When Delayed Threshold Sensing Cannot Help: An Exact Scalar Acquisition-Timing Addendum

Prepared for Brian Theory
AI-assisted mathematical development and internal review

9 November 2026 Version 1.1

Abstract

We compare initial-time and delayed threshold acquisition for mandatory, zero-error prediction of sampled signs of a disturbed scalar system. Inputs are fixed in advance, all observations are paid, and each prediction precedes its transition. For $\dot{x} = -x + u + d$ on $[-2, 2]$, sampling interval $\log 2$, $|d| \le 1/10$, and input alphabet $\{-1, 0, 1\}$, we characterize the optimum for every positive finite horizon, nonempty initial set, and finite threshold library. Delayed and initial-time acquisition have the same worst-case probe budget: either both are infeasible or both equal $\lceil \log_2 K \rceil$, where K is the number of disturbance-independent output words. The key obstruction is a first-difference classification and an indistinguishable pair of zero-disturbance trajectories. Adding the single input value $1/2$ gives a two-step instance with positive-width initial uncertainty in which initial acquisition is impossible but one delayed probe is necessary and sufficient. Within the same ordered-target, equal-cost contract, a strict improvement between two finite optima is impossible. These are specialized geometric results and an integration-ready addendum, not a new general theory of active diagnosis or symbolic abstraction. External priority remains unresolved.

1 Purpose and relation to existing acquisition work

The parent technical report *From Observable States to Acquirable Predictive Models* distinguishes a mathematically sufficient predictive representation from the information a permitted experiment can acquire [1]. Its Section 7 already treats alternating observations and known actions through finite-horizon support planning. The fact that dynamics can change distinguishability is therefore not a new claim here. Nor do we reinterpret its topological query-admissibility rule as a physical safety condition.

This addendum isolates a narrower question: when does waiting help a fixed threshold library under a fixed control word and strict forecast deadlines? Our substantive candidate increment is the exact obstruction for one specified input alphabet, with a sharp counterexample after adding one input. Binary decision trees, indistinguishability arguments, interval reachability and support propagation are classical ingredients. The proposed placement is a technical addendum to the parent report, rather than a separate broad paper.

Symbolic control already provides rigorous finite abstractions of continuous systems [2, 3]. We instead prove the acquisition result directly on real intervals: no finite grid is required, and no approximate bisimulation or new abstraction-error theorem is claimed. Our earlier stochastic trace calculations are outside this note; bounded disturbance does not define a probability law. In particular the standard approximate trace bound studied in [4] is not a contribution of this addendum.

Active diagnosis already uses inputs to separate bounded-error output sets; Scott et al. give a sufficient disjoint-set condition in their Theorem 3 [5]. Adaptive and preset distinguishing experiments in finite-state machines are also classical [6, 7]. These comparisons matter: our controls are fixed in *both* acquisition classes, and our batch comparator permits adaptive queries at one physical instant, unlike a preset state-changing experiment. Ordered search trees likewise predate this work [8]; the elementary maximum-depth bound below is proved directly, not inferred from a weighted expected-cost result.

The nearby *Diagnosing While Repairing* companion [9] studies origin labels under coordinate resets and compares adaptive with fixed repair order. Its state-changing testing mechanism is an antecedent, not a theorem proved again here under new terminology. Here the distinction is sensing time under the same fixed inputs, adversarial continuous disturbances, and mandatory next-step sign reports. Every threshold is always permitted. The obstruction uses integer centers and contraction geometry, not support-dependent query legality or a choice of repair order. A targeted primary-source review found no exact duplicate of this classification among the examined sources; that bounded search does not establish originality. Public working-note release and a journal priority claim are separate decisions.

2 The complete operational contract

Let $X = [-2, 2]$, $H \geq 1$ be an integer, and $B \subseteq X$ be an arbitrary nonempty set. No interval, closure or finite-cardinality assumption is imposed on B. The continuous state satisfies

$$\dot{x}(t) = -x(t) + u(t) + d(t), \qquad |d(t)| \leq D = \tfrac{1}{10}. \tag{1}$$

The disturbance is an arbitrary measurable function subject to this bound. At $t_k = k \log 2$ the input for the next interval is u_k, held constant. The entire word $(u_0, \ldots, u_{H-1})$ is announced before acquisition and cannot depend on observations. Initially $u_k \in \{-1, 0, 1\}$; Section 5 adds $1/2$. Both alphabets keep X invariant, since the extreme vector fields point inward at its endpoints.

Variation of constants gives the exact sampled system

$$x_{k+1} = \tfrac{1}{2}(x_k + u_k) + w_k, \qquad w_k \in [-W, W], \quad W = \tfrac{1}{20}. \tag{2}$$

Every w_k in this interval is realized by the constant disturbance $d(t) = 2w_k$ on that sampling interval. Conversely every admissible disturbance induces such a w_k. Thus all finite sequences in $[-W, W]^H$ are admissible.

The target is $q(x) = \mathbf{1}\{x \geq 0\}$. A finite supplied library $E \subset \mathbb{R}$ contains thresholds $e_r(x) = \mathbf{1}\{x \geq r\}$, $r \in E$. The library may be empty. Every query costs the same $c > 0$, is noiseless, instantaneous and nonmutating, and is available at every checkpoint. There is no support-dependent admissibility restriction: an uninformative query is allowed but can be omitted. Multiple queries at one checkpoint see the same state. The possible-state set is conditioned on each purchased answer before the next decision. Neither $q(x_k)$ nor any state, crossing event or future output is observed for free.

An *online* policy may query at checkpoint k, then must report $q(x_{k+1})$ before transition k, for $k = 0, \ldots, H-1$. Query choices and waiting may depend only on prior paid answers, the known model and the clock. No between-checkpoint sensing is available. A *batch* policy queries only at checkpoint zero and reports the whole H-bit word before the first transition. "Batch" here permits adaptive query selection at that single instant; it does not mean a nonadaptive list of tests. There is no abstention. Every report must be correct for every $x_0 \in B$ and every disturbance sequence. Let

b_{on} and b_{bat} be the minimum hard worst-case number of probes, with infinity denoting infeasibility. Monetary optima are cb_{on} and cb_{bat}. We allow private randomization only if correctness is almost sure for each admissible execution and the probe cap is hard; it will not improve these optima.

3 Reachability and the information lower bound

For $1 \le n \le H$, put

$$c_n = -\sum_{i=0}^{n-1} 2^i u_i, \quad \rho_n = D(2^n - 1), \quad I_n = [c_n - \rho_n, c_n + \rho_n). \tag{3}$$

The point-reachable interval from x is exactly

$$2^{-n}(x - c_n) + [-D(1 - 2^{-n}), D(1 - 2^{-n})]. \tag{4}$$

Consequently its sign is always zero if $x < c_n - \rho_n$, always one if $x \ge c_n + \rho_n$, and can take both values otherwise. The lower endpoint of I_n is included because reaching zero yields label one; its upper endpoint is excluded because all reachable states there are nonnegative.

Lemma 1 (Physical obstruction)**.** *If $B \cap I_n \neq \varnothing$ for some n, no online policy can meet the contract, even if it is given the exact current state before each forecast. If B avoids every I_n, each $x \in B$ has a unique disturbance-independent word $v(x) \in \{0,1\}^H$.*

Proof. Fix $x \in B \cap I_n$ and write its attainable checkpoint-$(n-1)$ interval as $J = [L, U]$, allowing a singleton when $n = 1$. Write $u = u_{n-1}$. The union of next intervals contains both labels exactly when

$$L < -u + D \qquad \text{and} \qquad U \ge -u - D.$$

Equivalently $J \cap [-u - D, -u + D) \neq \varnothing$. Choose an attainable z in that intersection and a physical disturbance prefix reaching it. From this same prefix, the next interval $[(z+u)/2 - W, (z+u)/2 + W]$ has a negative member and a nonnegative member. The two continuations agree on all information before the deadline, including the exact current state, but require opposite next reports. Thus they cannot both be predicted without error. Formula (4) proves the second assertion. For randomized protocols, the two identical pre-deadline information laws likewise cannot support both reports with probability one. □

When the obstruction is absent, let $K = |v(B)|$. Each coordinate of $v(x)$ is nondecreasing in x, so the distinct words occur in order without recurrence and $1 \le K \le H + 1$. Define the initial signature $\sigma_E(x) = (e_r(x))_{r \in E}$.

Lemma 2 (Ordered batch optimum)**.** *Suppose B avoids all I_n. Batch acquisition is feasible if and only if v is constant on each fiber of σ_E restricted to B. If feasible, $b_{\mathrm{bat}} = \lceil \log_2 K \rceil$, and every online policy needs at least $\lceil \log_2 K \rceil$ probes in the worst case.*

Proof. Identical signatures produce identical batch transcripts. For sufficiency, sort the finite thresholds. Their signature atoms are ordered intervals, intersected with B, and each nonempty atom is word-homogeneous by hypothesis. The atoms for a given word occur consecutively. Between any two adjacent word groups there is therefore a supplied threshold separating the entire groups, including their endpoint conventions. If $K = 1$, stop without a query. Otherwise query the threshold between the $\lfloor K/2 \rfloor$-th and next groups, and recurse. This balanced tree has height $\lceil \log_2 K \rceil$. No compactness of B is needed.

A deterministic protocol with at most b paid binary answers has at most 2^b leaves. Fixed controls and the clock add no observed branches. Each leaf determines all reported bits, so K distinct required words need at least K leaves, also for an online protocol that interleaves reports and queries. For a randomized protocol, choose one state for each word and a fixed admissible disturbance path for each. The intersection of these finitely many probability-one correctness events has probability one. Fix a random tape in that intersection; the deterministic leaf bound applies. □

Corollary 3 (No finite-to-finite improvement)**.** *Whenever the robust-word condition and the ordered equal-cost threshold contract above hold and batch acquisition is feasible,* $b_{\mathrm{on}} = b_{\mathrm{bat}} = \lceil \log_2 K \rceil$. *This remains true for any fixed real input word for which the declared state-domain model is valid.*

Proof. The batch tree is an online policy and attains the online lower bound. □

This statement is about hard worst-case cost. Expected-cost objectives, unequal test costs, nonadaptive batch tests, different deadlines and adaptive controls are different problems.

4 Complete equality for the original alphabet

Theorem 4. *Under the contract of Section 2 with alphabet* $\{-1, 0, 1\}$*, for every positive finite* H*, nonempty* $B \subseteq [-2, 2]$ *and finite* $E \subset \mathbb{R}$*,*

$$b_{\mathrm{on}} = b_{\mathrm{bat}} = \begin{cases} \infty, & B \cap \bigcup_{n=1}^{H} I_n \neq \varnothing, \\ \infty, & \textit{otherwise, some initial signature fiber has different words}, \\ \lceil \log_2 K \rceil, & \textit{otherwise}. \end{cases}$$

Proof. The physical-obstruction case is Lemma 1. Otherwise, choose $x < y$ with different robust words and let n be the first differing position. Then

$$x < c_n - \rho_n, \qquad y \geq c_n + \rho_n. \tag{5}$$

Every c_n is an integer. For $n \geq 5$, $\rho_n \geq 31/10$, which is incompatible with (5) inside $[-2, 2]$. For $n = 4$, $\rho_n = 3/2$ and (5) implies $-1/2 < c_4 \leq 1/2$, hence $c_4 = 0$. A sum of distinct powers of two with coefficients in $\{-1, 0, 1\}$ can vanish only when all coefficients are zero: its greatest nonzero weight exceeds the sum of the smaller weights. All four inputs are therefore zero. Robust opposite fourth bits then imply opposite first bits, a contradiction.

For $n = 2, 3$, (5) permits only centers $-1, 0, 1$. Center zero again forces all preceding inputs zero and an earlier difference. For $s \in \{-1, 1\}$ the complete representations of $c_2 = s$ are

$$(-s, 0), \quad (s, -s),$$

and those of $c_3 = s$ are

$$(-s, 0, 0), \quad (s, -s, 0), \quad (s, s, -s).$$

The first representation repeats center s from step one; the second three-step representation repeats it from step two. A later band with the same center contains the earlier band, so robust opposite later bits were already opposite earlier. The only possible first-difference cases are

$$n = 1; \qquad n = 2 : \ (s, -s); \qquad n = 3 : \ (s, s, -s). \tag{6}$$

Now suppose $\sigma_E(x) = \sigma_E(y)$. Equivalently no threshold lies in $(x, y]$. In the last two cases of (6), $x < s < y$, and under zero disturbance both executions, before the n-th forecast deadline, obey

$$x_k = s + 2^{-k}(x - s), \qquad y_k = s + 2^{-k}(y - s), \quad 0 \le k < n. \tag{7}$$

Both states remain in $[x, y]$. Every supplied threshold thus gives the same answer on both paths at each available checkpoint. If a threshold equals x, both answers are one; a threshold equal to y would distinguish the initial states and is excluded. In case $n = 1$ only the initial checkpoint is available, so the conclusion holds directly. Induction on the paid transcript shows that adaptive query choice and waiting cannot distinguish the executions before the differing bit is due. Their reports must coincide although the required bits differ. Randomization also fails, by coupling the same private tape on both indistinguishable executions.

Thus a nonhomogeneous initial signature fiber makes online acquisition infeasible. If every fiber is homogeneous, Lemma 2 and Corollary 3 attain the claimed optimum. The finite classification proves the assertion for all horizons, not just the horizons used in the computational check. □

5 One added input permits a strict feasibility separation

Theorem 5. *Extend only the input alphabet to* $\{-1, 0, 1/2, 1\}$. *There is a two-step instance with* $b_{\mathrm{on}} = 1$ *and* $b_{\mathrm{bat}} = \infty$.

Proof. Choose announced inputs $(1/2, -1)$, library $E = \{1\}$ and

$$B = B_L \cup B_R, \qquad B_L = [21/20, 23/20), \quad B_R = [37/20, 39/20).$$

The bad bands are $[-3/5, -2/5)$ and $[6/5, 9/5)$. Both components avoid them: their robust words are respectively 10 and 11. At checkpoint zero e_1 is constantly one, so every batch transcript is identical and batch acquisition is impossible.

Without a probe, report the first bit one. The exact checkpoint-one sets are

$$F_{1/2}(B_L) = [29/40, 7/8), \qquad F_{1/2}(B_R) = [9/8, 51/40),$$

where $F_u(S) = \{(x + u)/2 + w : x \in S,\ |w| \le 1/20\}$. Now purchase e_1, which separates the two sets, and report its answer as the second forecast before input -1. The exact second-step intervals are

$$[-3/16, -1/80) \quad \text{and} \quad [1/80, 3/16),$$

so this policy succeeds for every admissible initial state and both disturbances. Zero probes cannot distinguish the two required words, proving optimality. The same impossibility applies to randomized zero-probe policies. □

The change is minimal in the *number of added input values*: with no additional input value a separation is impossible by Theorem 4, while one suffices. No minimum numerical perturbation, energy or metric distance is asserted. Waiting changes the informativeness of a fixed test, without free event detection, adaptive actuation or a missed prediction deadline. By Corollary 3, any strict improvement within this contract must have batch acquisition infeasible and online acquisition feasible; no strict improvement is possible when batch acquisition is feasible.

6 Reproduction, limits and release status

The companion standard-library Python program uses exact rational arithmetic. It checks all $1,092$ original-alphabet input words through horizon six, with initial-coordinate band endpoints and intervening midpoints. Among $7,182$ robust pairs with different words, the first-difference counts are $5,930$, 982 and 270 at steps one, two and three; every pair satisfies the classification and containment used above. An exact interval budget recursion for the extension checks that budget zero fails and budget one succeeds. Endpoint tests separately check ambiguity at initial $6/5$ and safe-positive equality at $9/5$. These finite checks corroborate, rather than replace, the analytical all-horizon and all-policy proofs. The package's reproducibility instructions and proof ledger identify the exact scripts, source hashes, review inputs and results.

All main results are on the exact continuous-state model. A grid-based overapproximation may introduce spurious infeasibility; its cost answer cannot be substituted without a property-specific exactness proof. No abstraction error is present in the displayed interval calculations. No intersample safety, noisy or delayed sensor theorem, empirical validation, topological acquisition constraint, cognitive property or consciousness conclusion is established. The disturbance is adversarial; almost-sure prediction under a specified stochastic law is a different contract.

The result is deliberately presented as a scoped addendum. A targeted comparison with the parent acquisition report, nearby self-diagnosis work and active-testing literature accompanies this version. Internal AI proof and hostile-review records do not constitute human peer review or establish priority. Public release is a separate centrally coordinated action; this local package itself is unpublished.

References

[1] Brian Theory. *From Observable States to Acquirable Predictive Models*. Revised technical report, 7 October 2026, Sections 4, 7 and Appendix B. https://www.b-theory.com/manuscripts-html/acquirable-predictive-models/.

[2] G. Pola, A. Girard and P. Tabuada. Approximately bisimilar symbolic models for nonlinear control systems. *Automatica* 44(10), 2508–2516, 2008. https://arxiv.org/abs/0706.0246.

[3] G. Reissig, A. Weber and M. Rungger. Feedback refinement relations for the synthesis of symbolic controllers. *IEEE Transactions on Automatic Control* 62(4), 1781–1796, 2017. https://doi.org/10.1109/TAC.2016.2593947.

[4] G. Bian and A. Abate. On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. *FoSSaCS*, 321–337, 2017. https://arxiv.org/abs/1701.04547.

[5] J. K. Scott, R. Findeisen, R. D. Braatz and D. M. Raimondo. Input design for guaranteed fault diagnosis using zonotopes. *Automatica* 50, 1580–1589, 2014, Theorem 3. https://doi.org/10.1016/j.automatica.2014.03.016.

[6] R. M. Hierons, G.-V. Jourdan, H. Ural and H. Yenigun. Checking sequence construction using adaptive and preset distinguishing sequences. *SEFM*, 2009, pp. 157–166. https://doi.org/10.1109/SEFM.2009.12.

[7] P. van den Bos and F. Vaandrager. State identification for labeled transition systems with inputs and outputs. Author preprint, 2019, Definition 20 and Theorem 35. https://arxiv.org/abs/1907.11034.

[8] T. C. Hu and A. C. Tucker. Optimal computer search trees and variable-length alphabetical codes. *SIAM Journal on Applied Mathematics* 21(4), 514–532, 1971. https://doi.org/10.1137/0121057. Publisher abstract/metadata consulted; full text not inspected in this audit.

[9] *Diagnosing While Repairing: Information Loss, Adaptive Order, and Noisy Readouts.* AI-assisted research development draft prepared for Brian Theory, companion manuscript v0.2.5, 8 November 2026; unpublished. Scope and adaptive-order sections compared with v0.2.4, 7 November 2026. A public persistent source location remains to be assigned by the author.