
ON THE LIMITS OF AN OBSERVED PRESENT

The MAYBE Paper

A finite theory of observation, unresolved states, and warranted abstention

Abstract. An observation may be precise while the conclusion attached to it remains undetermined. This monograph develops a finite mathematical framework for keeping those two facts separate. Its primitive objects are possible states, observation records, and propositions. A record removes states that contradict the evidence; a proposition is resolved only when it has the same truth value in every surviving state.

The resulting calculus has three outputs: YES, NO, and MAYBE. The third output is neither a numerical probability nor an evasive midpoint. It is an explicit certificate that the available evidence permits more than one answer. We derive this rule from observational equivalence, study its behavior under additional information, and compare it with Bayesian updating, entropy, and decisions with an abstention option.

Applications concern digital assets only at the level of observation. A chain identifier, a block number, bytecode, and an indexed pool price have different evidentiary scopes. A live network does not establish token deployment; token deployment does not establish utility; a present price does not identify a future return. These separations are formalized rather than left to typography.

The principal result is conditional and intentionally modest. If two admissible continuations share the complete observed record and disagree on the target proposition, no sound rule using that record alone can return a decisive answer. After thirty pages of definitions, calculations, counterexamples, and qualification, the permitted conclusion is MAYBE.

Observation $\neq$ Identification $\neq$ Prediction

(01)

The three tasks are distinct.

MAYBE Research / Discussion Paper 001 / Version 2.0 / 18 September 2026. An original conceptual monograph, not a peer-reviewed empirical study. No market dataset is estimated in this paper.

Reading map and notation

Conventions before conclusions

The argument has four movements. Pages 3-8 define evidence and the three-valued decision rule. Pages 9-20 examine probabilities, information, utility, and decisions. Pages 21-27 connect the formal language to the observation lab and its limitations. Pages 28-30 supply proofs, references, and the final result.

Throughout, Ω is a finite nonempty set of possible states. Finiteness keeps every infimum, sum, and conditional argument explicit. A proposition A is a subset of Ω . The realized state belongs either to A or its complement; the observer need not know which. R denotes a record, $S(R)$ the states compatible with it, and g an observation map.

Logical possibility is not the same as probabilistic weight. Unless a probability measure is introduced by an explicit assumption, no probability is assigned. In particular, two remaining states do not imply equal odds. A number chosen for an example is not an estimate of MAYBE, HyperEVM, or any real market.

Equations are numbered continuously. Definitions and propositions are internal to this finite framework unless a bibliographic source is identified. The phrase 'sound' means correct for every state admitted by the stated assumptions. It does not mean that those assumptions have been empirically validated.

Reading the document linearly is useful but not necessary. For the core argument, read pages 3, 5, 7, and 30. For the website's data semantics, read pages 22-24. For the distinction between a full derivation and an actual forecast, read pages 15-18.

$$\emptyset \neq S(R) \subseteq \Omega, \quad A \subseteq \Omega \quad (02)$$

Admissible states and the proposition under examination.

$$A^c = \Omega \setminus A, \quad |S| = \sum_{\omega \in S} 1 \quad (03)$$

Complement and cardinality conventions.

YES and NO concern propositions. OBSERVED, UNOBSERVED, and FAILED concern records. These vocabularies must not be interchanged.

1. The finite state space

Possibility precedes probability

Definition 1.1. A state describes every feature required to evaluate the target proposition. The finite set Ω collects the possibilities admitted before the current record is applied. A state can include a present condition and a future continuation; it need not be a snapshot of the current interface.

As a minimal example, take two independent binary coordinates: whether a network record is retrievable and whether a specified future use case succeeds. There are four states. Retrieving a record removes the two states in which retrieval fails. It leaves both outcomes for the use case available. This is a structural statement about the chosen model, not an empirical independence claim.

State spaces are modeling choices. Omitting a relevant continuation can produce false certainty; adding impossible continuations can produce unnecessary indecision. The formalism cannot certify the correctness of Ω from notation alone. It exposes the consequences of a declared space so that those choices can be inspected.

A refinement replaces a coarse state by several detailed states. A proposition defined on the coarse space may be lifted to the refined space through a projection. When the evidence does not distinguish the new details, refinement need not add knowledge. More variables can describe ignorance with greater resolution.

In the pages that follow, every conclusive answer is relative to Ω and the evidence policy. Absolute certainty about the adequacy of the model is neither assumed nor smuggled into the notation.

$$\Omega = \{0, 1\} \times \{0, 1\}, \quad |\Omega| = 4 \quad (04)$$

A finite illustrative space with two binary coordinates.

$$S = \{(1, 0), (1, 1)\}, \quad A = \{(0, 1), (1, 1)\} \quad (05)$$

A successful observation leaves both values of the target coordinate.

$$\pi : \tilde{\Omega} \rightarrow \Omega, \quad \tilde{A} = \pi^{-1}(A) \quad (06)$$

Lifting a proposition under state-space refinement.

The four-state example is synthetic. It contains no estimated probability and no actual token outcome.

2. The observation record

A record should survive its own retelling

Definition 2.1. A record contains an identifier, the measurement time, the retrieval time, the source, the requested target, the returned value, and a status. The target is essential: a number without an object of measurement is not a useful observation. The distinction between measurement and retrieval time is equally essential when sources are delayed.

A record is immutable once written. Refreshing an interface does not alter the original time of an old measurement. A cached value may be shown again, but it remains a repetition of the same record. Counting that repetition as an independent observation would inflate the apparent evidence without adding information.

Validation occurs before interpretation. A response may be syntactically valid JSON yet refer to the wrong chain. A token address may be a valid hexadecimal string yet identify an unrelated asset. The validation rule must therefore include both format checks and target checks.

Successful transport is weaker than successful measurement. An HTTP response can report an upstream failure, an empty result, or a method error. The client should preserve those distinctions rather than treating every completed request as an observed state.

The website implements a practical subset of this schema: record ID, retrieval time, block timestamp, chain ID, source, token status, optional market values, and issues. Each explicit run requests a fresh collection. Passive reads may reuse a brief cache with the original record identity. This difference is visible rather than concealed.

$$R = (i, t_m, t_r, s, q, v, \sigma), \quad t_m \leq t_r \quad (07)$$

Record fields; the time inequality assumes a consistent clock convention.

$$R_i = R_j \implies \text{new evidence}(R_i, R_j) = 0 \quad (08)$$

Replaying the same record does not create an independent measurement.

$$\text{valid}(R) = \text{format}(R) \wedge \text{target}(R) \quad (09)$$

Structural validity and target validity are separate checks.

A timestamp is a claim by a source, not a proof that every participating clock is synchronized.

3. Observational equivalence

The kernel of what can be known

Let g map each state to the record that an idealized deterministic observation would return. Two states are observationally equivalent when they produce the same record. This relation is reflexive, symmetric, and transitive, so it partitions Ω into equivalence classes.

Definition 3.1. The class of a realized state is its observational fiber. If the observer knows only $g(\omega)$, the fiber is the finest set that this observation alone can identify. A proposition that cuts across that fiber cannot be settled by the record.

Proposition 3.2. A proposition is identifiable from g on the whole state space if and only if it is a union of fibers. Proof: if it is a union of fibers, membership is constant on every possible observation. Conversely, if one fiber contains a state inside A and a state outside A , the same input record would require two different decisive answers.

This is the basic obstruction behind the document. Adding digits to a displayed value may leave the fiber unchanged. Reformatting the response may also leave it unchanged. The appearance of analytical depth does not split an equivalence class unless the new operation introduces a distinguishing observation.

A practical example is a network record shared by two possible project futures. Both futures can agree on every present block field and disagree on later utility. The present record does not distinguish them. Any conclusion selecting one must rely on an additional premise.

$$\omega \sim_g \omega' \iff g(\omega) = g(\omega') \quad (10)$$

Observational equivalence induced by a measurement map.

$$[\omega]_g = g^{-1}(\{g(\omega)\}) \quad (11)$$

The observational fiber.

$$A = g^{-1}(B), \quad B \subseteq g(\Omega) \quad (12)$$

An identifiable proposition is a union of complete fibers.

The proof is finite and deterministic. No claim about physical quantum measurement is required.

4. Admissible sets

Evidence as elimination rather than decoration

For each valid record R , let $E(R)$ be the subset of states compatible with that record under the declared observation model. Combining several records amounts to intersecting their compatible sets. This makes the inference inspectable: every removed state must conflict with at least one admitted observation.

Proposition 4.1. Adding a record cannot enlarge the admissible set when the model and evidence policy are fixed. This monotonicity is a property of set intersection, not a promise that every new record is correct. A bad record may shrink the set incorrectly.

If an intersection becomes empty, the system has a consistency problem. At least one record, assumption, target alignment, or modeling choice is incompatible with the rest. The appropriate output is an evidence-conflict state. Treating the empty set as a proof of both A and its complement would be formally vacuous and operationally misleading.

Observations may also be redundant. A new record whose compatible set already contains S adds nothing to the current identification problem. It may improve auditability or operational confidence without changing the set of admissible answers.

The distinction between redundant and refining observations clarifies why repeated clicks can be useful operationally but not logically decisive. A later block confirms that a network responds at a later time. It need not remove any continuation relevant to a question about future usefulness.

$$S_n = \Omega \cap \bigcap_{j=1}^n E(R_j) \quad (13)$$

The admissible set after n compatible records.

$$S_{n+1} = S_n \cap E(R_{n+1}) \subseteq S_n \quad (14)$$

Evidence refinement under a fixed model.

$$S_n = \emptyset \implies \text{EVIDENCE CONFLICT} \quad (15)$$

The empty-set case is handled before applying the answer rule.

A smaller set is more specific. It is not automatically more truthful.

5. The three-valued answer

A decisive answer needs agreement

Definition 5.1. For a nonempty admissible set S and a proposition A , output YES when every state in S belongs to A . Output NO when no state in S belongs to A . Output MAYBE when S contains at least one state of each kind.

This rule is deliberately asymmetric with respect to rhetorical confidence. The observer may strongly prefer one answer, yet the rule depends only on set membership. It asks what follows from the admitted evidence, not how persuasive a narrative feels.

Proposition 5.2. The three cases are mutually exclusive and exhaustive for nonempty S . Proof: the intersection S with A has cardinality between zero and the cardinality of S . The endpoints give NO and YES; every intermediate cardinality gives MAYBE.

The rule is not equivalent to a three-valued truth of the world itself. Within any single state, the proposition remains ordinarily true or false. MAYBE belongs to the observer's relation to a set of possible states. The distinction prevents the category error of treating incomplete knowledge as a new physical state of a token.

Once a proposition has been resolved, additional compatible evidence cannot reverse the answer under fixed assumptions. It may, however, reveal an inconsistency that forces the assumptions or record policy to be revised. Apparent reversals therefore require investigation rather than a claim that set intersection has changed its behavior.

$$\text{YES : } S \subseteq A, \quad \text{NO : } S \subseteq A^c \quad (16)$$

The two decisive cases.

$$\text{MAYBE : } 0 < |S \cap A| < |S| \quad (17)$$

Both answers remain represented among admissible states.

$$D(A, S) \in \{\text{YES, NO, MAYBE}\} \quad (18)$$

The answer alphabet.

MAYBE is an epistemic output. It is not a claim that reality has a third Boolean truth value.

6. Probability without invented odds

The interval induced by ignorance

A set of possible states does not specify a probability distribution. If only support is known, the compatible probabilistic models are all distributions supported on S . This family can express complete uncertainty about the relative weights of admissible states without pretending that they are equally likely.

Define lower and upper probabilities as the smallest and largest probability assigned to A across that family. If S contains a state inside A and a state outside A , point masses on those two states attain the endpoints zero and one. The entire interval remains possible.

Proposition 6.1. Under unrestricted support-compatible probabilities, the interval for A is $[1,1]$ in the YES case, $[0,0]$ in the NO case, and $[0,1]$ in the MAYBE case. The proof follows by constructing the relevant point masses and their convex mixtures.

This result is intentionally broad. Additional probabilistic restrictions may narrow the interval. Such restrictions must be named: a prior, an exchangeability condition, an empirical calibration procedure, or a bound on a parameter. Their numerical consequences should not be attributed to the observation alone.

In particular, the phrase 'it could go up or down' does not imply a fifty-fifty forecast. Exhausting a list of possibilities is not the same task as assigning their probabilities. The absence of a justified number is not repaired by selecting a visually balanced one.

$$\mathcal{P}(S) = \{p : p(\omega) \geq 0, \sum_{\omega \in S} p(\omega) = 1\} \quad (19)$$

All probability masses supported on the admissible set.

$$\underline{P}(A) = \inf_{p \in \mathcal{P}(S)} p(A), \quad \bar{P}(A) = \sup_{p \in \mathcal{P}(S)} p(A) \quad (20)$$

Lower and upper probabilities.

$$p_\lambda = \lambda \delta_{\omega_1} + (1 - \lambda) \delta_{\omega_0}, \quad 0 \leq \lambda \leq 1 \quad (21)$$

Mixing one state in A and one outside A realizes every probability in the interval.

The interval $[0,1]$ is a statement of missing identification, not an estimated confidence interval.

7. Bayesian updating, conditionally

A posterior inherits its premises

Introduce a prior mass function π on Ω and a likelihood $L(r|\omega)$ for the observed record. Bayes' rule yields a posterior whenever the normalizing denominator is positive. These are additional modeling objects; neither is supplied merely by drawing a probability symbol.

Under deterministic observation, the likelihood is an indicator of the relevant fiber. The posterior is then the prior restricted to that fiber and renormalized. Observation removes incompatible states, but relative weights among observationally identical states remain in the ratios assigned before the observation.

For two states with the same likelihood, the posterior odds equal the prior odds. This is an especially clear form of unresolved inference: the data has not distinguished the pair. A precise posterior can still be driven entirely by the analyst's starting allocation.

Bayesian calculations are therefore not rejected by this framework. They are labeled according to their assumptions. A posterior probability is a conditional output of a specified prior-likelihood pair. It should not be reported as a fact obtained directly from a network request.

If all prior mass lies outside the observed fiber, the normalizer is zero and the conditional update is undefined. This is a model-data conflict, not evidence that an arbitrary fallback probability has become correct. The model must be revised before a numerical posterior can be interpreted.

$$\pi(\omega \mid r) = \frac{L(r \mid \omega)\pi(\omega)}{\sum_{u \in \Omega} L(r \mid u)\pi(u)} \quad (22)$$

Bayes' rule when the denominator is positive.

$$L(r \mid \omega) = \mathbf{1}_{\{g(\omega) = r\}} \quad (23)$$

The likelihood of a deterministic observation.

$$\frac{\pi(\omega_1 \mid r)}{\pi(\omega_0 \mid r)} = \frac{\pi(\omega_1)}{\pi(\omega_0)} \quad (24)$$

Equal positive likelihoods preserve prior odds.

The posterior-odds identity assumes positive prior masses and a defined posterior.

8. Prior sensitivity

The conclusion may move while the data stands still

Consider two admissible states that have equal positive likelihood for the observed record. Give the first state prior mass q and the second mass $1-q$. The posterior probability of the first state remains q . Any value strictly between zero and one can therefore be produced without changing the record.

This two-state construction is enough to refute the claim that a posterior number is necessarily data-determined. It does not imply that all priors are equally defensible. It shows exactly where a defense of the selected prior would have to enter.

A sensitivity analysis evaluates the same target under a declared family of priors. The width between the largest and smallest posterior probabilities measures dependence on that family for the fixed likelihood and record. A wide range is not an implementation failure; it is a feature of the stated assumptions.

In more informative models, likelihood ratios can dominate bounded prior differences. That possibility does not rescue an uninformative record. A block number that is identical under two future continuations contributes a likelihood ratio of one for that pair, regardless of how many digits it contains.

The website does not display a posterior probability for success, utility, or appreciation because no prior-likelihood model for those targets has been estimated. Showing no such number is consistent with the mathematics. It is also more precise than displaying a probability whose premises cannot be inspected.

$$\pi_q(\omega_1) = q, \quad \pi_q(\omega_0) = 1 - q \quad (25)$$

A family of two-state priors.

$$L(r \mid \omega_1) = L(r \mid \omega_0) > 0 \implies P_q(A \mid r) = q \quad (26)$$

The observation leaves the chosen weight unchanged.

$$W(r) = \sup_{\pi \in \Pi} P_{\pi}(A \mid r) - \inf_{\pi \in \Pi} P_{\pi}(A \mid r) \quad (27)$$

Posterior sensitivity over a declared family of priors.

Sensitivity is always relative to a chosen model family. It is not a universal error bar.

9. Entropy is not a conclusion

A scalar summary cannot replace a proposition

For a specified finite probability mass function, Shannon entropy summarizes dispersion. It is maximized by the uniform distribution on a fixed support and vanishes at a point mass. These standard properties concern a probability model, not a bare set of possibilities.

The maximum-entropy principle selects a distribution subject to explicit constraints. With only normalization on a finite support, it selects equal weights. That selection is a modeling convention; it does not demonstrate that an actual market or project has equal frequencies across those states.

Entropy and logical resolution answer different questions. A distribution may place almost all mass on one state while retaining another state with positive mass that reverses A. Logical resolution over its support remains MAYBE even when entropy is very small.

Conversely, many states can all agree on A. Entropy may be large because other coordinates vary, while A itself is fully resolved. A display of uncertainty without a target proposition can therefore obscure exactly what is unknown.

The binary entropy curve illustrates this separation. Its maximum occurs at one-half, but MAYBE is not the name of that maximum. MAYBE applies to every nontrivial admissible split under the set-based rule, irrespective of whether a separate probabilistic model is balanced or highly concentrated. Shannon's formulation supplies the information measure; the epistemic distinction here is our own finite construction [1].

$$H(p) = - \sum_{\omega \in S} p(\omega) \log p(\omega), \quad 0 \log 0 = 0 \quad (28)$$

Entropy for a specified probability mass function.

$$0 \leq H(p) \leq \log |S| \quad (29)$$

The finite-support entropy bounds.

$$h(q) = -q \log q - (1 - q) \log(1 - q) \quad (30)$$

Binary entropy; the endpoint values are defined by continuity.

Natural logarithms give entropy in nats. Entropy is not a token valuation.

10. Processing cannot create evidence

A more elaborate display can be less informative

Let X be a random state, Y its observed signal, and Z a deterministic transformation of that signal. Every value of Z collects one or more values of Y . The transformation can preserve information or discard distinctions, but it cannot separate states already identified by the same Y .

In the finite probabilistic setting, the data-processing inequality expresses this ordering through mutual information. The statement requires a specified joint distribution and the Markov structure X to Y to Z . It is not a claim that every numerical summary of information has the same monotonicity.

There is also a purely set-based version. If two states share the same Y , they necessarily share $f(Y)$. Consequently, fibers of the processed observation are unions of fibers of the original observation. Formatting a record into a chart cannot add a new distinction about the underlying state.

This has an immediate interface consequence. A waveform, a scanning animation, and a large equation can make a page legible or memorable. They cannot increase what the returned fields establish. Their contribution is presentation, not evidence.

A useful display can nevertheless improve human reasoning by exposing source, time, target, and missingness. That improvement comes from reducing misinterpretation of available information. It should not be confused with collecting a new observation. The visual system is successful when it makes the limits of the record easier to see.

$$X \longrightarrow Y \longrightarrow Z, \quad Z = f(Y) \quad (31)$$

A state, a signal, and a processed display.

$$I(X; Z) \leq I(X; Y) \quad (32)$$

Data processing for the stated finite Markov model.

$$g(\omega) = g(\omega') \implies f(g(\omega)) = f(g(\omega')) \quad (33)$$

The deterministic fiber argument.

The finite information inequality is standard; the website application is an interpretation, not an empirical result.

11. Utility requires a context

Useful for whom, for what, and at what cost?

The predicate 'useful' is incomplete without a context. A service may save time for one task and add friction to another. A digital asset may have a technical function without producing an advantage for a particular user. The framework therefore treats utility as a contextual comparison rather than a permanent scalar attached to a name.

Let $u(a, \omega; c)$ denote a specified outcome score for action a , state ω , and context c . A useful-state proposition can be defined by comparing that score with a baseline and a threshold. Neither the score nor the threshold is inferred from a token symbol.

An empirical utility claim would require a measurement protocol: target users, task, baseline, horizon, costs, and an evaluation rule. Without those choices, the phrase 'utility demonstrated' has no single testable meaning. Naming a feature is not the same as measuring its net contribution.

The familiar superposition-like brand expression can be read as a notation for unresolved alternatives. It is not a wavefunction measured in a laboratory, and its coefficients are not estimated physical amplitudes. The actual logic in this paper operates on admissible sets.

Under the finite model, usefulness is resolved only if every admissible state exceeds the declared baseline and threshold. If some do and some do not, the answer remains MAYBE. This is demanding by design: the word 'useful' should not obtain its evidence from its own repetition.

$$\Delta u(\omega; c) = u(a, \omega; c) - u(a_0, \omega; c) \quad (34)$$

Utility relative to a declared baseline action.

$$A_c = \{\omega : \Delta u(\omega; c) \geq \tau_c\} \quad (35)$$

A contextual usefulness proposition with a declared threshold.

$$\min_{\omega \in S} \Delta u(\omega; c) \geq \tau_c \implies \mathcal{D}(A_c, S) = \text{YES} \quad (36)$$

A sufficient and necessary condition for YES under this definition.

The model does not provide a utility score for MAYBE. It specifies what such a claim would need to mean.

12. Directional exhaustion

An exhaustive theorem with no directional preference

Theorem 12.1. If a current and a later price are finite real numbers measured in the same unit, exactly one of three relations holds: the later price is greater, less, or equal. This is the trichotomy property of the real order applied to a pair of prices.

Proof. Let d be the later price minus the current price. Every real number is positive, negative, or zero, and these cases are mutually exclusive. Substituting the definition of d gives the three alternatives. No market assumption is needed.

The theorem is complete as an enumeration and empty as a directional forecast. It gives probability one to the union of all alternatives under any probability model in which those alternatives are defined. It gives no individual alternative a specified probability.

A missing later price is outside the theorem's premise. Delisting, an unavailable indexer, incompatible units, or a nonnumeric response cannot be treated as equality. The set of operational outcomes must include missingness separately when the instrument may fail.

Remark 12.2. A statement can be mathematically correct, beautifully typeset, and useless for selecting a trade. These properties are not contradictory. The theorem's role is to reveal the difference between covering the outcome space and narrowing it. The conclusion 'one of these will occur' should not be marketed as knowledge of which one.

$$\Delta P = P_{t+1} - P_t \in \mathbb{R} \quad (37)$$

A finite, consistently denominated price difference.

$$\{\Delta P > 0\} \cup \{\Delta P = 0\} \cup \{\Delta P < 0\} = \Omega_P \quad (38)$$

The three directions exhaust the domain where both prices exist.

$$P(\Delta P > 0) + P(\Delta P = 0) + P(\Delta P < 0) = 1 \quad (39)$$

Normalization constrains the sum, not its individual terms.

Theorem 12.1 proves that the next measured price has a direction. It does not tell us the direction.

13. A future is not identified by a snapshot

A constructive non-identification argument

Fix a record ending at time t with current price p greater than zero. Construct two continuations that share every recorded value through t . In one, the next price is p plus epsilon. In the other, it is p minus epsilon, where epsilon is positive and smaller than p .

Both continuations are compatible with the snapshot unless an additional model rules one out. They disagree on the proposition that the next return is positive. The target is therefore not identifiable from that snapshot in the unrestricted continuation space.

This is a counterexample to universal prediction from the given record, not a proof that every statistical forecast is worthless. A forecasting model may impose dynamics, estimate conditional frequencies, and report calibrated uncertainty. Those steps introduce assumptions and data beyond the snapshot.

The strength of the argument lies in its economy. It does not need a complicated adversarial market. Two continuations suffice. If a proposed inference cannot distinguish them, adding rhetorical certainty to the output does not remove the ambiguity.

The same construction applies to future utility and chain expansion. A present record can be copied into two admissible futures with different later outcomes. A promise about one future requires evidence or a commitment that is not contained in the copied record. The website reports what was retrieved and leaves these targets unresolved.

$$P_{t+1}^{(+)} = p + \varepsilon, \quad P_{t+1}^{(-)} = p - \varepsilon, \quad 0 < \varepsilon < p \quad (40)$$

Two positive-price continuations.

$$R_t(\omega_+) = R_t(\omega_-), \quad \mathbf{1}_A(\omega_+) \neq \mathbf{1}_A(\omega_-) \quad (41)$$

Identical records and opposite answers.

$$\mathcal{D}(A, S(R_t)) = \text{MAYBE} \quad (42)$$

The answer in any admissible set containing both continuations.

Non-identification is relative to the chosen observation and model class. It is not a claim of universal unknowability.

14. Bounds do not appear by notation

The difference between an interval and evidence for it

Suppose a model supplies lower and upper bounds on the next price. A return interval follows by transforming those bounds through the current positive price. The algebra is straightforward. The hard part is justifying the original price bounds.

A narrow interval can be visually persuasive while resting on an arbitrary range. To make a bound informative, one must specify whether it is logical, empirical, probabilistic, or imposed as a scenario. Those meanings are not interchangeable.

A scenario interval says 'consider these possibilities.' A confidence interval concerns the behavior of a statistical procedure under stated sampling assumptions. A prediction interval concerns a future random quantity under a model. A deterministic guarantee excludes all outcomes beyond its endpoints. Using the same bracket notation does not make these claims equivalent.

If no finite upper bound is supported, it should not be manufactured for layout convenience. If a current price is unavailable, a normalized return may not be defined. If the observed price belongs to one pool, even the current denominator has a venue and a time.

The MAYBE framework therefore permits intervals only with their provenance. An interval produced by explicitly chosen scenarios may be useful for exploration; it must remain labeled as a scenario. In this monograph, the calculation below is an algebraic transformation of assumed bounds, not an estimated range for any token.

$$0 < P_t, \quad L \leq P_{t+1} \leq U \quad (43)$$

Assumed finite bounds, not measured forecasts.

$$r_{t+1} = \frac{P_{t+1} - P_t}{P_t} \quad (44)$$

Simple return when the current price is positive.

$$\frac{L}{P_t} - 1 \leq r_{t+1} \leq \frac{U}{P_t} - 1 \quad (45)$$

The resulting return interval preserves the assumptions of the price interval.

Brackets can organize uncertainty. They cannot establish their own endpoints.

15. A model of motion

Conditional calculations remain conditional

To illustrate how assumptions enter a dynamic calculation, let X be a real-valued process with independent increments taking values plus a and minus a with equal probability. After n steps, the expected increment is zero and its variance is n times a squared.

The derivation follows by linearity of expectation and additivity of variance for independent increments. If independence is removed, covariance terms reappear. If equal probabilities are removed, a drift appears. The attractive simplicity of the formula is purchased by those assumptions.

The process is an abstract coordinate, not necessarily a price. An unconstrained additive walk may become negative, which is inappropriate for many price definitions. Defining a positive process by exponentiating X resolves positivity but changes the relation between expected coordinate and expected level.

Indeed, a zero expected change in X does not imply a zero expected change in its exponential. Convexity matters. Even in this elementary toy model, the variable being averaged must be specified carefully before a result is interpreted.

This page is included to make the scientific style accountable to its assumptions. We can write a motion model, derive its moments, and identify its limitations. None of those steps estimates the process governing MAYBE. Without such estimation and validation, the model supplies an example of conditional reasoning and no operational forecast.

$$X_n = X_0 + \sum_{k=1}^n \xi_k, \quad P(\xi_k = a) = P(\xi_k = -a) = \frac{1}{2} \quad (46)$$

A synthetic symmetric random walk.

$$\mathbb{E}[X_n - X_0] = 0, \quad \text{Var}(X_n - X_0) = na^2 \quad (47)$$

Moments under independent increments.

$$P_n = e^{X_n}, \quad \mathbb{E}[P_n] = e^{X_0}(\cosh a)^n \quad (48)$$

Exponentiation yields a positive process with a different mean behavior.

All parameters are illustrative. No calibration, backtest, or performance claim is implied.

16. The value of another observation

An experiment matters through the decision it can change

More observations are not automatically more useful. Their value depends on the action set, the loss function, the distribution of states, and the information already available. An observation that cannot change the preferred action may have no decision value in a particular problem.

Under a specified finite probability model, compare the minimum expected loss before seeing a signal with the expected minimum conditional loss after seeing it. Their difference is the expected value of information. With no observation cost and permission to ignore the signal, this value is nonnegative.

The proof is elementary: after receiving the signal, the decision maker can always use the action selected before receiving it. Optimizing over a larger class of signal-dependent choices cannot do worse in expectation under the same model.

This does not imply that every data collection is worthwhile. Acquisition costs, latency, privacy costs, and model errors can outweigh the expected improvement. It also does not imply that an observation about one target is valuable for another target.

In the lab, retrieving the latest block is useful for checking network responsiveness and maintaining an auditable record. It need not improve a decision about future token utility. The general comparison of experiments is a classical subject [2]; this page uses only the simple finite decision argument stated here.

$$V_0 = \min_a \mathbb{E}[\ell(a, \omega)] \quad (49)$$

Minimum expected loss without the additional signal.

$$V_1 = \mathbb{E}_Y[\min_a \mathbb{E}[\ell(a, \omega) \mid Y]] \quad (50)$$

Expected optimized loss after observing Y .

$$\text{VOI} = V_0 - V_1 \geq 0, \quad \text{net value} = \text{VOI} - c \quad (51)$$

Information value before and after an explicit acquisition cost.

The nonnegativity statement assumes the same valid model on both sides and the option to ignore the signal.

17. Abstention as a decision

MAYBE can be a selected action rather than a failure

Consider a binary proposition with three available actions: assert YES, assert NO, or abstain. Assign loss zero to a correct assertion, loss one to an incorrect assertion, and a fixed loss c to abstention, where c lies strictly between zero and one-half.

If a posterior probability p is available, the expected losses are $1-p$, p , and c respectively. Abstention minimizes expected loss whenever p lies between c and $1-c$, with ties permitted at the boundaries. This is a decision-theoretic criterion, distinct from logical resolution.

Now remove the point probability and retain only the interval $[0,1]$. Under worst-case expected loss, either decisive assertion has loss one, while abstention has loss c . The robust choice is therefore abstention. This conclusion follows from the declared loss structure rather than from a preference for indecision.

Different costs produce different decisions. A time-critical system may rationally act despite unresolved truth. The act of choosing does not retroactively turn an uncertain proposition into a known one. Decision and identification remain separate tasks.

The word MAYBE can label the epistemic result while an application follows its own action policy. In this website, unresolved future claims are not forced into a decisive marketing answer. That is a product choice consistent with the finite logic, not a theorem that all real-world agents must wait forever.

$$L_{\text{YES}}(p) = 1 - p, \quad L_{\text{NO}}(p) = p, \quad L_{\text{MAYBE}}(p) = c \quad (52)$$

Expected losses under a specified posterior and abstention cost.

$$c \leq p \leq 1 - c \implies \text{MAYBE} \in \arg \min_a L_a(p) \quad (53)$$

The posterior region where abstention is optimal.

$$\sup_{p \in [0, 1]} L_{\text{YES}}(p) = 1 > c = \sup_{p \in [0, 1]} L_{\text{MAYBE}}(p) \quad (54)$$

Worst-case loss when the probability is unrestricted.

Logical MAYBE and optimal abstention may coincide, but they are defined by different rules.

18. Robustness to an assumption change

A conclusion should reveal what holds it up

Let a model class contain several admissible state sets, each produced by a different defensible assumption package. A conclusion is robust across that class only when it survives every member. A decisive answer in one narrow model may disappear when another plausible model is admitted.

The union of the model-specific admissible sets represents all states allowed by at least one model. Applying the three-valued rule to this union yields a conservative answer. If any model admits a counterexample, the union prevents an unconditional YES.

Care is needed when models use different state descriptions. Their states must first be embedded in a common space or connected by a declared correspondence. Taking a union of incompatible encodings is not a meaningful robustness calculation.

A useful audit identifies which assumption excludes each counterexample. If an answer depends on a particular growth premise, a guaranteed integration, or an unmeasured utility threshold, that dependence belongs beside the answer. Hiding it in a footnote does not make the conclusion less conditional.

This page provides an operational test for increasingly elaborate narratives. Ask whether the extra complexity introduces distinguishing evidence or merely narrows the model by assumption. Both moves can be mathematically legitimate, but only the first should be described as having observed more.

$$S_{\text{rob}} = \bigcup_{m \in \mathcal{M}} S_m \quad (55)$$

The robust admissible set on a common state space.

$$S_{\text{rob}} \subseteq A \iff S_m \subseteq A \text{ for all } m \quad (56)$$

YES across a model family requires agreement in every model.

$$\exists m, \omega \in S_m \cap A^c \implies \mathcal{D}(A, S_{\text{rob}}) \neq \text{YES} \quad (57)$$

One admitted counterexample blocks a universal affirmative answer.

Robustness is relative to the inspected family. An unconsidered model remains unconsidered.

19. Networks as a product state

Availability, deployment, and identity are different coordinates

A network can be available while a particular contract is absent. A contract can be present while its project identity is unverified. A project can be authentic while a future use case remains uncertain. These are separate coordinates and should be represented separately.

Let the network coordinate record availability, the deployment coordinate record code presence at an address, and the identity coordinate record whether the association with the named project has been established. The product space allows combinations that a single green status badge would conceal.

HyperEVM is the designated genesis network for the MAYBE presentation. Reading chain ID 999 confirms the target of the network query when the source behaves as assumed. It does not by itself supply a MAYBE contract address or prove a deployment.

Other networks in the interface represent possible states rather than scheduled commitments. The framework does not convert a displayed network name into a roadmap. A future deployment must be supported by a concrete target, a chain-specific record, and an identity association.

The same logic applies to bytecode. Nonempty code is evidence of code at the queried address and block. It does not demonstrate source verification, safe behavior, ownership, liquidity locking, or the authenticity of a social account. Each of those claims needs its own evidence policy.

$$\Omega_C = \Omega_N \times \Omega_D \times \Omega_I \times \Omega_U \quad (58)$$

Network, deployment, identity, and utility coordinates.

$$\text{network available} \not\Rightarrow \text{token deployed} \quad (59)$$

Availability does not entail a specific deployment.

$$\text{code}(a, b) \neq \text{empty} \not\Rightarrow \text{identity}(a) = \text{MAYBE} \quad (60)$$

Code presence does not establish project identity.

A network label is a target of inquiry. It is not an endorsement or a deployment announcement.

20. Missing values are not zeros

Failure must remain a first-class result

A missing price and a measured price of zero are different observations. Substituting zero for missingness changes the meaning of arithmetic, charts, and status messages. It can manufacture a return, a liquidity figure, or an apparent collapse that was never measured.

Introduce an explicit missing symbol outside the numerical domain. Arithmetic functions are then defined only when their required inputs are numeric, unless a separate missing-data rule is declared. The interface should carry missingness through to its output rather than coercing it away.

Missingness has causes. A contract may not be configured; a source may be unavailable; a pool may not be indexed; a response may fail validation. Collapsing these causes into one blank field makes an operational failure indistinguishable from an unasked question.

Partial success also matters. If the block and chain ID are valid while gas-price retrieval fails, the valid block should remain an observed result with a specific missing gas field. A failure in an optional field need not erase the entire record.

Conversely, a successful optional request cannot repair an invalid core target. Receiving a market response for another chain does not validate the intended chain query. The collection layer and the display layer both need these boundaries. A clear explanation of what was not measured is part of a correct result.

$$\mathcal{V} = \mathbb{R}_{\geq 0} \cup \{ \perp \}, \quad \perp \neq 0 \quad (61)$$

An explicit missing value is outside the numeric domain.

$$f(x_1, \dots, x_k) \text{ defined only if } x_i \neq \perp \text{ for all } i \quad (62)$$

A conservative rule for arithmetic on incomplete inputs.

$$R = (R_{\text{core}}, R_{\text{optional}}), \quad \sigma \in \{\text{observed, partial, failed}\} \quad (63)$$

A record can retain valid core data while reporting an optional failure.

UNCONFIGURED is a configuration state. *FAILED* is a request outcome. Neither means a token is worth zero.

21. Time, staleness, and finality

A fresh request does not make every source current

An observation has more than one time. The chain block has a source timestamp; the server has a collection time; the viewer has a display time. Their differences expose delay and staleness, subject to clock uncertainty. These times should not be silently substituted for one another.

A new request can return the same block as a previous request. That does not make it a cache bug: the source may not yet have advanced. The record identity and collection time can still differ. The correct test of freshness is whether a new collection occurred, not whether every measured value changed.

Finality is a further property. A block returned as latest is not automatically a proof of irreversible settlement under every chain's rules. This site retrieves a latest-block observation; it does not implement a cross-network finality verifier.

Repeated observations produce a time series of records, not a complete history of all intervening states. Gaps between samples can hide events. A sequence of successful queries also does not prove continuous availability during the intervals.

The practical interface labels saved records when they age and allows the user to run another observation. It preserves timestamps in the archive and avoids presenting old values as newly collected. This is a modest but necessary form of temporal honesty.

$$a = t_{\text{display}} - t_{\text{collect}}, \quad d = t_{\text{collect}} - t_{\text{block}} \quad (64)$$

Display age and apparent source delay.

$$R_i \neq R_j \not\Rightarrow v_i \neq v_j \quad (65)$$

Different collection records need not contain different values.

$$\{R(t_1), \dots, R(t_n)\} \neq \{R(t) : t_1 \leq t \leq t_n\} \quad (66)$$

Discrete samples are not continuous surveillance.

The delay interpretation assumes sufficiently comparable clocks. Latest-block retrieval is not a finality proof.

22. The observation protocol

A concrete instrument with a limited remit

The lab begins with a user action. It requests a fresh server collection and shows a running state. The server queries the HyperEVM endpoint for chain ID, latest block, and gas price. It validates the chain target and numeric fields before assembling a record.

If a nonzero valid contract address is configured, the server separately requests bytecode and standard ERC-20 supply metadata. The metadata calls can fail even when code exists; the record preserves that distinction. Without a configured address, no token-specific measurement is invented.

Market data is a separate indexed observation. Candidate pools are restricted to the intended chain and the configured token as base asset. Among eligible entries, the interface selects the pool with the largest reported liquidity. This deterministic selection is not a claim that one pool represents every market.

Timeouts are explicit. A failed collection does not become an observed state merely because the interface stops spinning. The client releases its running state, offers another attempt, and preserves a readable explanation. Valid partial fields remain available with their issues.

The browser archive stores at most one hundred records locally. It is a convenience for the current browser, not a public registry or an independently verified database. JSON export preserves the currently selected records. The protocol measures the present under the source assumptions described in the limitations.

request → retrieve → validate → record → display (67)

The observation pipeline.

$Q = \{q : \text{chain}(q) = 999, \text{base}(q) = a\}$ (68)

Conceptual eligible-pool set; the indexer's chain label is mapped to HyperEVM.

$q^* \in \arg \max_{q \in Q} \text{reported liquidity}(q)$ (69)

The selected indexed pool, when eligible records exist.

Sources: HyperEVM documentation [4] and DEX Screener API documentation [5]. Source availability and correctness remain external assumptions.

23. A complete synthetic example

Four states, two records, one unresolved proposition

Consider four states, labeled w_1 through w_4 . The proposition A is true in w_1 and w_3 , and false in w_2 and w_4 . The first record admits w_1 , w_2 , and w_3 . A second record admits w_1 , w_2 , and w_4 . Their intersection leaves w_1 and w_2 .

Neither record is useless: each removes a state, and together they remove half the original space. Nevertheless, the two survivors disagree on A . The three-valued rule therefore returns MAYBE. Evidence has increased without resolving the selected proposition.

A third possible record admitting only w_1 would produce YES. A different third record admitting only w_2 would produce NO. A record admitting only w_3 would conflict with the first two and produce an empty intersection, which is handled as evidence conflict rather than an answer.

Under an additional uniform prior on the four original states and deterministic record compatibility, the first posterior probability of A is two-thirds and the posterior after two records is one-half. That probability can move downward even though A has not changed and the evidence set has become smaller.

This example separates four quantities often confused in a dashboard: number of records, number of surviving states, probability under a chosen model, and logical resolution of a proposition. None is a substitute for the others. The calculation is fully specified and intentionally contains no actual market data.

$$\Omega = \{w_1, w_2, w_3, w_4\}, \quad A = \{w_1, w_3\} \quad (70)$$

The synthetic state space and target.

$$E_1 = \{w_1, w_2, w_3\}, \quad E_2 = \{w_1, w_2, w_4\}, \quad S_2 = \{w_1, w_2\} \quad (71)$$

Two records narrow the admissible set.

$$P(A \mid E_1) = \frac{2}{3}, \quad P(A \mid E_1 \cap E_2) = \frac{1}{2}, \quad D(A, S_2) = \text{MAYBE} \quad (72)$$

Uniform-prior probabilities and the independent logical answer.

Every state and probability on this page is synthetic. This is a worked example, not a backtest.

24. An audit of certainty

What would have to change the answer?

An unresolved answer should be accompanied by a route to resolution whenever that route can be specified. For a finite model, one can enumerate which surviving states disagree with a proposed answer and ask what observation would distinguish them.

Let S intersect both A and its complement. Any observation map that is constant across an opposing pair fails to distinguish that pair. A resolving experiment must return different observations for every opposing pair that remains relevant to the realized result, or otherwise exclude one side through justified assumptions.

This condition does not guarantee that a practical instrument exists. It identifies the informational requirement. Some targets concern future events that have not yet occurred; some are poorly defined; some depend on inaccessible variables. The correct next step may be to refine the question rather than to add more requests to the same endpoint.

An audit can be organized into four questions: What is the proposition? Which record supports it? Which admitted counterexample survives? What new measurement would eliminate that counterexample? These questions are more useful than asking whether the interface feels sufficiently scientific.

The MAYBE answer is thus compatible with a serious research program. It records a boundary and invites a better experiment. What it refuses is the substitution of effort for evidence. Thirty pages of analysis may clarify the boundary without moving it.

$$\mathcal{K}(A, S) = \{(\omega_1, \omega_0) : \omega_1 \in S \cap A, \omega_0 \in S \cap A^c\} \quad (73)$$

The set of opposing admissible pairs.

$$h(\omega_1) \neq h(\omega_0) \quad \text{for all } (\omega_1, \omega_0) \in \mathcal{K}(A, S) \quad (74)$$

A sufficient condition for a deterministic experiment to separate the proposition on S .

$$|\mathcal{K}(A, S)| = |S \cap A| \cdot |S \cap A^c| \quad (75)$$

The number of unresolved opposing pairs in a finite model.

A useful question specifies what evidence could change its answer.

25. Assumptions and limitations

The model must remain visible

The mathematical results are conditional on a finite state space and a declared compatibility rule. They do not establish that the selected space contains every relevant real-world possibility. An omitted state can make an apparently decisive conclusion unsound outside the model.

The record model also idealizes source behavior. RPC services and indexers may lag, rate-limit, disagree, or return incorrect data. Validation catches specified failures, not every possible failure. This paper does not supply a cryptographic proof for every displayed field or an independent audit of the upstream services.

No empirical distribution of MAYBE returns, usage, holders, liquidity, or adoption is estimated. There is no fitted forecast, simulated profit claim, measured quantum system, or promised chain deployment. The examples are finite constructions designed to clarify inference.

The probability sections add assumptions explicitly. Bayesian priors, likelihoods, loss functions, and random-walk increments are illustrative objects. Their mathematical correctness does not validate their suitability for a real token. The information-theoretic references provide background, not an endorsement of the project.

Finally, a truthful interface can still be misunderstood. A dramatic visual style may lend authority to a limited observation. For that reason, the source, time, target, and missing fields belong in the main experience. The proper role of the visual language is to make the distinction between known and unknown easier to read.

valid deduction + unsupported premise $\neq$ validated prediction (76)

A valid derivation does not validate its premises.

model uncertainty $\neq$ measurement uncertainty $\neq$ future uncertainty (77)

Three different limitations that should not be collapsed.

This is a conceptual research-style document. Its formal results concern the stated finite model, not financial performance.

Appendix A. Proofs and edge cases

The small details that keep MAYBE honest

Proposition A.1 (persistence). Let T be a nonempty subset of S . If the answer to A on S is YES, the answer on T is YES. If the answer on S is NO, the answer on T is NO. Proof: subset inclusion is transitive, and disjointness is inherited by subsets. The nonempty condition excludes evidence conflict.

Proposition A.2 (resolution of every proposition). Every proposition on Ω is resolved by an admissible set S if and only if S is a singleton. Proof: a singleton belongs wholly to A or its complement. Conversely, if S contains two different states, choose A to contain one and not the other; that proposition receives MAYBE.

Proposition A.3 (coarsening). Suppose $h=f$ composed with g . If A is identifiable from h , then A is identifiable from g . Proof: each h -fiber is a union of g -fibers. A union of h -fibers is therefore a union of g -fibers. The converse can fail when f merges observations that disagree on A .

Edge case A.4. The empty admissible set is not assigned YES or NO. In classical universal quantification, both 'all states satisfy A ' and 'all states satisfy its complement' hold vacuously on the empty set. The operational rule avoids that ambiguity by checking consistency before resolution.

These proofs use only elementary finite-set reasoning. Their brevity is intentional. The weight of the argument rests on faithful modeling and record interpretation, not on making an elementary result look difficult.

$$\emptyset \neq T \subseteq S \subseteq A \implies T \subseteq A \quad (78)$$

Persistence of an affirmative answer.

$$|S| = 1 \iff \text{every } A \subseteq \Omega \text{ is resolved on } S \quad (79)$$

A singleton resolves every proposition on the state space.

$$h = f \circ g, \quad A = h^{-1}(B) \implies A = g^{-1}(f^{-1}(B)) \quad (80)$$

Identifiability under a more informative observation.

The empty-set guard is part of the answer rule, not an exception added after the result.

Appendix B. References and reproducibility

Sources, provenance, and what was not measured

[1] C. E. Shannon (1948). A Mathematical Theory of Communication. Bell System Technical Journal, 27, 379-423 and 623-656. Background for entropy and information. Primary-source access: Nokia Bell Labs, 'A Mathematical Theory of Communication.'

[2] D. Blackwell (1953). Equivalent Comparisons of Experiments. The Annals of Mathematical Statistics, 24(2), 265-272. DOI: 10.1214/aoms/1177729032. Background for comparing information structures in decision problems.

[3] E. T. Jaynes (1957). Information Theory and Statistical Mechanics. Physical Review, 106, 620-630. DOI: 10.1103/PhysRev.106.620. Background for inference under explicitly stated constraints; cited for context rather than physical claims about assets.

[4] Hyperliquid documentation. HyperEVM. Official documentation for the designated network and RPC interface. <https://hyperliquid.gitbook.io/hyperliquid-docs/for-developers/hyperevm>

[5] DEX Screener documentation. API reference. Source description for indexed token-pair observations. <https://docs.dexscreener.com/api/reference>

Reproducibility statement. The manuscript source contains every paragraph and numbered expression. The PDF and web reader use that same source. All worked examples specify their finite states or model parameters; no private dataset, price history, or undisclosed empirical sample is used.

Provenance statement. MAYBE Research names the project's conceptual writing series, not an academic affiliation. This document has not undergone journal peer review. The derivations are presented for inspection within their assumptions. The references do not endorse MAYBE or establish a token valuation.

source → {web reader, 30 page PDF}

(81)

One manuscript source, two reading formats.

Version 2.0 / 18 September 2026. References accessed during preparation. The source repository retains the PDF generator and validation procedure.

26. The final result

A theorem of warranted non-conclusion

Theorem 26.1 (MAYBE). Let Ω be a finite state space, let R be a record, and let $S(R)$ be a nonempty set of states compatible with that record. Let A be the target proposition. Suppose $S(R)$ contains one state in which A is true and another in which A is false. Then no decisive answer based only on R can be correct in every admissible state.

Proof. A rule using only R receives the same input in both states. If it returns YES, it is incorrect in the state outside A . If it returns NO, it is incorrect in the state inside A . Thus every decisive answer fails in at least one admissible state. The sound three-valued rule returns MAYBE.

The theorem does not declare all questions permanently unanswerable. A better observation can remove an opposing state. A sharper definition can identify a different proposition. A defensible assumption can narrow the model, provided the conclusion remains conditional on it. What cannot be done is to obtain the missing distinction by asserting that it has already been observed.

We have introduced states, records, fibers, compatible sets, probability families, posteriors, entropy, decisions, dynamics, missing values, temporal limits, and an operational protocol. The notation has become considerably more elaborate. The evidence for the unresolved target has not become stronger merely because the notation has.

The present may be observed. A future may be modeled. A conclusion must still be earned. Under the assumptions of the theorem, the complete answer is:

$$\exists \omega_1 \in S(R) \cap A, \quad \exists \omega_0 \in S(R) \cap A^c \quad (82)$$

The two admissible witnesses.

$$g(\omega_1) = g(\omega_0) = R, \quad \mathbf{1}_A(\omega_1) \neq \mathbf{1}_A(\omega_0) \quad (83)$$

The record cannot distinguish their opposing answers.

$$\mathcal{D}(A, S(R)) = \text{MAYBE} \quad (84)$$

The warranted conclusion.

MAYBE.