theta) = \log 2$ for every $0 < \theta < 1$. □

6.8 Interpretation

Theorem 6.13 describes a fixed-parameter property of the long-trajectory limit

$$n \to \infty.$$

For finite n, the bridge entropy should not be interpreted as a finite-length singularity. The limiting statement is more precise:

$$\boxed{\text{for every fixed } \theta \in (0,1), \text{ midpoint reconstructibility is destroyed as } n \to \infty.}$$

At $\theta = 0$ and $\theta = 1$, the midpoint is exactly determined by the initial state and the deterministic rule. For every fixed $0 < \theta < 1$, however small the noise, the midpoint becomes asymptotically maximally uncertain after conditioning only on the endpoints.

Definition 6.14 (Finite-length singularity for the BSC bridge curve)**.** *Fix $n \geq 2$ and $t \in \{1, \dots, n-1\}$. In this finite BSC example, a* finite-length singularity *of the bridge curve means a point $\theta_0 \in (0,1)$ at which the map*

$$\theta \longmapsto H(X_t \mid X_0, X_n)$$

fails to be real-analytic. Endpoint support degeneracies at $\theta = 0$ and $\theta = 1$ are treated separately and are not called finite-length phase transitions in this paper. A genuine phase-transition theory would require a separate limiting or scaling framework, such as $n \to \infty$, growing state spaces, or another nontrivial scaling limit.

Proposition 6.15 (Fixed-length regularity of the BSC bridge curve)**.** *Fix $n \geq 2$ and $t \in \{1, \dots, n-1\}$. For the binary symmetric channel, the bridge curve*

$$\theta \longmapsto H(X_t \mid X_0, X_n)$$

is real-analytic on the open interval $(0,1)$. The deterministic endpoint values at $\theta = 0$ and $\theta = 1$ are zero. Therefore the discontinuity of

$$\Phi_\infty(\theta) = \lim_{n\to\infty} H(X_{\lfloor n/2 \rfloor} \mid X_0, X_n)$$

is a long-trajectory limiting effect, not a finite-length singularity of any fixed finite trajectory.

Proof. For $0 < \theta < 1$, the BSC has full support. Hence all endpoint conditioning events appearing in the bridge formula have positive probability. The explicit formula of Proposition 6.9 writes the bridge entropy as

$$s_n H_{\mathrm{b}}\left(\frac{q_t\, q_{n-t}}{s_n}\right) + q_n H_{\mathrm{b}}\left(\frac{q_t\, s_{n-t}}{q_n}\right),$$

where

$$s_r = \frac{1 + (1-2\theta)^r}{2}, \qquad q_r = \frac{1 - (1-2\theta)^r}{2}.$$

On $(0, 1)$, the denominators appearing in the displayed expression are positive on their support, and the arguments of H_{b} lie in $(0, 1)$. The functions s_r and q_r are polynomial functions of θ, and $H_{\mathrm{b}}(p)$ is real-analytic for $p \in (0, 1)$. Therefore the bridge curve is real-analytic on $(0, 1)$.

At $\theta = 0$ and $\theta = 1$, Lemma 6.12 gives exact deterministic reconstruction, so the bridge entropy is zero. Thus any jump described by Theorem 6.13 arises only after taking $n \to \infty$. It is not a finite-length singularity. □

Remark 6.16 (Non-uniformity near the deterministic boundary)**.** *The limiting discontinuity is caused by non-uniform convergence in the parameter θ. For every fixed $\theta \in (0, 1)$, the powers*

$$(1-2\theta)^{\lfloor n/2 \rfloor}$$

decay to zero as $n \to \infty$. This convergence is not uniform as $\theta \to 0$ or $\theta \to 1$. Hence the limiting bridge entropy may jump at the deterministic boundary even though each fixed trajectory length is a finite probabilistic object with no finite-length singularity.

Remark 6.17 (Order of limits)**.** *The theorem fixes $\theta \in (0, 1)$ first and then sends $n \to \infty$. It does not assert a uniform statement over joint scaling regimes $\theta = \theta_n \to 0$ or $\theta_n \to 1$. Near $\theta = 0$, the relevant powers involve*

$$\alpha_n = 1 - 2\theta_n.$$

For example, if $n\theta_n \to 0$, then

$$(1-2\theta_n)^{\lfloor n/2 \rfloor} \to 1,$$

so the chain remains close to deterministic over the observation window. If $n\theta_n \to \infty$, then

$$(1-2\theta_n)^{\lfloor n/2 \rfloor} \to 0,$$

and the fixed-interior behavior is recovered. A complete analysis of such joint-scaling regimes belongs to a separate scaling theory and is not part of the finite CEOT1 computation.

This phenomenon is invisible to the endpoint-only normalized diagnostic as a per-step endpoint diagnostic. Since the state space is fixed and binary,

$$H(X_n \mid X_0) \le 1,$$

and therefore

$$\lim_{n \to \infty} \frac{H(X_n \mid X_0)}{n} = 0.$$

By contrast, Proposition 6.6 shows that the unnormalized endpoint entropy still separates deterministic boundary parameters from fixed stochastic interior parameters. Bridge entropy therefore detects midpoint reconstruction uncertainty that endpoint-only normalized diagnostic cannot see; it is not claimed to be the only endpoint-sensitive diagnostic.

6.9 Transition to the General Outlook

The binary symmetric channel is deliberately elementary. Its value is not that it produces a complicated formula, but that it separates three notions that endpoint analysis tends to conflate: local conditional uncertainty contribution, endpoint uncertainty, and bridge-state reconstructibility. The final section summarizes the finite CEOT framework and explains which extensions require genuinely new analytic input.

7 Discussion and Outlook

This final section summarizes what the finite theory establishes, explains why endpoint entropy is not sufficient for a state-relative categorical obstruction calculus, and indicates the main directions in which the present finite framework should be extended. The point of the paper is not that the entropy identities used above are new as identities in information theory. They are consequences of the chain rule, conditional entropy, and Markov factorization. The point is that, once placed on stochastic morphisms and their factorizations, these quantities form a coherent layered obstruction calculus.

7.1 What the Finite Theory Establishes

The finite theory developed in this paper establishes four layers of state-relative numerical entropy obstruction invariants:

$$\mathcal{E}_\mu(K), \qquad \Delta_\mu(K,L), \qquad \mathcal{D}^{(n)}_{\mu_0}, \qquad \Phi^{(n)}_{t,\mu_0}.$$

They correspond, respectively, to single morphisms, two-step factorizations, full trajectories, and bridge states.

At the first level, Theorem 2.13 shows that

$$\mathcal{E}_\mu(K) = 0 \quad \Longleftrightarrow \quad K \text{ is deterministic on } \operatorname{supp}(\mu).$$

Thus kernel entropy detects the exact obstruction to visible deterministic realization, but only on the part of the source actually seen by the input law. Corollary 2.18 gives the full-support version:

$$\mathcal{E}_\mu(K) = 0 \quad \Longleftrightarrow \quad K = \delta_f$$

for a unique function $f : X \to Y$. This is the single-morphism obstruction.

At the second level, Theorem 3.3 identifies the two-step entropy defect as

$$\Delta_\mu(K,L) = H(Y \mid X, Z).$$

Therefore $\Delta_\mu(K,L)$ detects by vanishing whether the intermediate state in the factorization

$$X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$$

can be reconstructed from the endpoints. Corollary 3.6 then gives the exact criterion for endpoint-reconstructible two-step factorization. Example 3.11 also shows that endpoint-reconstructible factorization is strictly weaker than deterministic composition: the intermediate randomness may be genuine, but not lost.

At the third level, Theorem 4.7 gives the global balance law

$$\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0}.$$

This states that the accumulated local conditional uncertainty along a finite trajectory splits into endpoint-visible uncertainty and hidden path uncertainty. In expanded form,

$$H(X_n \mid X_0) + H(X_1, \ldots, X_{n-1} \mid X_0, X_n) = \sum_{k=0}^{n-1} H(X_{k+1} \mid X_k).$$

Thus $\mathcal{D}^{(n)}_{\mu_0}$ is the obstruction to reconstructing the whole hidden trajectory from its endpoints.

At the fourth level, Definition 5.1 introduces bridge entropy

$$\Phi^{(n)}_{t,\mu_0} = H(X_t \mid X_0, X_n).$$

This is a one-time marginal obstruction: it asks only whether the state at time t can be reconstructed from the endpoints. Proposition 5.7 gives the corresponding zero-entropy reconstruction criterion, and Proposition 5.17 gives the explicit finite bridge kernel under the positive-probability support condition. Example 5.27 shows that bridge entropy is not determined by the bare endpoint composite. It is a state-relative invariant of a marked two-block stochastic factorization.

The resulting layered obstruction calculus can be summarized as follows:

$$\begin{array}{ccl}
\mathcal{E}_\mu(K) & : & \text{failure of visible deterministic realization,} \\
\Delta_\mu(K, L) & : & \text{failure of endpoint two-step reconstruction,} \\
\mathcal{D}^{(n)}_{\mu_0} & : & \text{failure of full hidden path reconstruction,} \\
\Phi^{(n)}_{t,\mu_0} & : & \text{failure of bridge-state reconstruction.}
\end{array}$$

This layered calculus is the finite state-relative core of the CEOT program.

7.2 Why Endpoint Entropy Is Insufficient

A central lesson of the paper is that endpoint entropy alone is not enough. The endpoint composite

$$K_{0:n} : X_0 \rightsquigarrow X_n$$

forgets the factorization through the intermediate spaces

$$X_1, \ldots, X_{n-1}.$$

Consequently, quantities depending only on $K_{0:n}$ cannot distinguish factorizations that hide different amounts of intermediate information.

This limitation is already visible in the finite time-homogeneous case. If $X_k = X$ and $K_k = K$, then

$$H(X_n \mid X_0) \leq \log |X|.$$

Hence

$$\lim_{n\to\infty} \frac{H(X_n \mid X_0)}{n} = 0.$$

The endpoint-only normalized diagnostic is therefore trivial on fixed finite state spaces. Proposition 4.19 shows that, under finite irreducibility, the path defect rate instead converges to the stationary one-step conditional entropy rate of the Markov chain:

$$\lim_{n\to\infty} \frac{\mathcal{D}^{(n)}_{\mu_0}}{n} = \sum_x \pi(x) H(K(x, \cdot)).$$

Thus hidden path uncertainty, not endpoint-only normalized diagnostic, captures the conditional uncertainty accumulated along the trajectory.

The binary symmetric channel in Section 6 makes the same point more sharply. For every fixed finite state space, the endpoint-only normalized diagnostic vanishes. Nevertheless, the midpoint bridge entropy has a nontrivial limiting discontinuity:

$$\Phi_\infty(\theta) = \begin{cases} 0, & \theta = 0, \\ \log 2, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}$$

Thus every fixed positive stochasticity level destroys asymptotic midpoint reconstructibility in the long-trajectory limit, although this is not a uniform statement over joint limits in which the parameter approaches a deterministic boundary as the trajectory length grows. The endpoint-only normalized diagnostic does not detect this fixed-parameter long-trajectory reconstruction discontinuity as a per-step endpoint effect. This does not contradict the fact that the unnormalized endpoint entropy has its own boundary/interior long-time limit.

This is the conceptual reason CEOT emphasizes factorization-sensitive quantities. The obstruction is not only in where a stochastic morphism begins and ends, but also in how it factors.

7.3 Relation to Classical Information Theory

The mathematical identities used in this paper are classical. The equality

$$\Delta_\mu(K, L) = H(Y \mid X, Z)$$

is a chain-rule identity for the Markov chain $X \overset{K}{\rightsquigarrow} Y \overset{L}{\rightsquigarrow} Z$. The path balance law is likewise a chain-rule expansion of

$$H(X_1, \ldots, X_n \mid X_0)$$

in two different ways. The bridge kernel formula is the finite-state Markov bridge formula, closely related to Doob-transform constructions, reciprocal processes, and Schrödinger bridge problems [6, 7, 8, 9, 10, 11].

The contribution of CEOT is therefore not a new entropy calculus at the level of formal identities. Its contribution is the placement of these entropy quantities as state-relative numerical obstructions, invariant under relabeling isomorphism of the corresponding pointed stochastic structure. In ordinary information theory, one usually begins with random variables and asks how much information they carry about one another. In the present paper, one begins with stochastic morphisms and their compositions, and asks which categorical data can be reconstructed after composition, collapse, or endpoint conditioning.

This shift changes the interpretation of the same entropy expressions. The formulas remain standard finite entropy formulas; CEOT changes the pointed stochastic data to which they are attached and the reconstruction problems whose zero sets they classify. The conditional entropy

$$H(Y \mid X, Z)$$

becomes an obstruction to recovering the intermediate object in a stochastic factorization. The conditional entropy

$$H(X_1, \ldots, X_{n-1} \mid X_0, X_n)$$

becomes an obstruction to recovering a hidden categorical trajectory. The bridge entropy

$$H(X_t \mid X_0, X_n)$$

becomes a marked-factorization-sensitive obstruction to reconstructing an internal state. This is the numerical obstruction interpretation developed in the paper.

The categorical aspect is especially visible in Example 5.27. Two factorizations may have the same endpoint composite but different bridge entropy. Therefore bridge entropy is not a property of the endpoint stochastic matrix alone. It is a property of the pointed marked factorization. This is precisely the type of phenomenon that categorical organization is designed to retain.

7.4 Limits of the Present Paper

The present theory is intentionally finite and discrete. This restriction is not cosmetic; it avoids several serious issues that arise immediately in a measurable or continuous setting. In particular, this paper does not treat:

(1) measurable Markov kernels on standard Borel spaces;

(2) continuous-state stochastic morphisms;

(3) differential entropy or infinite entropy;

(4) disintegration beyond the finite case;

(5) growing state spaces and thermodynamic limits;

(6) stochastic functors from general categories into Markov categories;

(7) finite numerical reconstruction certificates beyond finite trajectories.

These exclusions are deliberate. In the finite setting, all conditional laws used in the paper exist on their positive-probability supports, all entropies are finite, all pointwise bridge kernels are explicit, and all reconstruction criteria reduce to elementary support conditions. This makes the obstruction mechanism transparent.

There is also a conceptual limitation. The fixed-parameter long-trajectory reconstruction discontinuity for the binary symmetric channel is a long-trajectory limiting discontinuity, not a thermodynamic phase transition in the usual statistical-mechanical sense and not a singularity at any fixed finite trajectory length. A genuine phase-transition theory for CEOT should involve growing state spaces, scaling limits, or nontrivial limiting families of stochastic categories. This belongs to a later stage of the theory.

Finally, the present paper studies stochastic morphisms inside a fixed finite ambient category. It does not yet study numerical entropy obstruction invariants for functors

$$F : \mathcal{C} \to \mathbf{FinStoch}_{\neq\varnothing}$$

from a nontrivial source category $\mathcal{C}$. Such functorial versions would be closer to an abstract categorical reconstruction-obstruction framework, but they should be developed only after the finite morphism-level theory is stable.

7.5 Future Directions

The finite theory suggests several natural extensions.

Measurable CEOT. The first extension is from finite stochastic matrices to Markov kernels on standard Borel spaces. In that setting, one must handle regular conditional probabilities, disintegration, null sets, absolute continuity, and possible infinite entropy. Bridge kernels need not be given by finite sums; they require measurable disintegration. A measurable CEOT would therefore need hypotheses ensuring that bridge reconstruction and entropy defects remain well-defined.

Growing state spaces and scaling limits. The binary symmetric channel shows a fixed-parameter long-trajectory reconstruction discontinuity in a long-trajectory limit, but the endpoint-only normalized diagnostic remains trivial. To obtain genuinely nontrivial entropy-rate and scaling-discontinuity phenomena, one should allow state spaces to grow:

$$|X_n| \to \infty,$$

or introduce refinement limits of stochastic categories. The relevant question is not simply whether the definitions generalize, but whether bridge entropy and path defects exhibit nontrivial scaling laws.

Finite reconstruction certificates for stochastic functors. A more categorical extension is to study functors

$$F : \mathcal{C} \to \mathbf{FinStoch}_{\neq\varnothing}$$

and define finite reconstruction certificates measuring failure of exact deterministic realization. One possible direction is to compare a stochastic functor with deterministic subfunctors or Dirac lifts. Numerical entropy defects could then measure failures of compatibility between composition in $\mathcal{C}$ and endpoint-reconstructible stochastic reconstruction in $\mathbf{FinStoch}_{\neq\varnothing}$.

Applications to hidden-state dynamics and algorithms. The quantities $\mathcal{D}^{(n)}_{\mu_0}$ and $\Phi^{(n)}_{t,\mu_0}$ are natural diagnostics for hidden-state loss. Potential applications include randomized algorithms, stochastic computation, latent-state inference, filtering, and coarse-grained Markov models. In such settings, endpoint behavior may be easy to observe while intermediate states are hidden. CEOT provides a language for separating endpoint uncertainty from hidden trajectory uncertainty.

Comparison with algorithmic reconstruction obstructions. The reconstruction-obstruction viewpoint is also compatible with algorithmic settings in which one studies whether a local update, path, or lifted state can be compressed without loss. In such contexts, entropy defects may serve as quantitative certificates of non-reconstructibility. This suggests possible connections with lifted-state dynamic programming, stochastic path optimization, and compression-sensitive inference systems.

7.6 Conclusion

This paper establishes the finite state-relative first layer of the CEOT program. The central objects are not new probability spaces, but finite numerical reconstruction certificates attached to pointed stochastic categorical structure. The resulting theory should be read as a disciplined finite reconstruction calculus: exact, support-relative, diagram-invariant, and deliberately not a cohomological obstruction theory, an abstract Markov-categorical entropy functor, or a thermodynamic phase-transition theory. The finite results can be compressed into five statements:

$$\mathcal{E}_\mu(K) = 0 \quad \Longleftrightarrow \quad K \text{ is deterministic on } \operatorname{supp}(\mu),$$

$$\Delta_\mu(K,L) = H(Y \mid X,Z),$$

$$\mathcal{C}^{(n)}_{\mu_0} + \mathcal{D}^{(n)}_{\mu_0} = \mathcal{S}^{(n)}_{\mu_0},$$

$$\mathcal{D}^{(n)}_{\mu_0} = \sum_{k=1}^{n-1} \Delta_k^{\mathrm{comp}}, \qquad \Delta_k^{\mathrm{comp}} = \Delta_{\mu_0\to\mu_k}(K_{0:k}, K_k), \qquad \mu_k = \mu_0 K_{0:k},$$

and, for the binary symmetric channel,

$$\Phi_\infty(\theta) = \begin{cases} 0, & \theta = 0, \\ \log 2, & 0 < \theta < 1, \\ 0, & \theta = 1. \end{cases}$$

Together, these results show that entropy can be organized as a state-relative exact numerical obstruction invariant and uncertainty score for visible deterministic realization, factorization, path reconstruction, and bridge reconstruction.

The next stage of CEOT should not merely repeat these definitions in more general categories. It should identify settings where factorization-sensitive entropy has nontrivial scaling, where bridge reconstruction undergoes genuine limiting transitions, and where stochastic categorical structure imposes obstructions that are invisible from endpoint data alone.

Remark 7.1 (Final terminology check)**.** *The terms "entropy obstruction", "path defect", and "bridge obstruction" are used only in the finite numerical sense developed in this paper: they denote conditional-entropy invariants with exact zero-set reconstruction criteria. They do not denote cohomology classes, derived obstruction classes, or thermodynamic entropy production.*

References

[1] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[2] T. M. Cover and J. A. Thomas. *Elements of Information Theory.* Second edition. Wiley-Interscience, 2006.

[3] J. G. Kemeny and J. L. Snell. *Finite Markov Chains.* Springer, 1976.

[4] J. R. Norris. *Markov Chains.* Cambridge University Press, 1997.

[5] D. A. Levin, Y. Peres, and E. L. Wilmer. *Markov Chains and Mixing Times.* Second edition. American Mathematical Society, 2017.

[6] J. L. Doob. *Classical Potential Theory and Its Probabilistic Counterpart.* Springer, 1984.

[7] B. Jamison. Reciprocal processes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 30:65–86, 1974.

[8] B. Jamison. The Markov processes of Schrödinger. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 32:323–331, 1975.

[9] A. Beghi. On the relative entropy of discrete-time Markov processes with given end-point densities. *IEEE Transactions on Information Theory*, 42(5):1529–1535, 1996.

[10] C. Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems*, 34(4):1533–1574, 2014.

[11] Y. Chen, T. T. Georgiou, and M. Pavon. On the relation between optimal transport and Schrödinger bridges: a stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169(2):671–691, 2016.

[12] B. Fong. Causal theories: A categorical perspective on Bayesian networks. Master's thesis, University of Oxford, 2012.

[13] T. Fritz. A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239, 2020.

[14] K. Cho and B. Jacobs. Disintegration and Bayesian inversion via string diagrams. *Mathematical Structures in Computer Science*, 29(7):938–971, 2019.