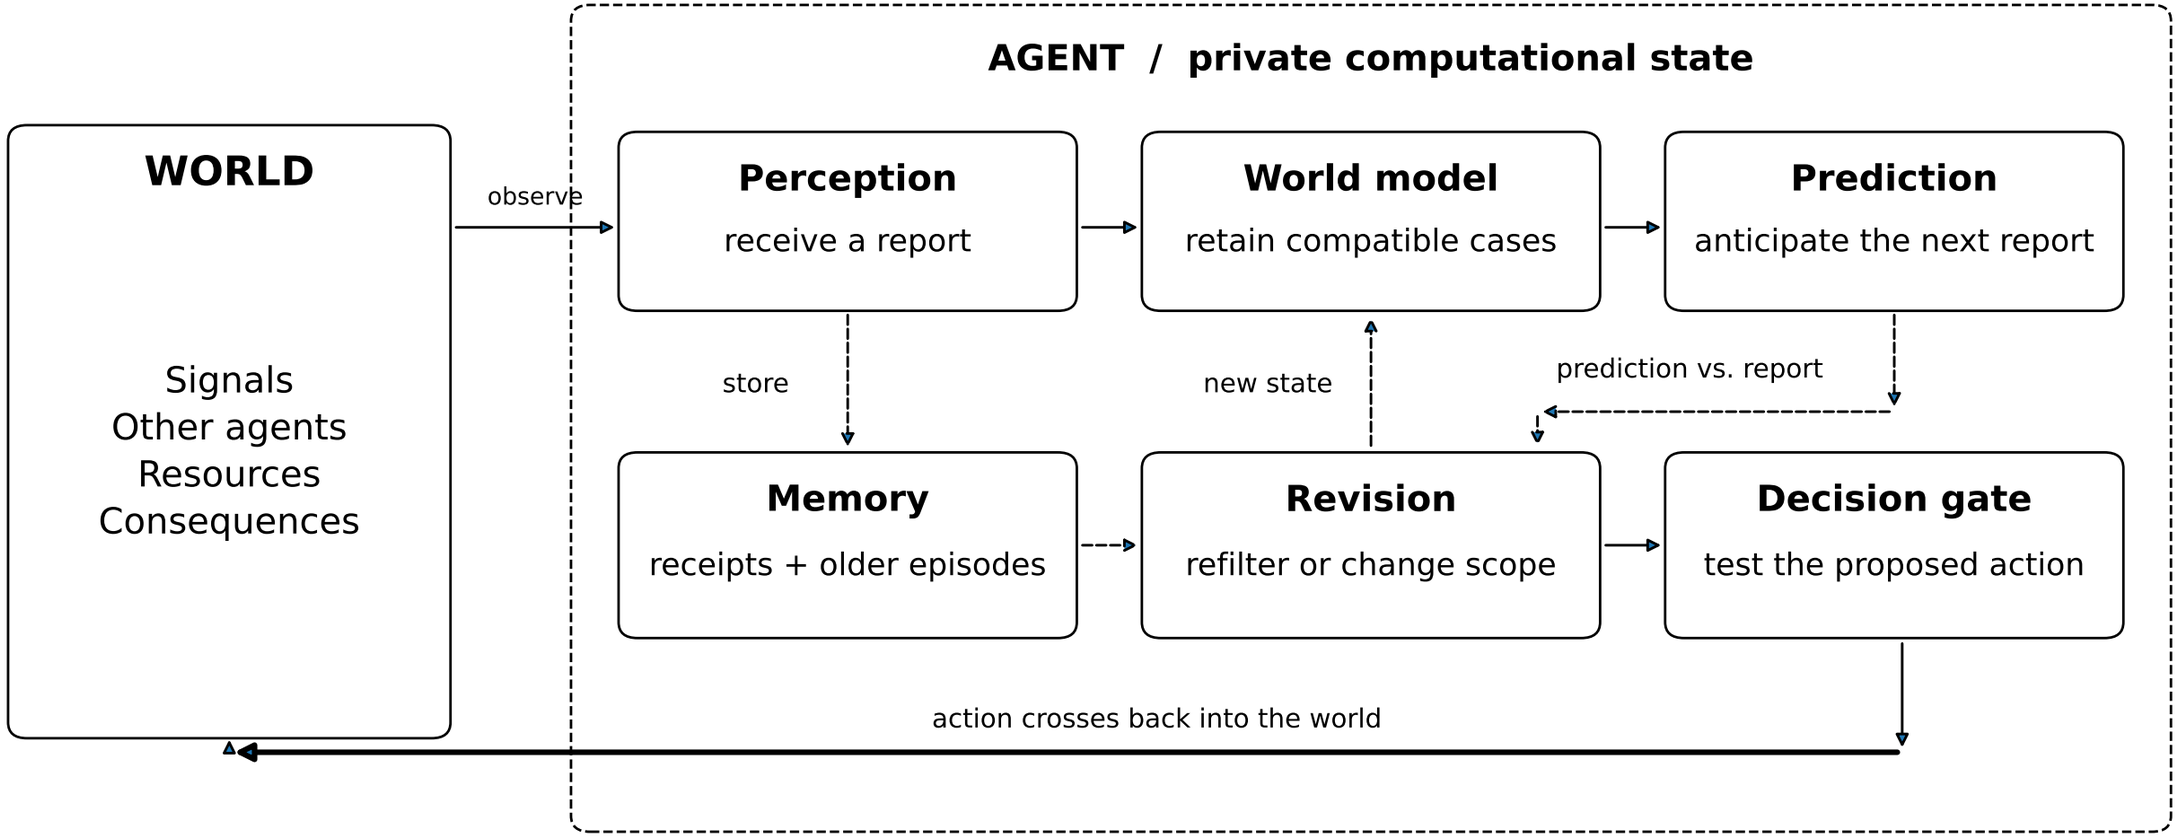
An agent acts across a boundary

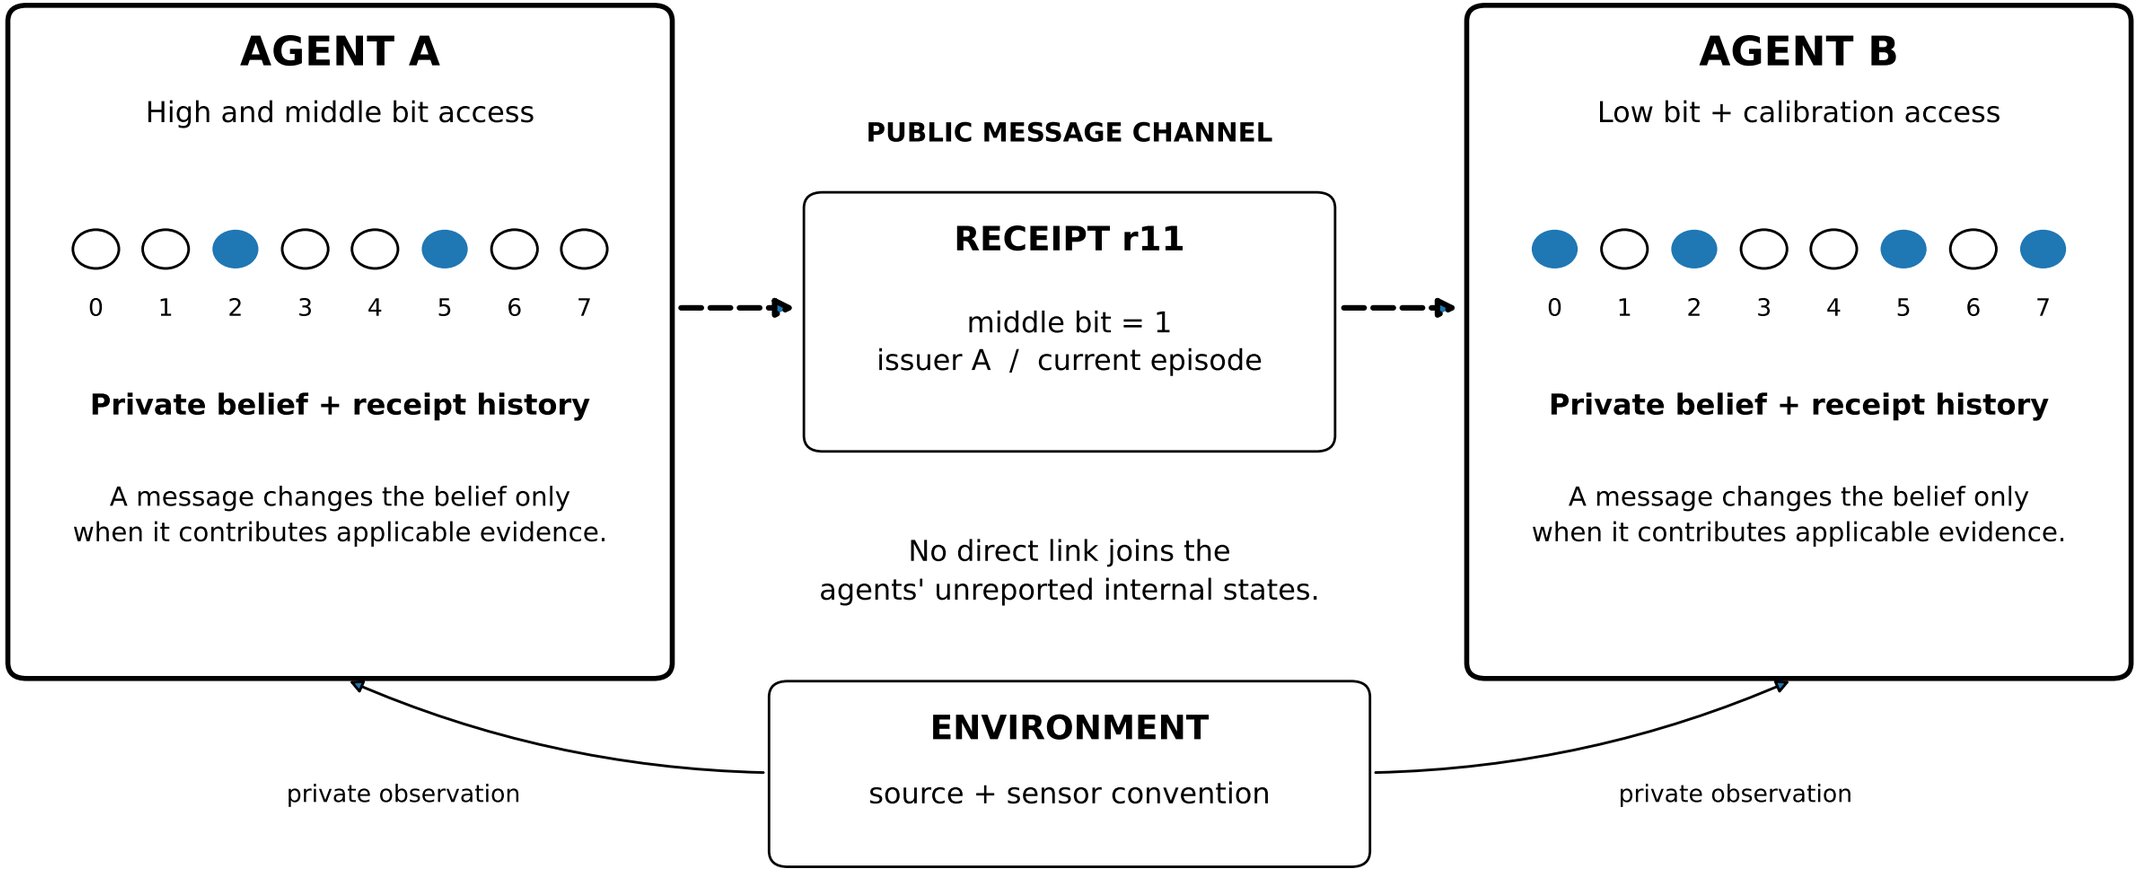
Outside: observations and consequences. Inside: retained state, predictions, comparisons, and decisions.



Solid arrows carry observations or actions. Dashed arrows carry internal revision or remembered evidence. These are functional roles, not measured neural modules. The toy implementation makes belief, memory, and action checks explicit.

Two agents do not share a mind

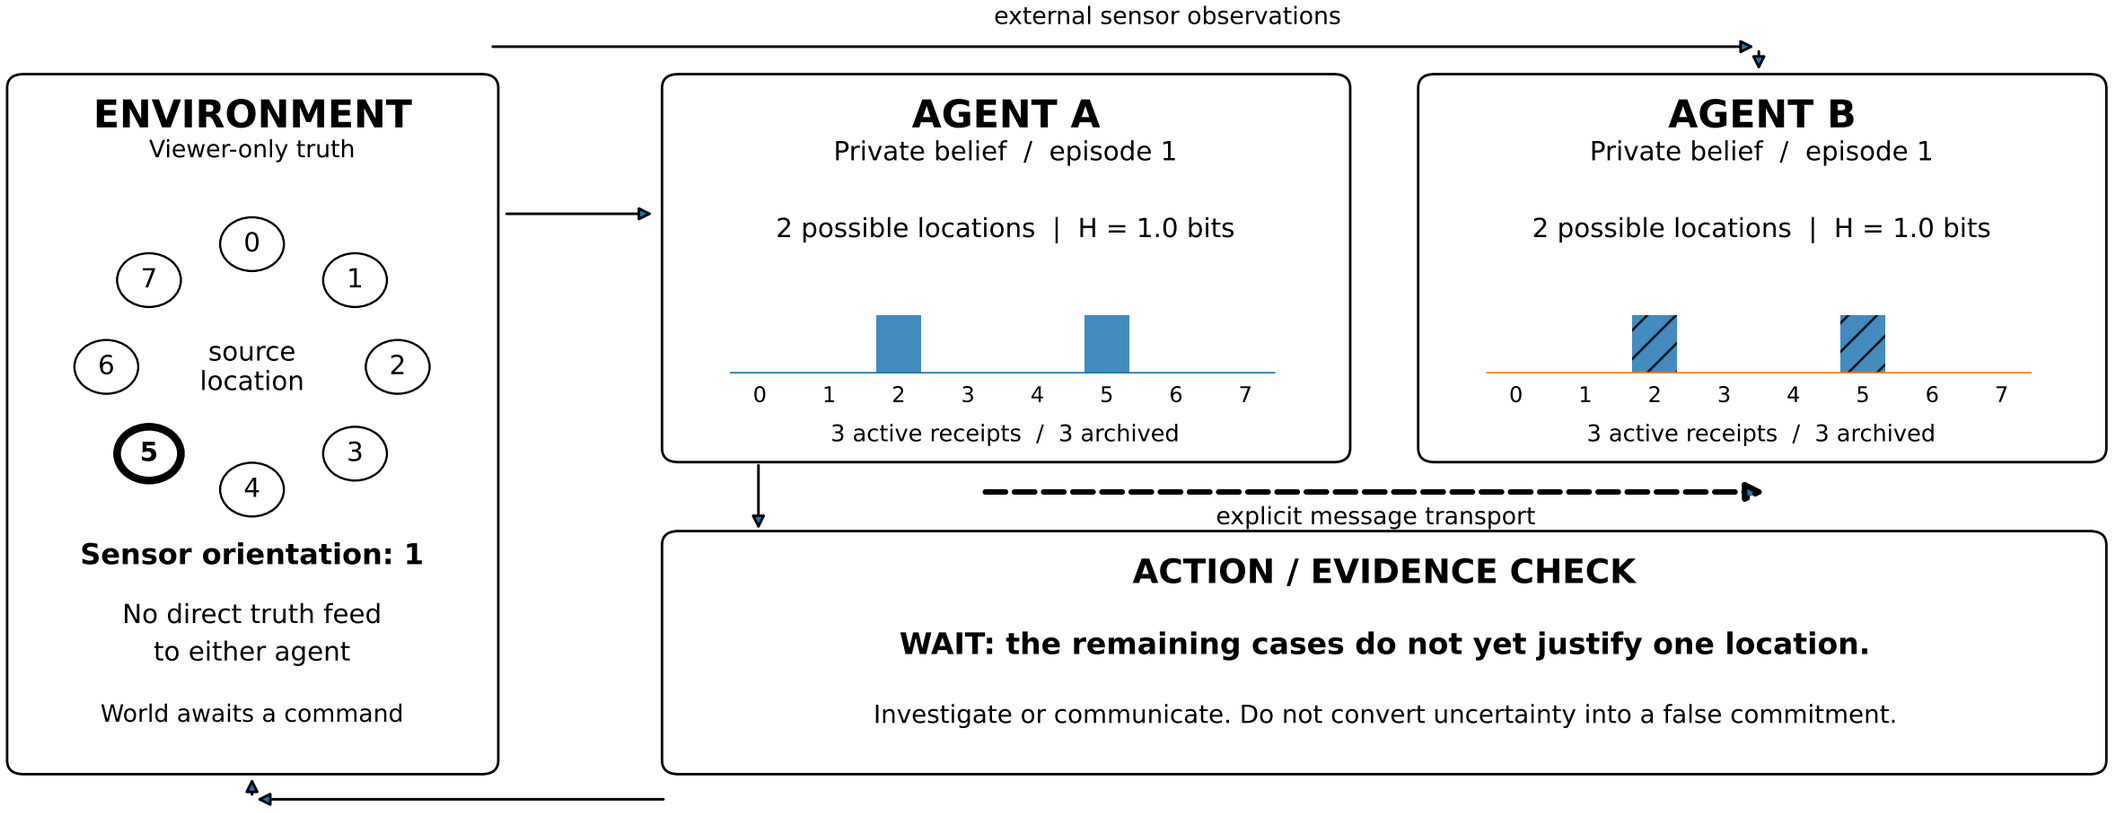
Distinct evidence produces distinct private states. Messages can bring those states into agreement.



At event 11, A retains locations {2, 5}; B still retains {0, 2, 5, 7}. A must send its middle-bit receipt before B can make the same update. Reliable transport and truthful reports are explicit assumptions, not achievements inferred from the diagram.

A live view of internal and external state

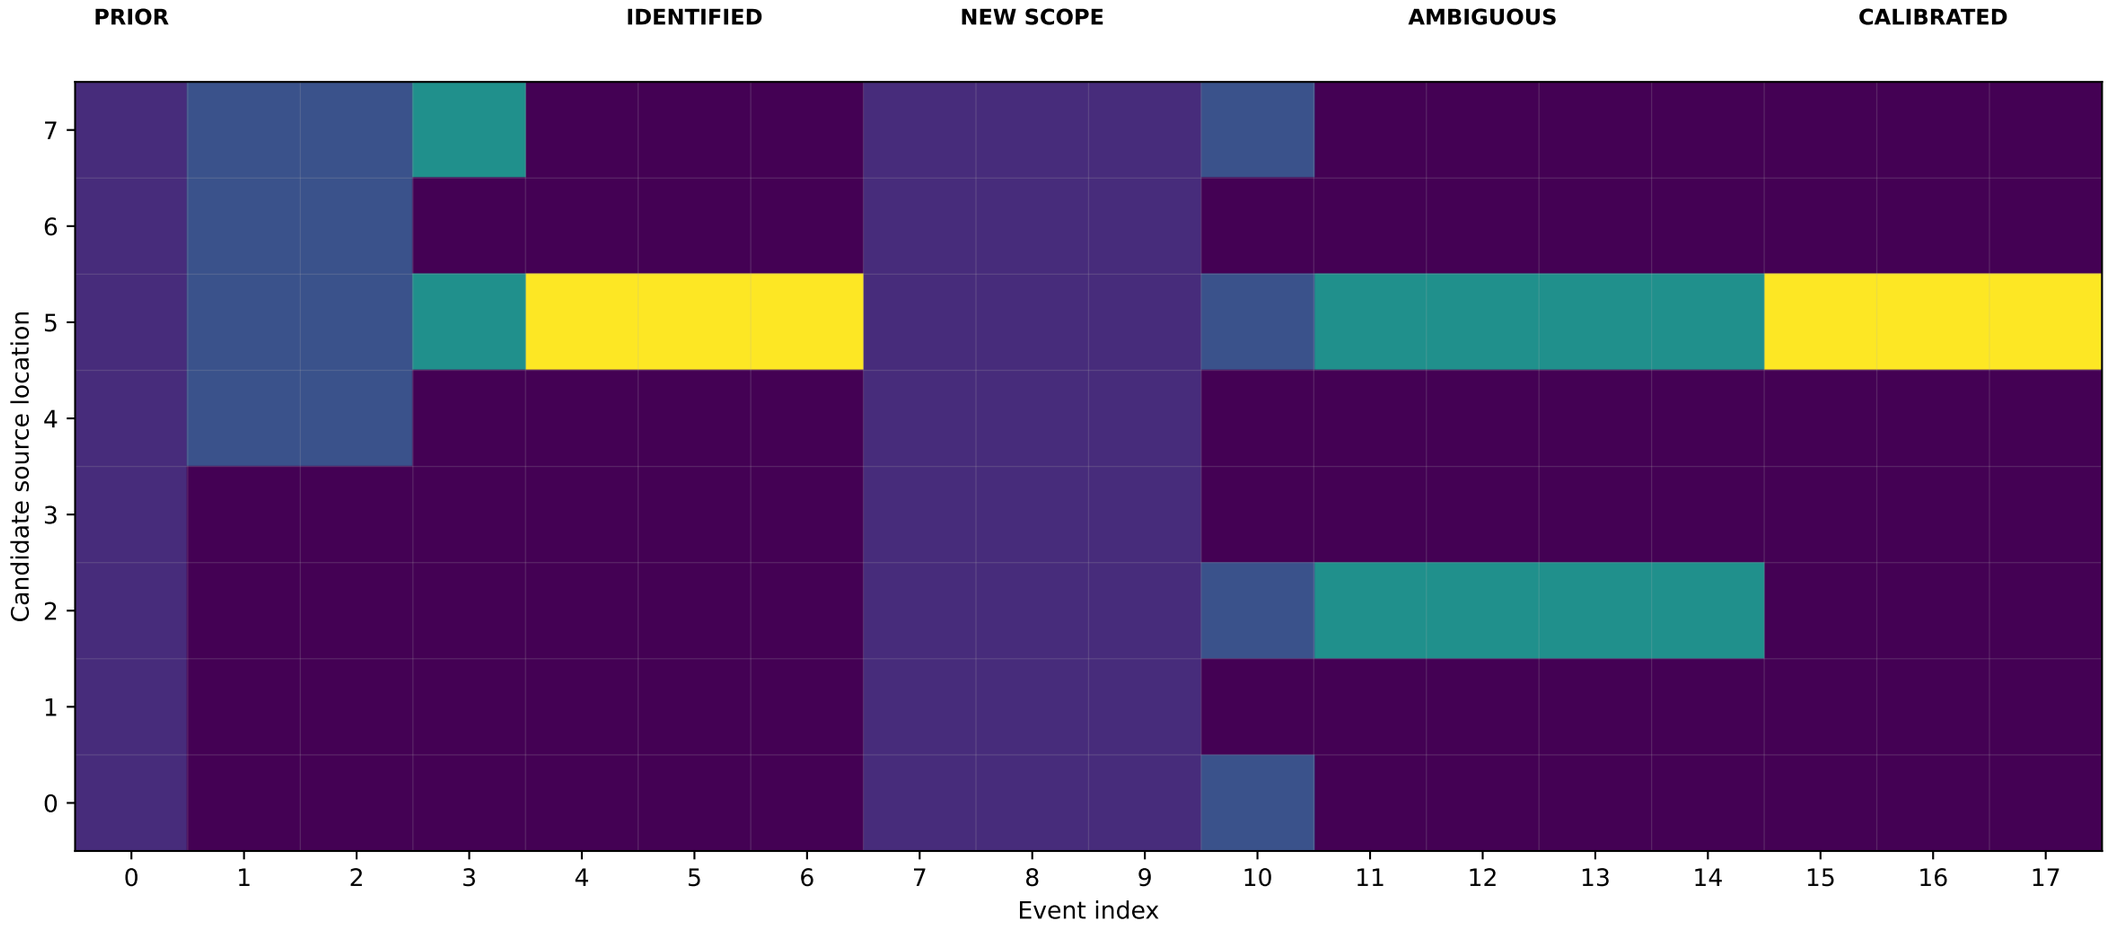
Event 12: both agents agree about the uncertainty, but neither can yet select the right action.



The animation uses this same layout at all 18 events. Histogram heights are computed probabilities, not decorative activations. The outlined world location is shown to the reader only; the agents receive bit reports and explicit messages.

The representation changes over time

A's complete belief over source location, not a single confidence score.

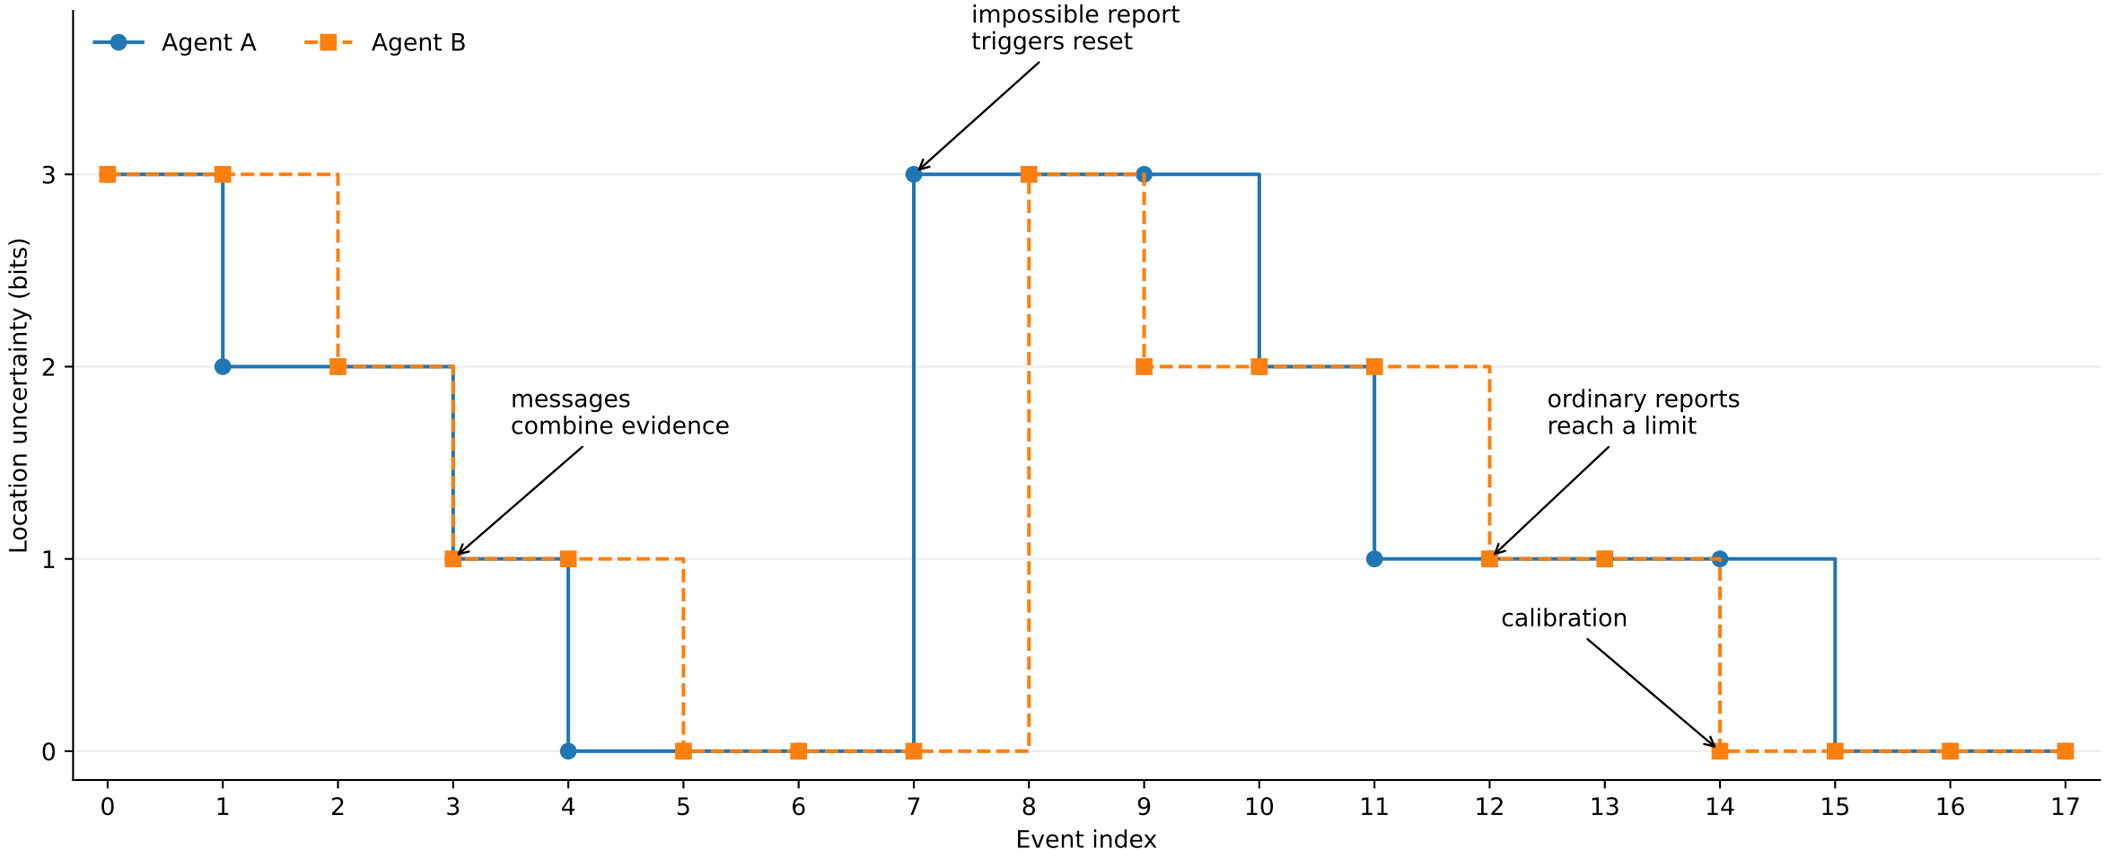


Each column sums to one. Initial observations narrow the support. At event 7, a contradiction triggers a supplied model-reset rule and uncertainty expands. Calibration later resolves what ordinary readings cannot. Intensity is probability from 0 to 1.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

Uncertainty contracts, expands, then resolves

Private state can diverge temporarily even when both agents use the same update rule.

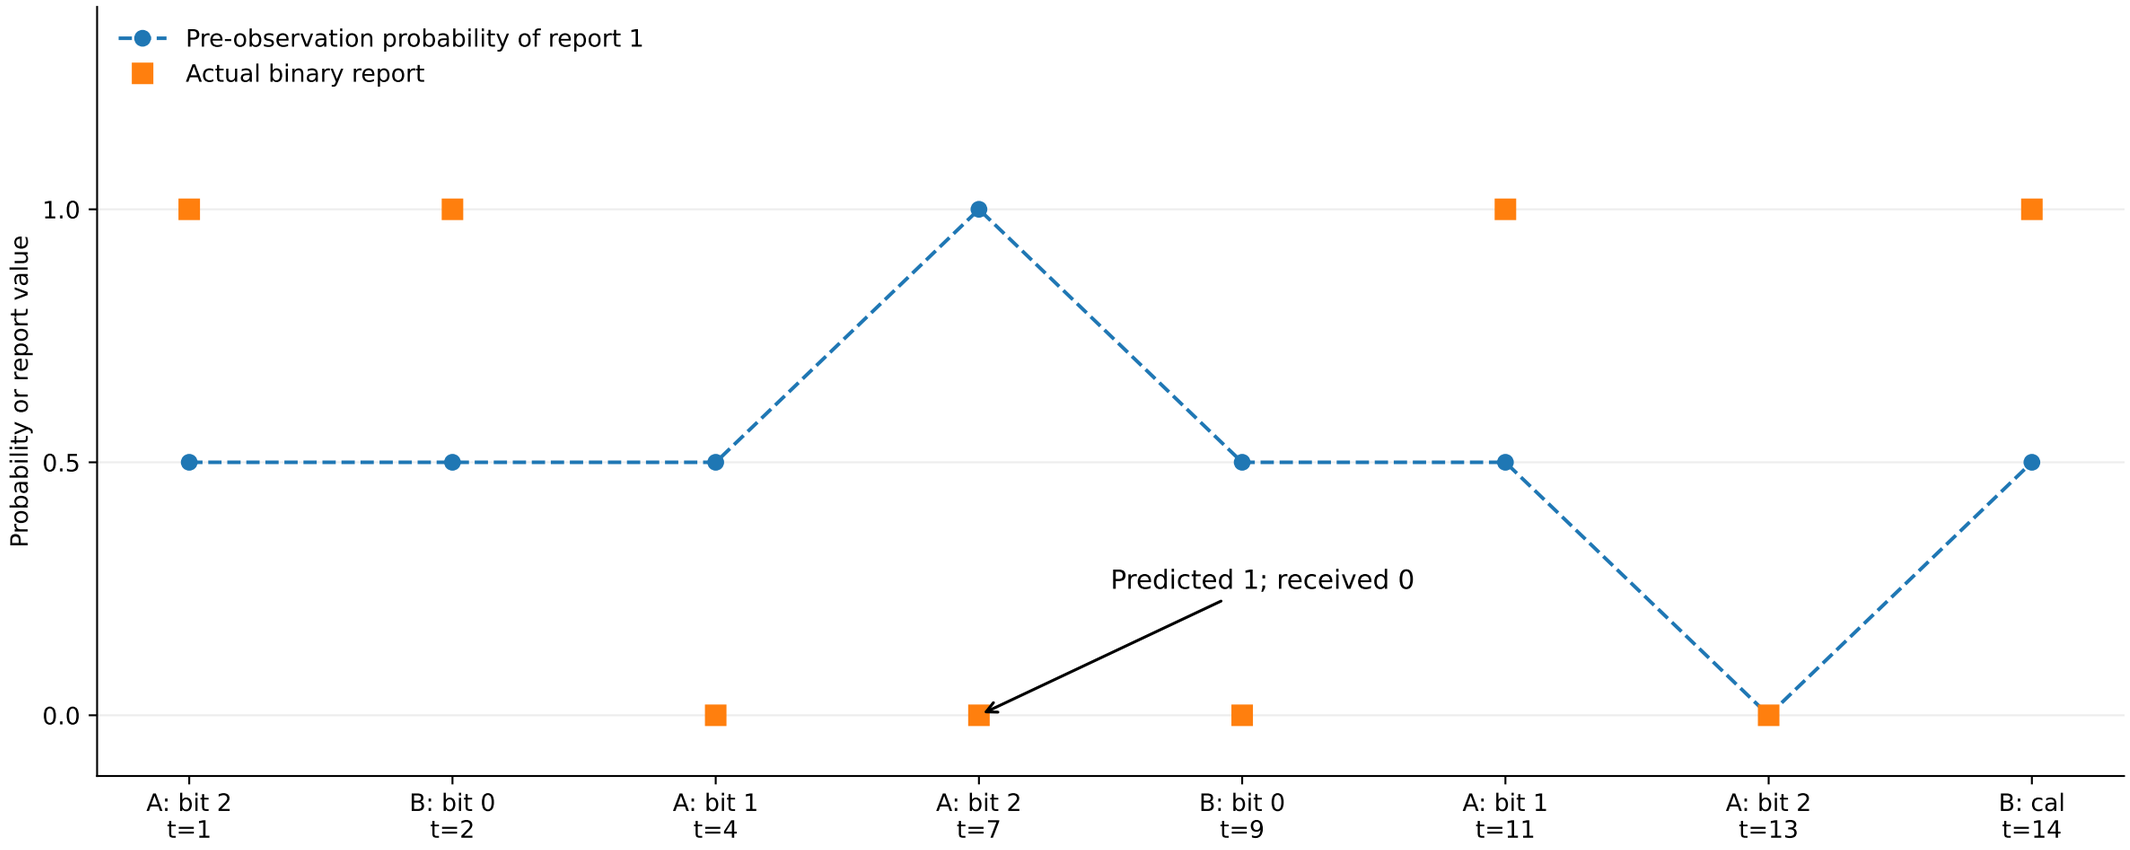


Entropy is calculated from each exact location distribution: H = -sum p log2 p. The rise at events 7-8 is the supplied response to a model contradiction, not evidence that more data inherently makes Bayesian inference worse. The stale belief at event 6 illustrates a missing premise.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

A prediction error is an internal event

The sensor supplies a report. The agent compares it with a prediction made before receiving it.

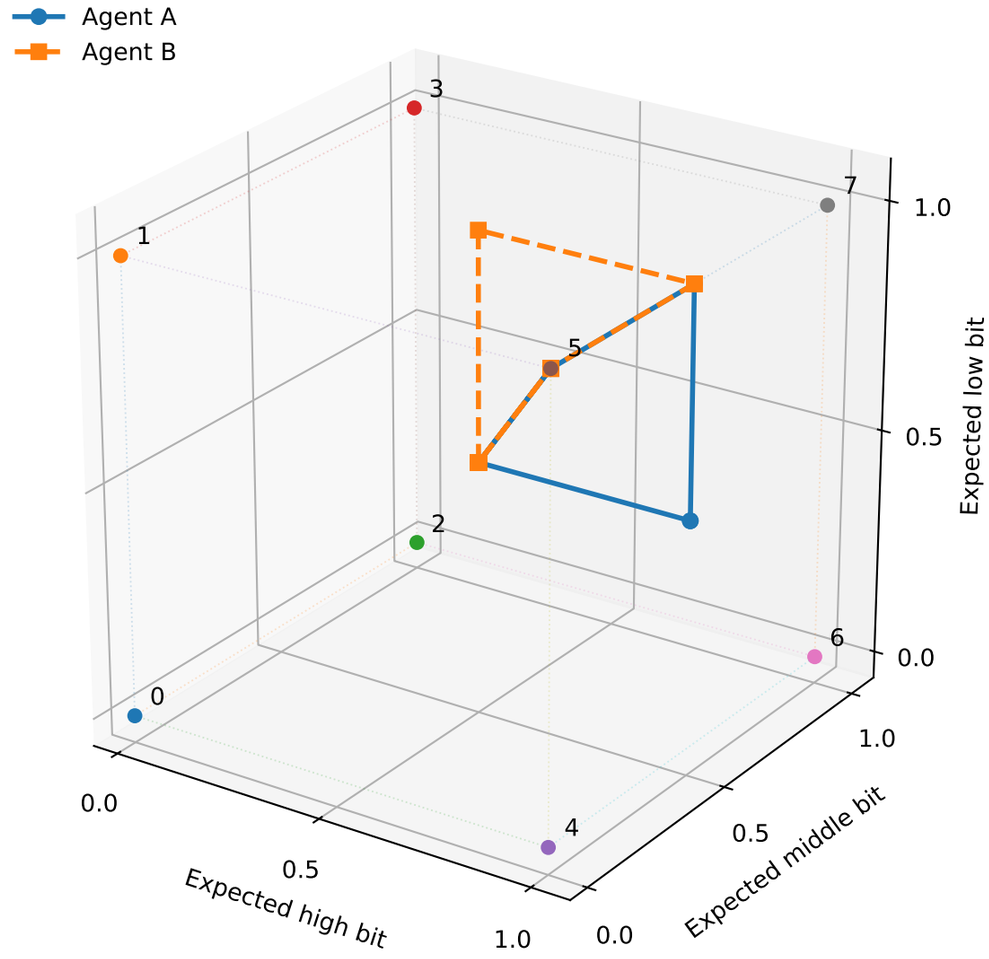


At event 7, A predicts report 1 with probability one and receives 0. That is a contradiction within this noiseless finite model. It triggers model expansion, not a conclusion that one particular physical cause has been identified. Calibration is a different observation channel.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

An actual three-dimensional state trajectory

The coordinates are expected source bits, calculated from each agent's probability distribution.



READ THE GEOMETRY

Corners = definite locations.

Interior points = belief centroids.

The paths move as evidence arrives or the model resets.

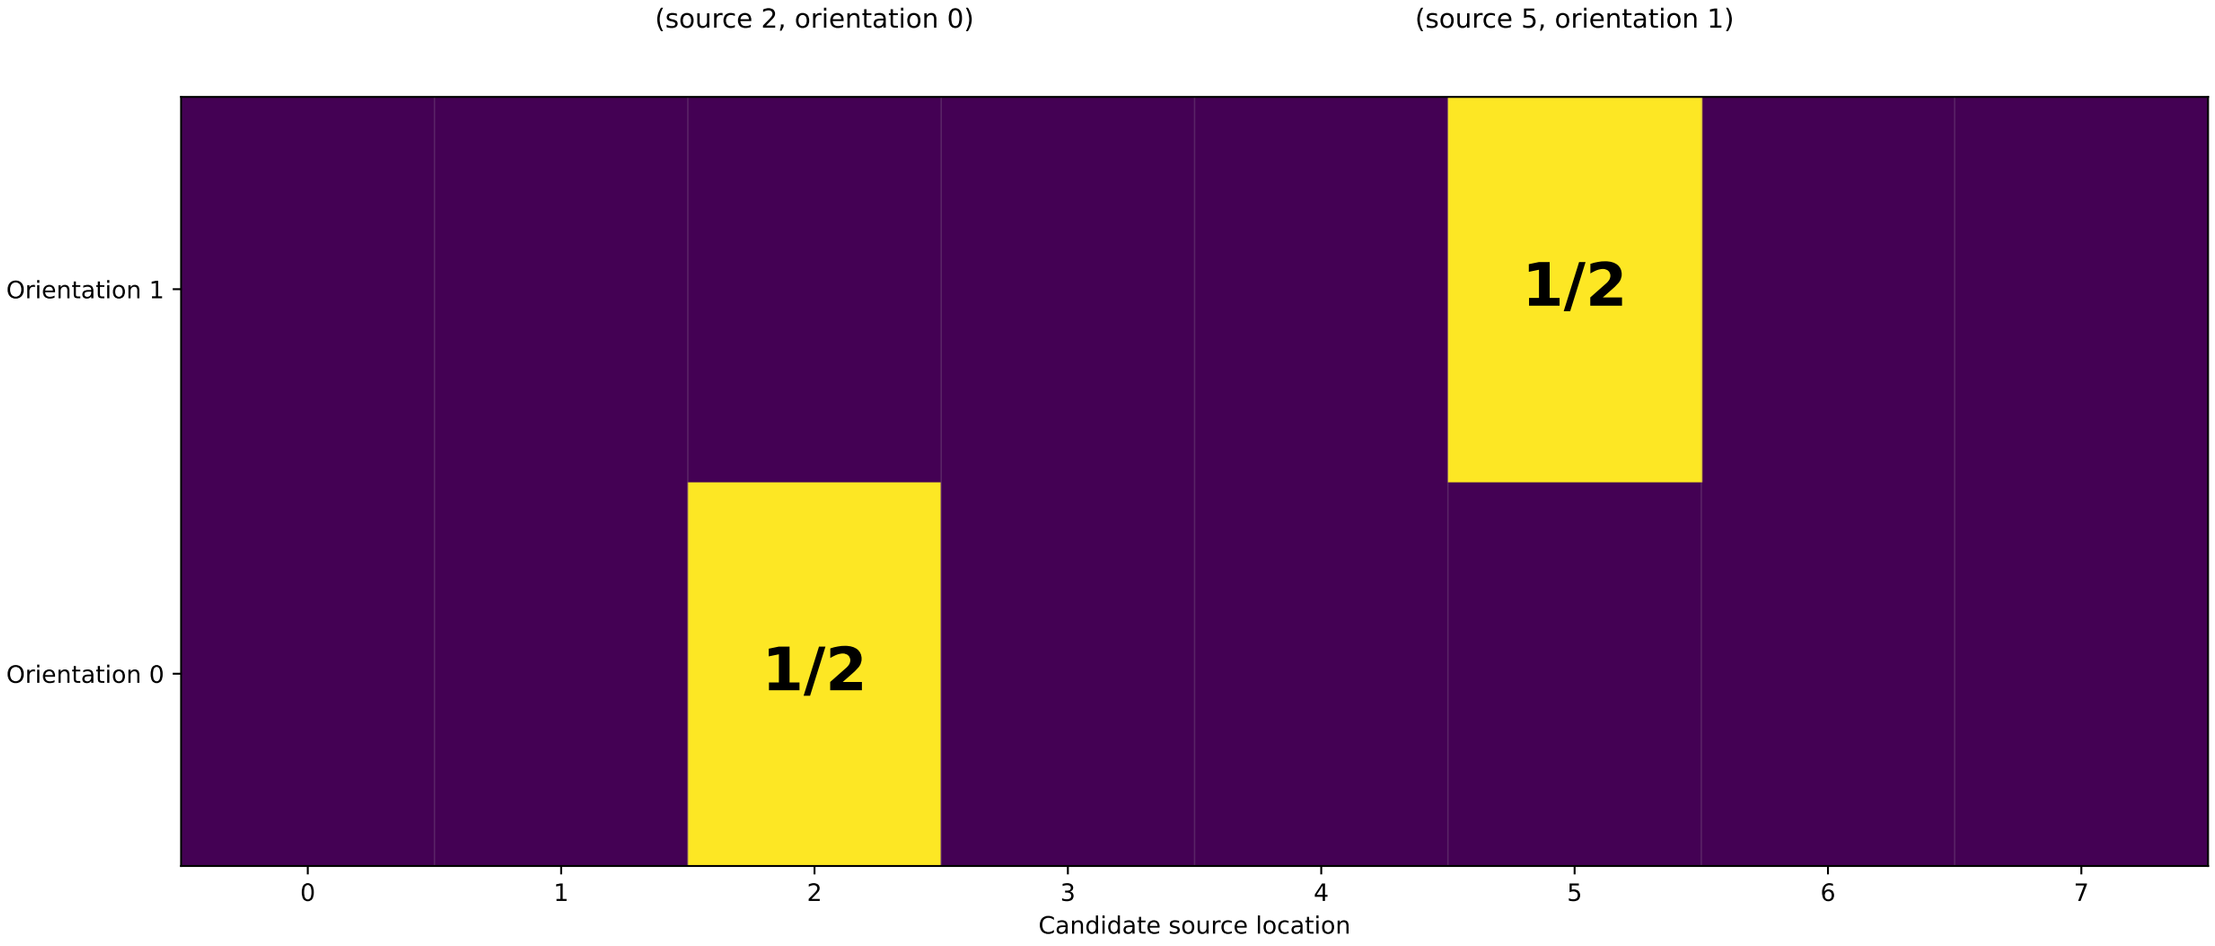
A shared point need not mean a shared full representation.

This is a genuine 3D projection of the explicit belief state, not an invented neural embedding. It is lossy: the uniform prior and the later 50/50 belief over locations 2 and 5 have the same centroid. The full belief heatmap distinguishes them.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

Two worlds produce the same ordinary reports

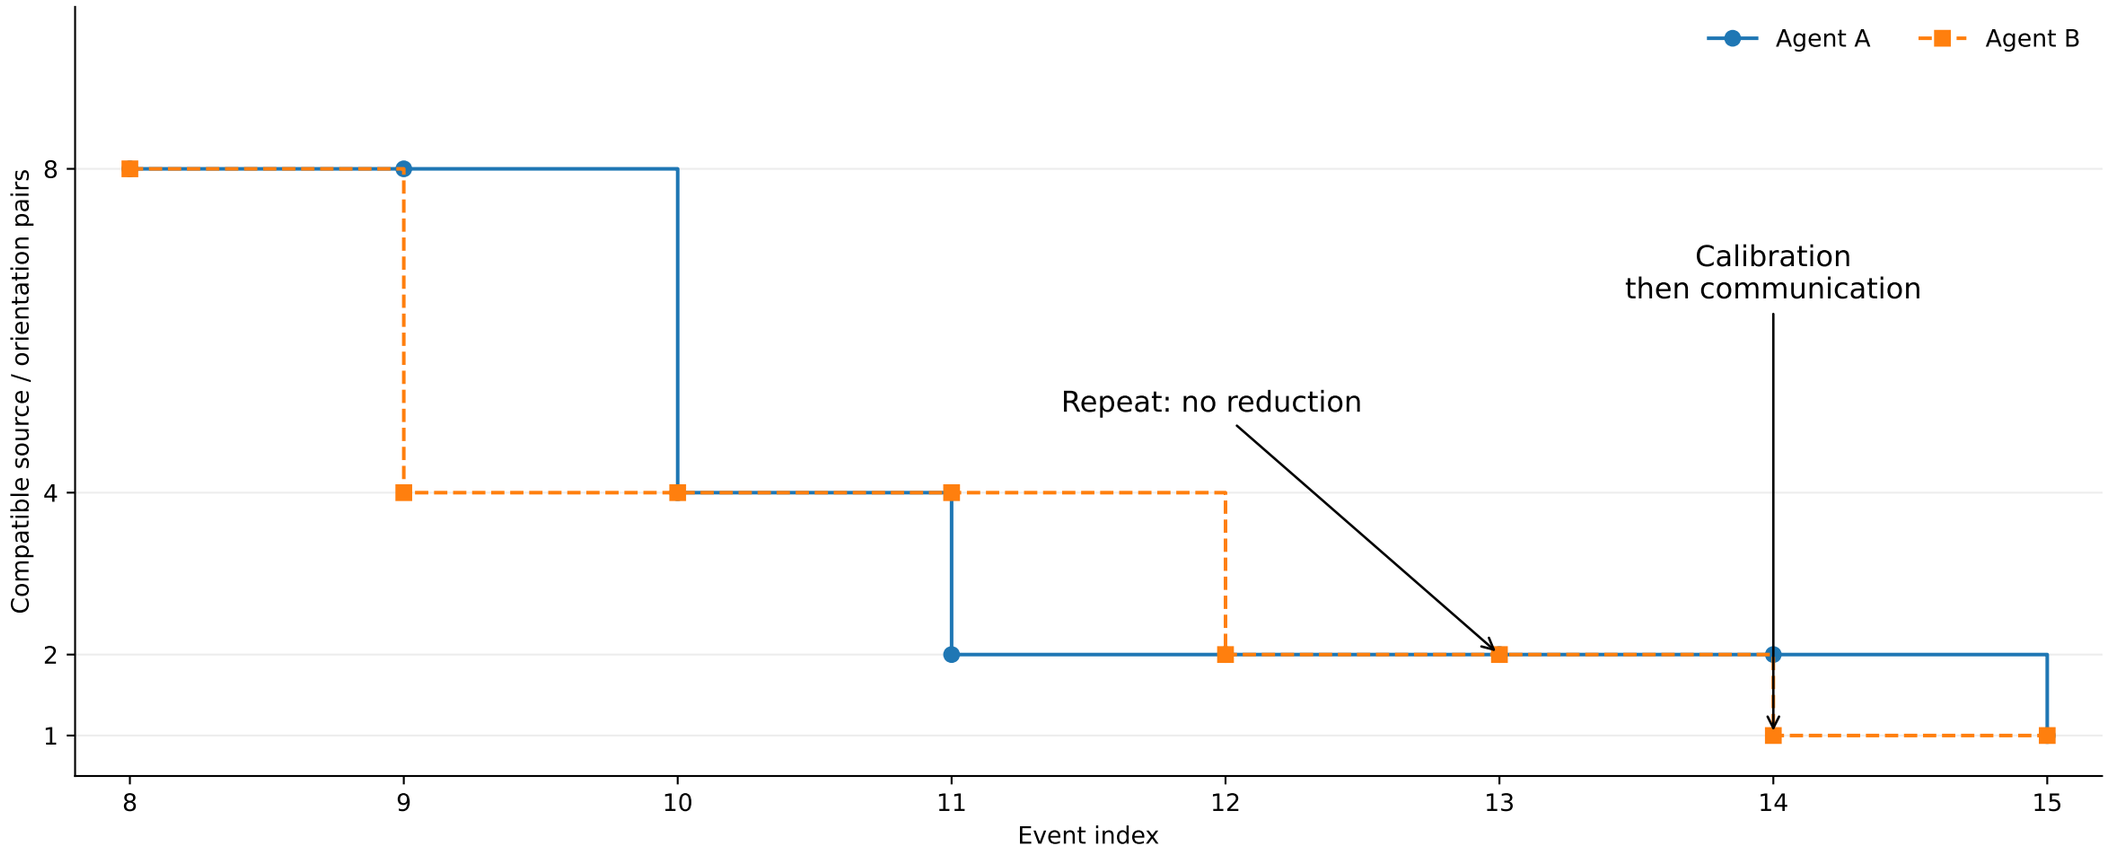
At event 12, both explanations fit every reported source bit.



Location 2 has ordinary bit pattern 010 with orientation 0. Location 5 has bit pattern 101 with orientation 1, producing the same reported 010. Both posterior entries are exactly 1/2. A repeated ordinary reading cannot distinguish them.

More of the same is not a new distinction

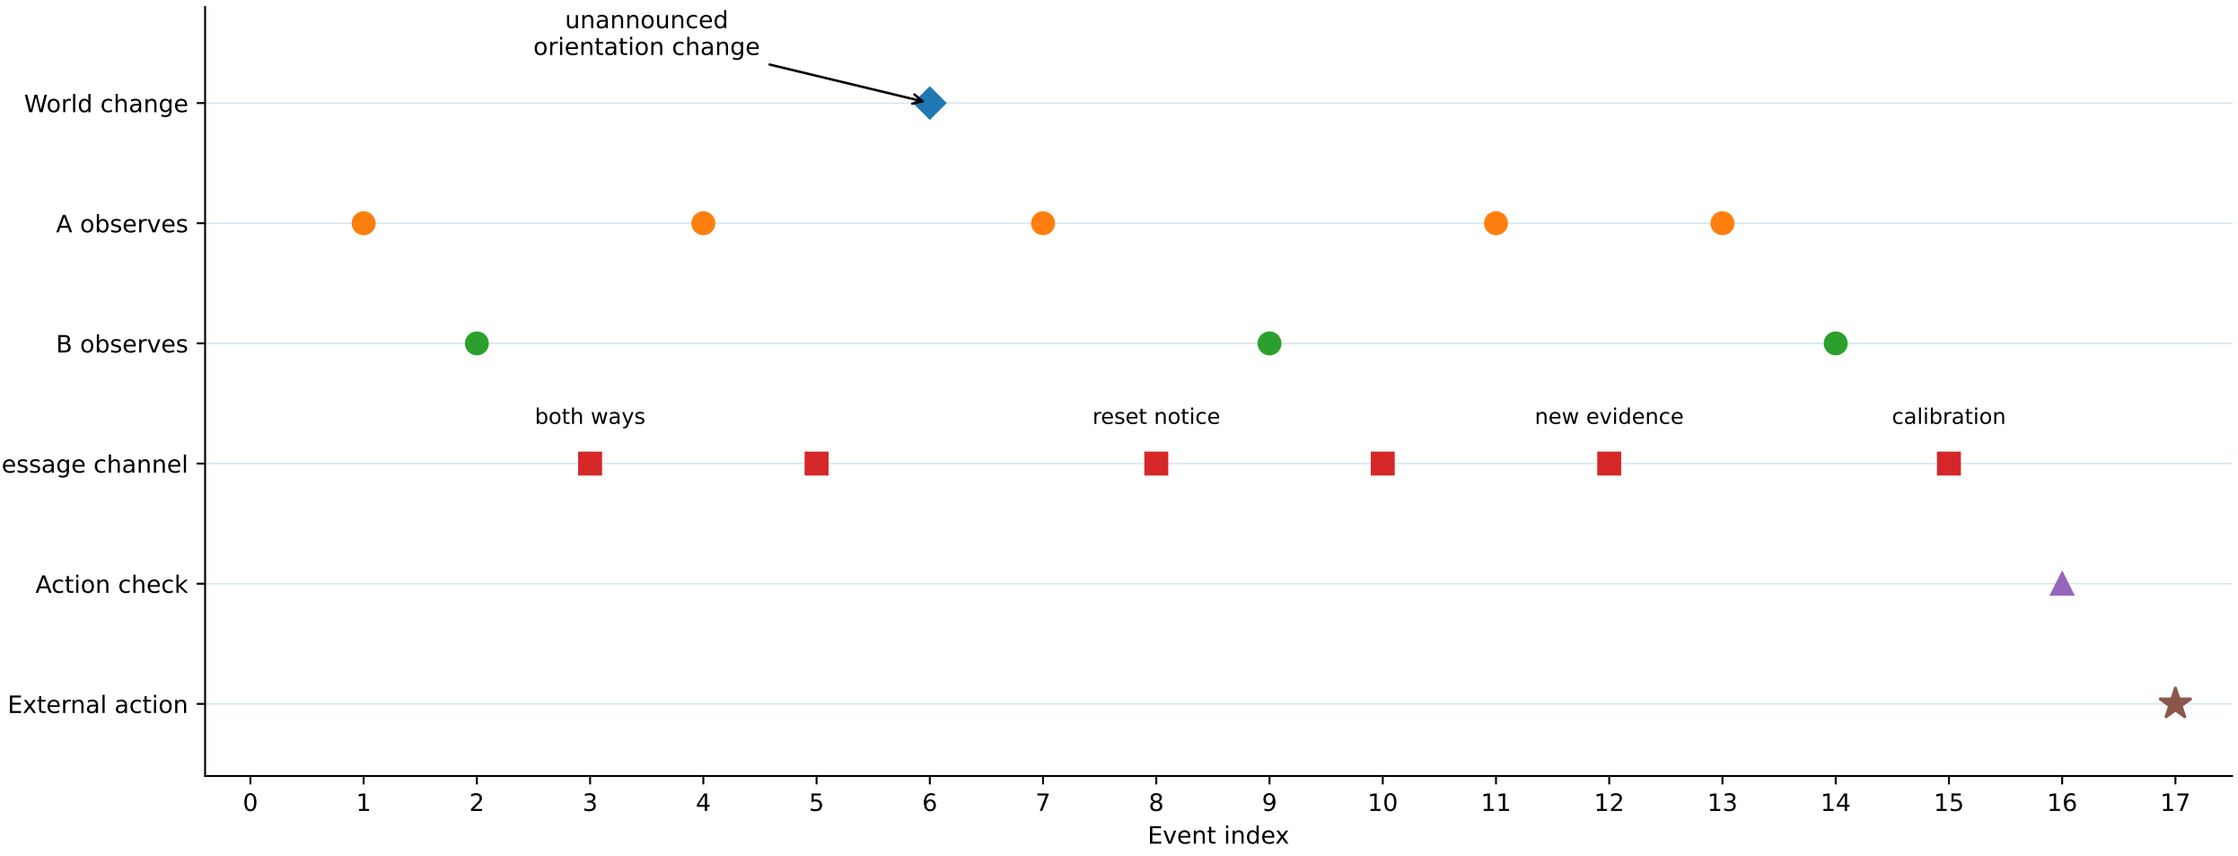
A new evidence channel resolves a problem that repeated ordinary observations leave unchanged.



These are joint (source, orientation) hypotheses after the reset. The repeated high-bit reading at event 13 leaves both possibilities intact. The calibration at event 14 reduces B to one case, and its message lets A make the same update at event 15.

External events and internal reactions are interleaved

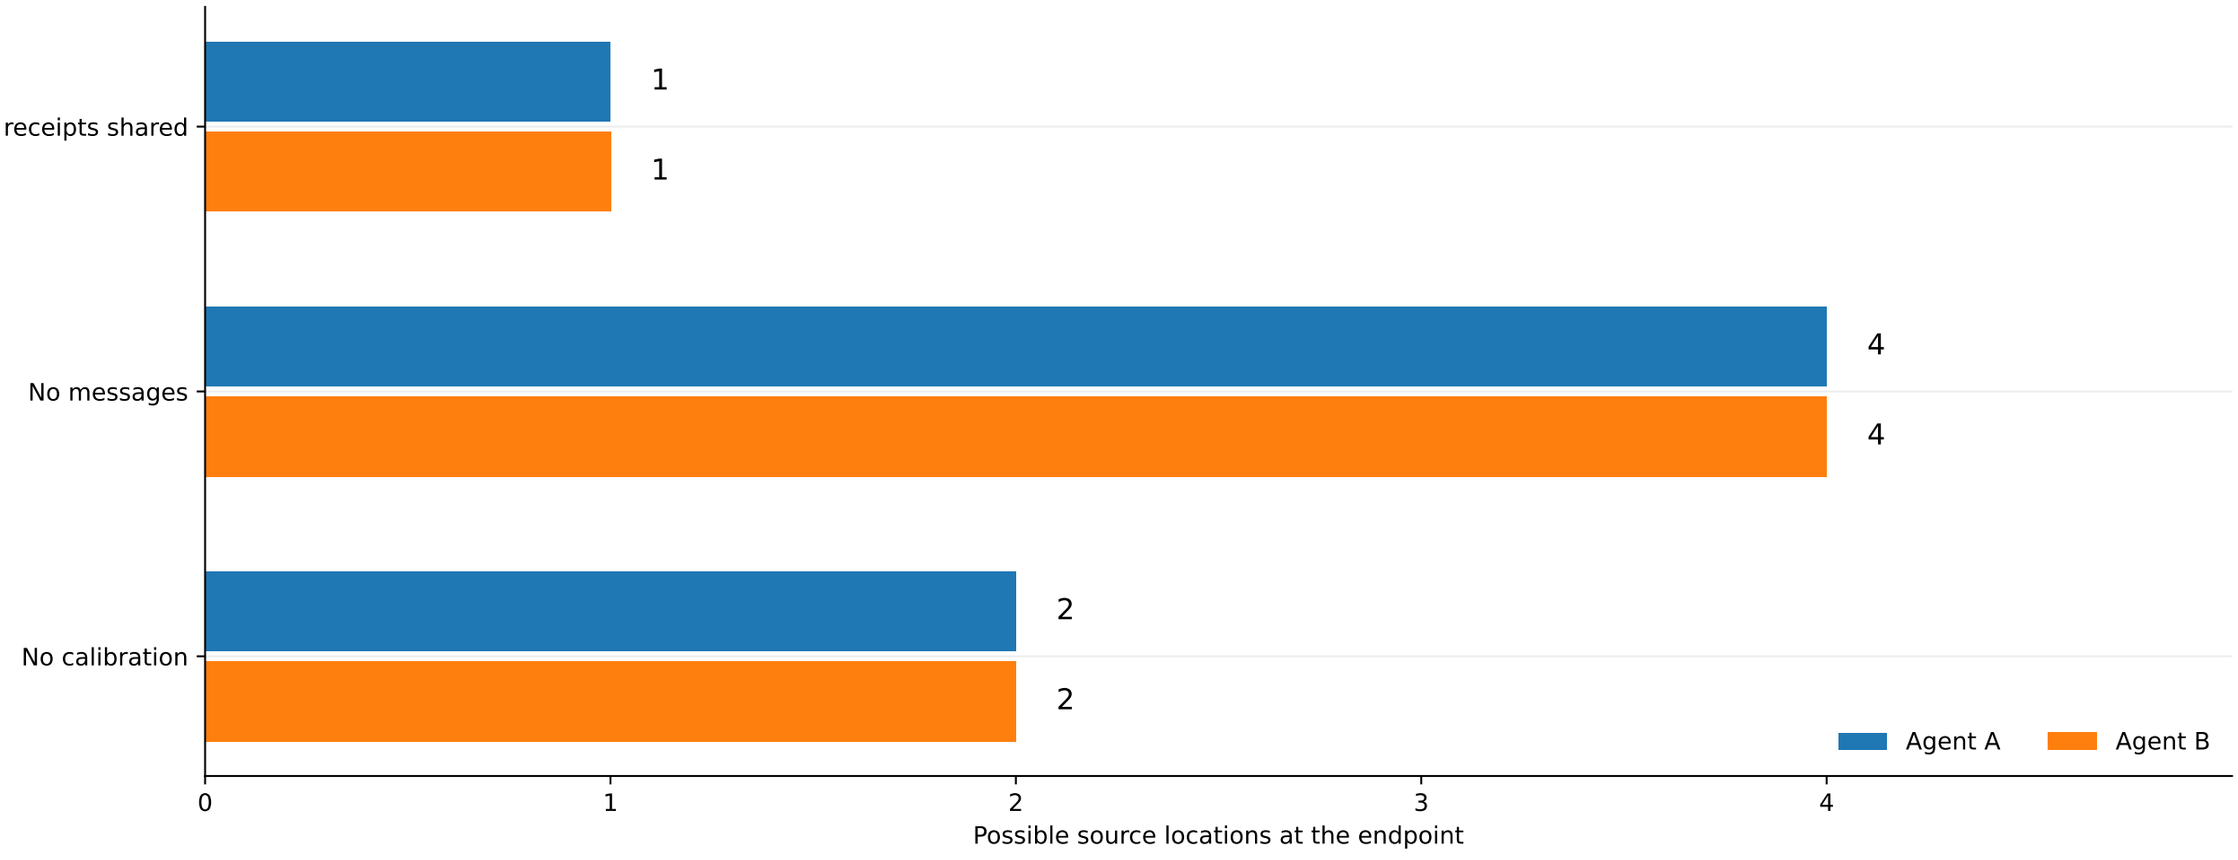
An event raster reveals who observes, who communicates, and when an action actually reaches the world.



Circles are scheduled observable operations, squares are messages, and the final star is the external service command. Event 6 changes the world without directly changing either belief. The next report is what causes A to revise. Time here is event order, not measured seconds.

Remove a channel; the internal endpoint changes

Controlled access comparison using the same post-change source and the same available private observations.

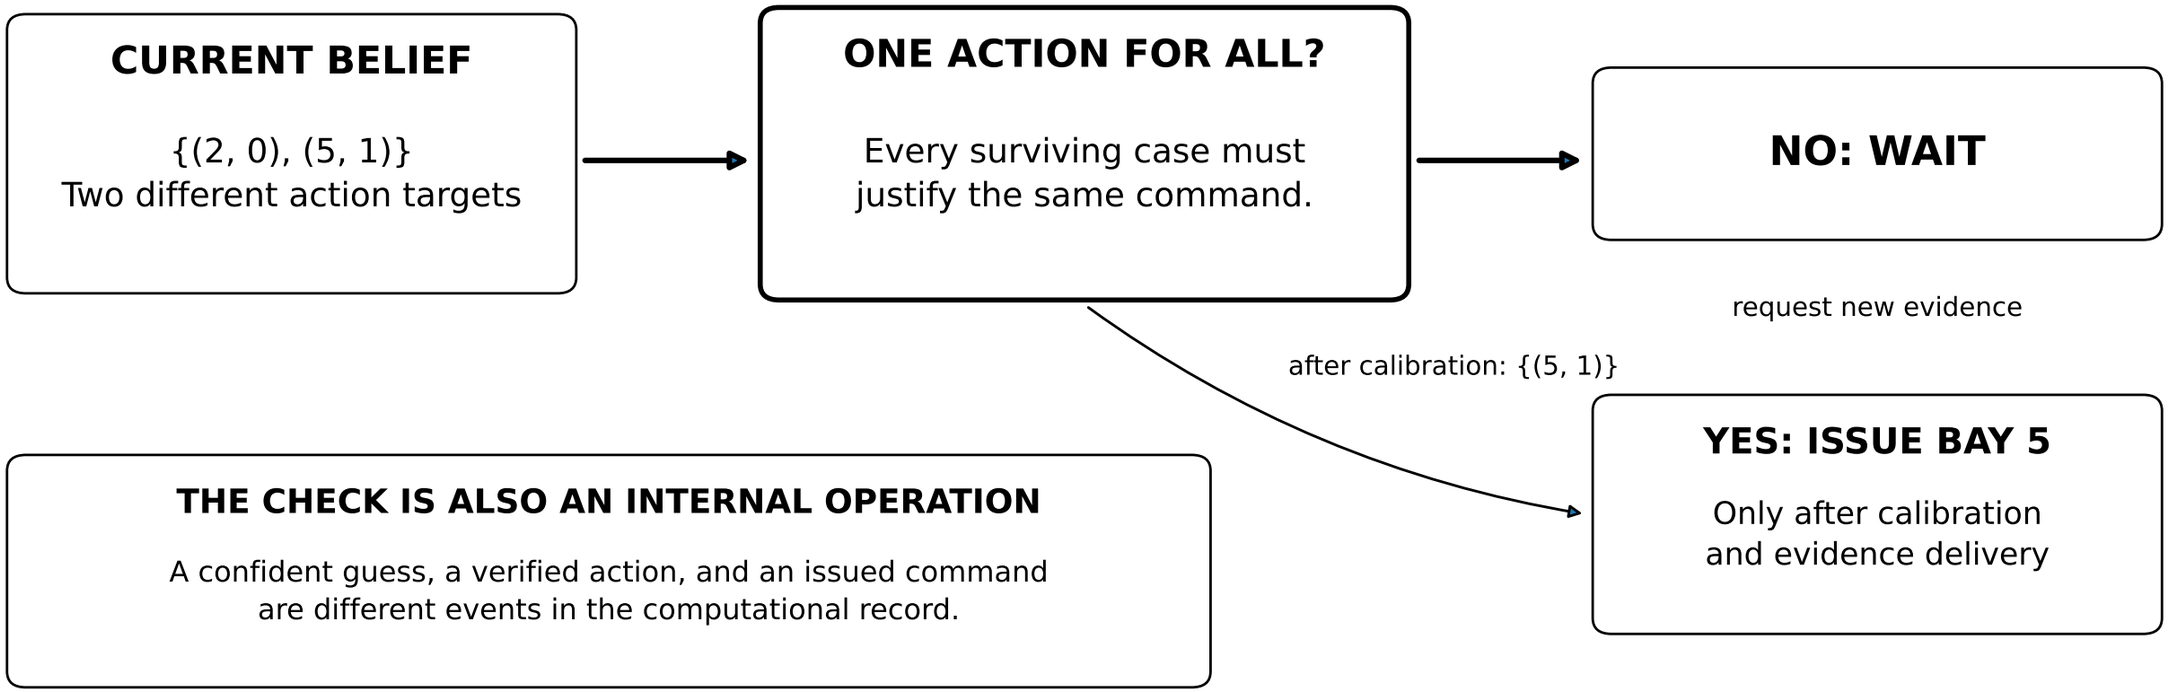


Without messages, A has high and middle bits while B has the low bit and calibration. Each retains four possible source locations. With all messages but no calibration, both retain two. These are exact information-access comparisons, not training runs or general performance scores.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

Waiting is part of the action policy

Agreement about uncertainty is not permission to choose an arbitrary member of it.



The gate checks every case in the agent's current support. This is not a probability threshold or a moral authority claim. At event 12, locations 2 and 5 remain possible; after calibration, only 5 remains. Correctness still depends on the actual world belonging to the model family.

Companion concept: Self-Improving Intelligence v1.0, Sections 7 and 31. New illustrative model: code/model.py.

State updating is not the same as self-improvement

The animation implements the upper loop. The lower loop is a proposed extension, not a hidden result.

FAST LOOP / executed in this model



Persistent belief change can occur while the update algorithm stays fixed.

SLOW LOOP / proposed learning extension



The toy agents update beliefs and replace a contradicted model scope using supplied rules. They do not learn a new query policy, alter source code, or improve their own learner. The manuscript studies additional finite mechanisms for compiler and rule acquisition; those experiments are not being rerun here.

Source: Self-Improving Intelligence v1.0, Sections 1, 23, 31. Atlas distinction; no claim of observed neural self-improvement.

What these displays mean

A visual model can be expressive without claiming access to an unmeasured mind.

WORLD	A fixed location j in 0,...,7 within each episode. The sensor convention changes once between investigations.
OBSERVATION	$o = \mathrm{bit}_q(j)$ XOR θ A calibration directly reports the orientation.
INTERNAL STATE	Uniform probability on the cases compatible with active receipts. Old receipts are archived when the supplied reset rule is used.
COMMUNICATION	Reliable delivery, honest receipts, explicit episode notices. Sharing is not independent remeasurement.
DECISION	Issue a command only when all remaining locations agree. The scripted schedule makes one final external commitment.
LIMIT	No neural training, real activation telemetry, consciousness claim, or general guarantee for arbitrary environmental change.

The source manuscript remains unchanged. The atlas extends its hidden-case and sensor-convention examples into a new scripted two-agent illustration. All frames, charts, and 3D coordinates are regenerated from code. No AI-generated pictorial asset is used.

References: Self-Improving Intelligence v1.0, Sections 7, 9, 24, 31. Canonical model, schedule, checks and CSV/JSON are included.