Part VI.B: Localic Systems, Machines and Probabilistic Descent

October 2026

1 Why localic systems?

Part VI.A studied strong monads on the arrow category $\mathbf{Set}^{\mathbf{2}}$. Garner's sheaf-comonad calculation ([4, arXiv v1, §9.3, Proposition 105]) identifies this category with the sheaves on the Sierpiński space. An arrow is therefore not only a map of witnesses to outcomes: it can also be understood as a system of local states varying over two specialisation contexts. The natural generalisation is the category $\mathcal{E} = \mathbf{Sh}(L)$ of sheaves on a locale L. The correct operational objects remain *strong* monads, equivalently $\mathcal{E}$-enriched monads for the cartesian closed enrichment of $\mathcal{E}$ over itself.

2 The first test: a finite cospan locale

Take the finite Alexandrov space with specialisation order $P = (p \to r \leftarrow q)$. Its opens are the upper sets of P, and its sheaf category is equivalent to the category of *covariant* functors $\mathbf{Set}^P$, whose objects are diagrams $A \xrightarrow{f} C \xleftarrow{g} B$. With opens $U = \{p, r\}$, $V = \{q, r\}$ and $W = U \cap V = \{r\}$, the corresponding sections satisfy

$$F(U) = A, \qquad F(V) = B, \qquad F(W) = C, \qquad F(X) = A \times_C B.$$

The covariant representable arities satisfy $y_p \times y_q \cong y_r$ and $1 \cong y_p \amalg_{y_r} y_q$. Thus the finite-colimit closure of representable arities already records the gluing of local contexts. The full subcategory of finitely presentable objects, equivalently the diagrams with finite components, is closed under finite products and contains the terminal object.

3 What is a machine internal to a sheaf topos?

Before studying probabilistic choice, it is helpful to recover familiar machines in this new ambient category. The guiding idea is simple: a machine no longer has just one set of states. Its state object is a sheaf S, so one has local states over each open, restriction between opens, and gluing of compatible local states. Machine operations must respect this sheaf structure. This is the operational content of insisting on strong, or equivalently self-enriched, theories and runners. We write $y = \mathrm{Id}_{\mathcal{E}}$ and $N(Z) = \mathbb{N}_{\mathcal{E}} \times Z$, the free strong monad on one unary operation. A runner from N to a strong monad R with residual state S is specified by a morphism $S \to R(S)$: the generating unary command acts on the state, while cartesian strength carries an arbitrary continuation along unchanged. This follows from freeness and the enriched Yoneda correspondence. Proposition 5.1 below identifies the whole measuring comonad with the cofree strong comonad, under the explicit duoidal and existence hypotheses. The Moore case, whose theory carries equations, is treated separately.

3.1 Deterministic unary machines: $\{N, y\}$

A residue-free implementation of the unary command is an endomorphism

$$d : S \longrightarrow S.$$

Over the Sierpiński space, S is one arrow $f : A \to B$, and d is a commuting square

$$\begin{array}{ccc} A & \xrightarrow{d_A} & A \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ B & \xrightarrow{d_B} & B. \end{array} \qquad f d_A = d_B f.$$

So the system has two state sets and two dynamics, constrained to agree under refinement. A trajectory of $a \in A$ is sent under f to the trajectory of $f(a) \in B$. Under the hypotheses of Proposition 5.1, the associated stream comonad is

$$\{N, y\}(S) \cong S^{\mathbb{N}_\mathcal{E}},$$

with counit evaluation at zero and comultiplication given by successive tails. For the cospan $S = (A \xrightarrow{f} C \xleftarrow{g} B)$ we similarly require three maps d_A, d_B, d_C with

$$f d_A = d_C f, \qquad g d_B = d_C g.$$

On global states $S(X) = A \times_C B$, the induced dynamics is

$$d_X(a, b) = (d_A(a), d_B(b)).$$

The compatibility equations make the right side a compatible pair. This is the first sense in which gluing and temporal evolution commute.

3.2 Counted unary runners: $\{N, N\}$

Now allow each source step to request some number of target unary steps. The generator is interpreted by

$$\rho : S \longrightarrow \mathbb{N}_\mathcal{E} \times S, \qquad \rho(s) = (k(s), d(s)).$$

Thus d is the updated state and k is the number of target operations required to implement one source operation. In a cospan the data consist of (d_A, k_A), (d_B, k_B), (d_C, k_C) and the conditions

$$\begin{aligned} f d_A &= d_C f, \quad k_A = k_C f, \\ g d_B &= d_C g, \quad k_B = k_C g. \end{aligned}$$

In particular, a global pair $(a, b) \in A \times_C B$ has an unambiguous operation count:

$$k_A(a) = k_C(fa) = k_C(gb) = k_B(b).$$

By Proposition 5.1, under its stated hypotheses, the measuring comonad has the counted-stream description

$$\{N, N\}(S) \cong (S \times \mathbb{N}_\mathcal{E})^{\mathbb{N}_\mathcal{E}}.$$

Its counit extracts only the state at the beginning of the stream. The sequence of counts is additional residual operational information. This makes $\{N, N\}$ a minimal example of a cost-aware machine; it is not yet an amortisation theorem, since we have imposed no potential inequalities.

3.3 Moore machines: $\{T_{\mathrm{Moore}}, y\}$

Fix internal input and output objects $I, O \in \mathcal{E}$. An internal Moore machine consists of

$$o : S \longrightarrow O, \qquad d : S \times I \longrightarrow S.$$

The map o is an observation; it leaves the state unchanged. For each input, d updates the state. The Moore theory has these operations and the observation equations giving precisely this comodel structure. Equivalently, by currying, this is a coalgebra for

$$H(S) = O \times S^I.$$

For the first examples take constant sheaves $I = \underline{A_0}$ and $O = \underline{B_0}$, where A_0 is an input alphabet and B_0 an output alphabet. On the Sierpiński space we obtain two ordinary Moore machines

$$(S_0, o_0, d_0), \qquad (S_1, o_1, d_1)$$

and a Moore-machine homomorphism $f : S_0 \to S_1$:

$$o_1 f = o_0, \qquad f d_0(s, a) = d_1(f(s), a).$$

Thus f can forget distinctions between states, but it cannot change the observed output or the response to an input word. Two states identified by f necessarily have the same complete output behaviour. For the cospan locale a Moore machine is a cospan *of Moore machines*

$$(A, o_A, d_A) \xrightarrow{f} (C, o_C, d_C) \xleftarrow{g} (B, o_B, d_B),$$

where both arrows preserve output and every input transition. Its global machine has states $A \times_C B$ and operations

$$o_X(a, b) = o_A(a) = o_C(fa) = o_B(b), \qquad d_X((a, b), i) = (d_A(a, i), d_B(b, i)).$$

For fixed constant I, O, the cofree Moore comonad has the familiar labelled-tree form

$$\{T_{\mathrm{Moore}}(I, O), y\}(S) \cong (S \times O)^{I^*},$$

subject to the same qualification about identifying an explicit cofree object with the Sweedler measuring object. The counit reads the state label at the empty word; derivatives $t \mapsto t_i$, with $t_i(w) = t(iw)$, are the input transitions. Since all these maps live in $\mathcal{E}$, restriction of a labelled tree commutes with taking derivatives.

A worked cospan Moore machine with genuinely different contexts. Take one input letter t, outputs $\{0, 1\}$, and

$$C = \{c_0, c_1\}, \quad A = \{a_0, a_0', a_1\}, \quad B = \{b_0, b_0', b_1\}.$$

The maps $f : A \to C$ and $g : B \to C$ send each subscript-j state to c_j. The output of each state is its subscript. Put $d_C(c_0) = c_1$, $d_C(c_1) = c_0$, and define

$$\begin{array}{lll} d_A(a_0) = a_1, & d_A(a_0') = a_1, & d_A(a_1) = a_0, \\ d_B(b_0) = b_1, & d_B(b_0') = b_1, & d_B(b_1) = b_0'. \end{array}$$

Both maps to C are Moore homomorphisms, but neither is invertible. The global state set is

$$A \times_C B = \{(a_0, b_0), (a_0, b_0'), (a_0', b_0), (a_0', b_0'), (a_1, b_1)\}.$$

All four states over c_0 move to (a_1, b_1), and that state moves to (a_0, b_0'). Thus gluing produces a five-state machine, not an isomorphic copy of any of its three local machines. Nevertheless the observable streams depend only on the subscript: the behavioural (Myhill–Nerode) quotient has two states, alternating 0 and 1. This illustrates both the enlargement of compatible global state space and the subsequent collapse under observability.

Varying interfaces. We can also let I, O be nonconstant sheaves. Then inputs themselves restrict across contexts, and the transition has to be a map $S \times I \to S$ of sheaves. The cofree behaviour expression uses the *internal* free monoid I^* and internal exponentials. One must not replace S^I by an arbitrary pointwise family $S_x^{I_x}$: exponentials in a presheaf topos are generally not pointwise. This is one reason to insist on enrichment from the beginning.

3.4 A common systems pattern

The three examples are instances of a common principle: a machine internal to $\mathcal{E}$ has local states and operations whose compatibility is expressed by naturality or, for a general space, by morphisms of sheaves. With fixed constant interfaces, restriction maps are machine homomorphisms. With varying interfaces, the strength and internal-hom structure also control how operations depend on the available local inputs and observations.

4 Strong polynomial lotteries

Fix rational weights throughout: let $\Delta_n(\mathbb{Q})$ be the set of nonnegative rational probability vectors with $n \geq 1$ indexed tickets. We retain ticket order, including zero-weight tickets, and define

$$\mathsf{rand}(Z) = \coprod_{n \geq 1} \coprod_{\mathbf{p} \in \Delta_n(\mathbb{Q})} Z^n.$$

The monad unit uses one ticket and multiplication flattens nested tickets with weights $p_i q_{ij}$. For any sheaf topos $\mathcal{E}$, this defines an internal polynomial strong monad

$$\mathsf{rand}_{\mathcal{E}}(F) = \coprod_{n \geq 1} \coprod_{\mathbf{p} \in \Delta_n(\mathbb{Q})} F^{\underline{n}},$$

where $\underline{n} = n \cdot 1$ and the sums are internal coproducts. The strength carries a fixed parameter into every ticket, and the monad maps respect it. The functor preserves pullbacks: this follows internally from the disjointness and pullback stability of coproducts and the pullback preservation of exponentials $(-)^{\underline{n}}$.

Proposition 4.1 (Finite cospan gluing)**.** *For maps $A \to C \leftarrow B$ of sets, there is a natural isomorphism*

$$\mathsf{rand}(A \times_C B) \cong \mathsf{rand}(A) \times_{\mathsf{rand}(C)} \mathsf{rand}(B).$$

Proof. Equality of two image lotteries over C forces equality of ticket count, probability vector, and each indexed observation in C. The corresponding tickets from A and B therefore pair uniquely to tickets in $A \times_C B$. □

This differs from ordinary finite distributions: compatible marginals do not determine a unique coupling. The polynomial lottery remembers a coupling through indexed tickets. This rigidity will matter in the monodromy example below (compare Remark 6.1).

5 The enriched Sweedler setting

For $\mathcal{E} = \mathbf{Sh}(L)$, the intended ambient category is a suitable category of accessible strong endofunctors, with functor composition and enriched Day convolution. The fixed-rank construction uses a small product-closed category $\mathcal{J}$ of enriched arities, identifies the chosen endofunctors via enriched left Kan extension, and checks local presentability, biclosed convolution, accessible composition, and coherent interchange. Under these hypotheses the general measuring-comonoid theorem of López Franco–Vasilakopoulou gives $\{T, R\}$ for strong monads T, R. We use this as a *conditional general criterion*; a rank-independent theorem and comparison of measuring objects across ranks are not assumed here. We keep the free-unary result distinct from source theories with equations. The notation in [2] uses $\star$ for our composition product $\circ$ and $\diamond$ for Day convolution; the duoidal interchange has direction

$$(F \circ G) \diamond (H \circ K) \longrightarrow (F \diamond H) \circ (G \diamond K).$$

Proposition 5.1 (Free unary measuring object, conditional on the ambient duoidal universe)**.** *Suppose the López Franco–Vasilakopoulou hypotheses hold in a category $\mathcal{F}$ of strong endofunctors, and the arity system contains* 1*, so its Day-convolution unit is $y = \mathrm{Id}$. If the cofree $\circ$-comonoid R^ν on R belongs to $\mathcal{F}$, then*

$$\{N, R\} \cong R^\nu, \qquad R^\nu(S) \cong \nu Z.\, S \times R(Z).$$

Proof sketch. The $\circ$-monoid structure on $[C, R]_\diamond$ comes from the $\circ$-comonoid structure of C, the $\circ$-monoid structure of R, and the duoidal interchange. Since N is the free $\circ$-monoid on y, monoid maps $N \to [C, R]_\diamond$ correspond to maps $y \to [C, R]_\diamond$. By convolution closure these correspond to maps $C \diamond y \cong C \to R$. Hence the measuring comonad represents the same functor of $\circ$-comonoids as the cofree comonoid on R. The convolution internal-hom convention is fixed by this adjunction; existence in the chosen accessibility universe is a separate requirement. □

For the free unary strong monad $N(Z) = \mathbb{N}_\mathcal{E} \times Z$, a runner into a strong monad R on residual state S amounts to a map $S \to R(S)$. By Proposition 5.1, provided the cofree strong comonad lies in the chosen measuring universe, the unary measuring object is the cofree comonad on the underlying strong endofunctor R:

$$\{N, R\}(S) \simeq \nu Z.\, S \times R(Z).$$

This gives streams for $R = \mathrm{Id}$, counted streams for $R = N$, and indexed probabilistic trees for $R = \mathsf{rand}_\mathcal{E}$. This identification uses the measuring universal property in Proposition 5.1; it does not extend automatically to source theories with equations, such as the Moore theory.

6 Beyond Alexandrov spaces: monodromy on the circle

We now take $X = S^1$, a non-Alexandrov topological space, and the nontrivial two-sheeted covering map $\pi : \widetilde{X} \to X$. Let F be its sheaf of local sections. On each simply connected short arc $U \subseteq X$, the restriction $F|_U$ admits two local sections a_U, b_U, but continuation once around the circle interchanges them. In particular,

$$\Gamma(F) = \varnothing.$$

Cover S^1 by two arcs U, V whose intersection has two components, and choose local trivialisations so that the transition function swaps the two sheets on one component but not the other. This represents the nontrivial monodromy cocycle.

6.1 Global sections and monodromy invariants

A locally constant sheaf G on a connected, locally simply connected space is classified by an action of $\pi_1(X)$ on a fibre G_0. Its global sections are precisely the fixed points:

$$\Gamma(G) \cong G_0^{\pi_1(X)}.$$

For the double cover, the fibre of F is the transposition set $\{0, 1\}$. Since every lottery has a positive number of tickets, a global ordered polynomial lottery requires global state sections for all its tickets. Thus $\Gamma(F) = \varnothing$ already implies $\Gamma(\mathsf{rand}_{\mathcal{E}} F) = \varnothing$; this vanishing is inherited from the absence of global states, and is not peculiar to lotteries. More strongly, $\mathsf{rand}(\tau)$ fixes no ordered lottery on the two-element fibre, and hence

$$\Gamma(\mathsf{rand}_{\mathcal{E}} F) = \mathsf{rand}(2)^{\mathbb{Z}/2} = \varnothing.$$

The local two-ticket expressions $(\frac{1}{2}, \frac{1}{2}; a_U, b_U)$ and $(\frac{1}{2}, \frac{1}{2}; b_U, a_U)$ illustrate the *stronger* ticket-sensitive fact: transporting around the circle exchanges their labels, so they are not equal in the polynomial lottery sheaf. This is not a second proof by gluing; it explains why these particular local lotteries do not match on overlaps.

6.2 An invariant probability distribution does glue

We use the same rational weights as for rand. Let $\mathsf{D}_{\mathcal{E}}$ denote the rational finite-distribution monad interpreted internally in $\mathcal{E}$; on locally constant finite sheaves over the circle it is obtained by applying the ordinary rational-distribution construction to each fibre and to monodromy. The aggregation map $\mathsf{rand}_{\mathcal{E}} \to \mathsf{D}_{\mathcal{E}}$ is compatible with monad multiplication. On a trivialising arc put

$$\mu_U = \tfrac{1}{2}\delta_{a_U} + \tfrac{1}{2}\delta_{b_U}.$$

This is invariant under exchange of the sheets. In fact it is the *unique* monodromy-invariant rational distribution on the two-sheet fibre, so

$$\boxed{\Gamma(\mathsf{D}_{\mathcal{E}} F) = \{\tfrac{1}{2}\delta_0 + \tfrac{1}{2}\delta_1\}, \qquad \Gamma(\mathsf{rand}_{\mathcal{E}} F) = \varnothing.}$$

Here the displayed distribution is understood as a global section of the distribution sheaf, not as a global section choosing either sheet.

Remark 6.1 (Connected global sections)**.** *For this connected space and constant discrete lottery shapes,* $\Gamma(\mathsf{rand}_{\mathcal{E}} F) \cong \mathsf{rand}(\Gamma F)$*. Thus the polynomial obstruction does* not *arise from failure of this comparison. For general disconnected locales, shape choices can vary between connected components. For non-locally-connected locales, infinite powers of constant sheaves require care (finite powers cause no such difficulty).*

6.3 Between ordered lotteries and distributions: symmetric lotteries

There is a natural intermediate construction on sets

$$\mathsf{Q}(Z) = \coprod_{n \geq 1} \big(\Delta_n(\mathbb{Q}) \times Z^n\big)/\mathfrak{S}_n,$$

where the symmetric group simultaneously permutes the weights and labels (the coproduct notation denotes orbits of the combined action). It identifies orderings but does not aggregate repeated outcomes. On the locally constant sheaves used here, $\mathsf{Q}_{\mathcal{E}}$ is formed by applying Q to each fibre and transporting the monodromy action through Q; equivalently, it is the local system associated with that induced action. Thus there is a sequence of monad morphisms

$$\mathsf{rand} \longrightarrow \mathsf{Q} \longrightarrow \mathsf{D}.$$

The class $[\frac{1}{2}, \frac{1}{2}; 0, 1]$ is fixed by swapping the two sheets and hence descends on the double cover. More generally, a symmetric lottery class is invariant precisely when, for each rational weight w, the number of tickets $(w, 0)$ equals the number of tickets $(w, 1)$. Hence $\Gamma(\mathsf{Q}_{\mathcal{E}}F)$ is infinite: zero-weight tickets provide additional padding classes, and positive weights can also be subdivided into more paired tickets. The three global-section sets form a useful hierarchy:

$$\Gamma(\mathsf{rand}_{\mathcal{E}}F) = \varnothing, \qquad |\Gamma(\mathsf{Q}_{\mathcal{E}}F)| = \infty, \qquad |\Gamma(\mathsf{D}_{\mathcal{E}}F)| = 1.$$

All the symmetric global lotteries aggregate to the unique equal-weight global distribution. However, the quotient no longer preserves pullbacks strictly. With $A = B = \{0, 1\}$ and $C = 1$, the pair of symmetric equal-weight lotteries on A and B admits distinct diagonal and anti-diagonal lifts to $\mathsf{Q}(A \times B)$. Thus the canonical pullback comparison is not injective. This is the familiar polynomial-versus-analytic trade-off: symmetric quotients allow monodromy-invariant lotteries, at the cost of losing unique ticketwise couplings. A groupoid-valued construction retaining the permutation automorphisms rather than identifying them is a candidate for further research; no measuring theorem for it is claimed here.

7 Machines on the circle: local dynamics and monodromy

The double cover gives examples of machines which exist perfectly well as sheaf objects even though they admit no globally selected state. Write F for the sheaf of local sections of the nontrivial two-sheeted cover of S^1. The fibre has two elements, but continuation around a generator of $\pi_1(S^1)$ exchanges them. This is a two-element local system with monodromy the transposition τ.

7.1 Deterministic dynamics without a global state

An endomorphism $d : F \to F$ of sheaves corresponds to an equivariant endomorphism of a two-element set with transposition action:

$$d\tau = \tau d.$$

There are exactly two: $d = 1$ and $d = \tau$. Both define $\{N, y\}$-coalgebraic dynamics on F, although $\Gamma(F) = \varnothing$. The second machine swaps the local sheets every time the unary command is executed. After one circuit in the base the sheet labels have exchanged, but the operation is still globally well defined because it commutes with that exchange. The stream sheaf $F^{\underline{\mathbb{N}}}$ records the complete local trajectory. A global point of that sheaf would have a global state as its initial value, so it has no global point. That does not prevent the sheaf machine or its local trajectories from existing.

7.2 Counted dynamics

A counted runner $F \to \underline{\mathbb{N}} \times F$ consists of one of the two equivariant dynamics above together with an equivariant function from the two-element transposition set to $\mathbb{N}$ with trivial action. Equivariance

forces the count to be the same on both sheets. Thus the entire family is described by

$$(d, k), \qquad d \in \{1, \tau\}, \quad k \in \mathbb{N}.$$

The count is a spatially coherent operational cost, even though it is not attached to a globally selected sheet.

7.3 Moore machines: what can be observed?

Take a constant finite input alphabet $I = \underline{A_0}$ and a constant discrete output alphabet $O = \underline{B_0}$. A Moore machine on F consists of an equivariant output map $o : F \to O$ and, for each input $a \in A_0$, an equivariant transition $d_a : F \to F$. Since monodromy interchanges the sheets, equivariance forces

$$o(\tau s) = o(s).$$

Hence o assigns the *same output* to both sheets. Each transition d_a is either the identity or the swap τ. The underlying state dynamics may be nontrivial, but constant-set observations cannot distinguish the sheets. For every input word the resulting observation remains the same. In this example the two locally distinct states are behaviourally indistinguishable to a constant-output Moore observer. The output behaviour of each local sheet is the same constant stream (or constant A_0^*-indexed output function), so the behaviour map from F factors through the terminal sheaf 1. Hence the minimal constant-interface Moore realisation is the one-state machine. This is a simple example of *global observable behaviour without a global chosen state*: the state sheaf is unpointed, while the behaviour map $F \to O^{I^*}$ picks out a global behaviour through its factorisation $F \to 1 \to O^{I^*}$. There is a richer option: let the output interface vary spatially. Taking $O = F$ and $o = 1_F$, with each d_a either 1_F or τ, gives a perfectly valid internal Moore machine. Now the output is itself a locally varying section; it can distinguish the two local sheets, even though there is no global output section choosing one of them. The point is not that the output has become a global element, but that observation is a morphism of sheaves. This illustrates why variable interfaces matter for our DOTS perspective. More generally, on a connected, locally simply connected space, a locally constant state sheaf with fibre S_0 and monodromy representation $\pi_1(X) \to \mathrm{Aut}(S_0)$ supports Moore operations exactly when those operations are equivariant for monodromy (with the corresponding equivariance condition on any varying input and output fibres). The double cover is the smallest nontrivial test.

7.4 Probabilistic transition and the distinction from distributions

The same monodromy test applies to an internal rand-coalgebra $F \to \mathsf{rand}(F)$. On the two-element fibre, an equivariant polynomial lottery ρ must satisfy

$$\rho(\tau s) = \mathsf{rand}(\tau)(\rho(s)).$$

For example, setting $\rho(s) = (\frac{1}{2}, \frac{1}{2}; s, \tau s)$ obeys this equation. It defines a global *transition operation on the sheaf* despite the absence of a global ordered lottery of states. This distinction is essential: a globally defined map $F \to \mathsf{rand}(F)$ need not determine a global section of $\mathsf{rand}(F)$. The monodromy obstruction from the preceding section concerns a global *state-independent* ordered lottery of sheet sections. The equivariant, state-dependent operation just displayed does exist. In contrast, the symmetric distribution $\frac{1}{2}\delta_s + \frac{1}{2}\delta_{\tau s}$ is independent of the chosen starting sheet and gives a global distribution section. In the polynomial setting its ordered counterpart changes when the starting sheet changes, but an equivariant state-dependent runner can still keep track of that change.

8 Safety and costs under descent

There are two notions of locality to distinguish. For the transitive two-sheet local system F, its only *locally constant* subsheaves are $\varnothing$ and F, since the two-element fibre has only trivial monodromy-invariant subsets. Nevertheless, general subsheaves are plentiful: via the étalé-space correspondence they are the open subsets $W \subseteq \widetilde{X}$. If $U \subsetneq S^1$ is a proper nonempty open arc, then $W = \pi^{-1}(U)$ determines a proper nonempty safety subsheaf $P_W \hookrightarrow F$ invariant under both 1_F and the deck involution τ. It says that a state is safe while situated over U. Such a predicate is *spatially localised*, rather than a global choice of a named sheet. More generally, locally constant safety predicates in a local system are unions of monodromy orbits. The choice of cost object is equally significant. Every sheaf morphism $F \to \underline{\mathbb{N}}$ is constant, because the covering space $\widetilde{X}$ is connected and the target is discrete. The same applies to a constant discrete rational cost object. By contrast, maps from F to the internal Dedekind reals correspond to continuous real-valued functions on $\widetilde{X}$; these may distinguish the sheets locally. Below we deliberately use the *constant, discretely ordered rational* cost object, compatible with our rational lottery weights. We make no claim here about a real-valued expectation monad. To obtain a nontrivial *locally constant* invariant predicate and a potential that actually changes, take

$$S = F \times \underline{K}, \qquad K = \{0, 1, 2\}, \qquad P = F \times \{0\} \hookrightarrow S.$$

The K-coordinate is a three-level resource or risk counter, while the monodromy swaps the sheet coordinate only. Its orbits on the six-element fibre $\{0, 1\} \times K$ are precisely the levels $j \in K$. Hence every locally constant safety predicate and every locally constant cost or potential on S must factor through K: the choice to make P, c and Φ depend only on j is forced, not a loss of generality. The double cover contributes the state-relative "stay/switch naming of lottery tickets. Define a state-dependent two-ticket polynomial lottery by

$$\rho(s, j) = \big(\tfrac{1}{2}, \tfrac{1}{2};\ (s, j), (\tau s, \max(j - 1, 0))\big).$$

It is equivariant: changing the input sheet by τ changes each successor sheet by τ, without altering ticket positions or weights. The safe subsheaf P is invariant, since both successors of a state with $j = 0$ again have level zero. Put $\Phi(s, j) = j$ and use the rational per-step cost

$$c(s, j) = \begin{cases} \frac{1}{2}, & j > 0, \\ 0, & j = 0. \end{cases}$$

Then in the constant ordered object $\underline{\mathbb{Q}_{\geq 0}}$ one has the exact expected-potential equation

$$c(s, j) + \tfrac{1}{2}\Phi(s, j) + \tfrac{1}{2}\Phi(\tau s, \max(j - 1, 0)) = \Phi(s, j).$$

For $j > 0$ the potential falls by one on the switch ticket and stays unchanged on the stay ticket; at $j = 0$ both the cost and potential vanish. Thus the amortised allowance is *zero*: the decrease in expected potential pays for the immediate cost. This is a concrete localic potential certificate rather than a static bound. All maps descend because they are equivariant for monodromy. The distinction between ticket *position* and an *absolute sheet* remains crucial. The two positions in ρ name outcomes relative to the current state ("stay" and "switch"); they do not pick out a globally named sheet. A wider descent theorem for safety predicates, costs and probabilities on arbitrary locales remains open.

9 How these examples fit the Sweedler picture

The deterministic, counted, and Moore machines give a progressive tutorial dictionary. A $\{N, y\}$-coalgebra is a sheaf with one endomorphism; a $\{N, N\}$-coalgebra is a sheaf with an endomorphism and an iteration count; a $\{T_{\mathrm{Moore}}, y\}$-coalgebra is a sheaf with observations and input-indexed dynamics. The cospan makes spatial compatibility a pair of commuting squares, whereas the circle makes it an equivariance condition against monodromy. The lottery case adds indexed probabilistic branching and distinguishes a morphism of probabilistic machines from a global lottery of states. These are operationally meaningful statements even without a completely unrestricted measuring-comonad theorem. The latter requires the enlarged accessible enriched-endofunctor universe, the biclosed convolution and interchange, and comparison of accessibility ranks. Those foundations should be referenced here rather than allowed to obscure the first examples.

10 What is proved and what remains

The explicit sheaf-machine descriptions, the internal strong polynomial lottery construction, its pullback preservation, and the monodromy calculations do not require an unrestricted general measuring theorem. The finite-localic enriched arity calculations provide tractable tests for the proposed duoidal universe. For tutorial purposes, it suffices to state the Sweedler construction under the general duoidal hypotheses and to retain the rank-comparison issue as a technical open question. Important future problems are: (i) simplify sufficient arity hypotheses for a uniform enriched Sweedler theorem; (ii) compare fixed-rank measuring objects; (iii) extend the descent analysis to arbitrary locales and sheaf-valued probabilistic safety and cost structures; and (iv) compare ordered polynomial, symmetric analytic, and groupoid-valued ticket semantics and their effects on pullbacks, measuring and safety certificates.

References

[1] R. Garner and I. López Franco, *Commutativity*, Journal of Pure and Applied Algebra 220 (2016), 1707–1751.

[2] I. López Franco and C. Vasilakopoulou, *Duoidal categories, measuring comonoids and enrichment*, arXiv:2005.01340.

[3] R. B. B. Lucyshyn-Wright, *Enriched algebraic theories and monads for a system of arities*, Theory and Applications of Categories 31 (2016), 101–137.

[4] R. Garner, *The costructure–cosemantics adjunction for comodels for computational effects*, Mathematical Structures in Computer Science (2022); arXiv:2011.14520, §9.3, Definition 104 and Proposition 105 in arXiv version v1 (30 November 2020; sheaf comonad and specialisation category).