Rights and Authorship Record

Instant Time (INSTa): Decisive Benchmark Tests and Direct Theorem Validation under Delayed Observation

Fares Alabdali

4/20/2026

This document is preserved as the original authorship record of the INSTa research work.
This version is archived for authorship, documentation, and citation purposes.

Archived on Zenodo

DOI: 10.5281/zenodo.19855723

Instant Time (INSTa): Decisive Benchmark Tests and Direct Theorem Validation under Delayed Observation

Fares Alabdali
Independent Researcher

Keywords: *delayed observation; operational window; delayed estimation; benchmark validation; finite-window sufficiency; practical robustness*

ABSTRACT

I developed INSTa for systems operating under delayed observation, where the quantity available at decision time is not, in general, a directly usable present state, but a causally available delayed information set. I therefore based the tested control law on a finite near-past window and a forward operational estimate. In this paper, I present the decisive benchmark evidence for standard INSTa, and I report observer-corrected results only where they are needed to delimit the role of the INSTa form. I begin with primary baseline tests: a Kalman-like delayed-estimation benchmark, a corrected matched-model Smith predictor benchmark, and predictive-control baselines based on an early LQR proxy and a later true MPC implementation. In the delayed mass-spring-damper setting, I obtained a steady-state error reduction from 0.820327 to 0.128582 against the Kalman-like baseline, whereas Smith predictor outperformed standard INSTa in the corrected matched-model-fixed-delay test. I then place these results within a broader suite covering mixed noise with parameter drift, abrupt delay jumps, actuator saturation, and missing observations. Finally, I report eight theorem-linked validation tests of the governing INSTa formulation: equation sanity, finite-window sufficiency, deviation-state robustness with a general robustness-bound interpretation, Theorem-1 validation, direct closed-loop stability validation, corrected CRLB validation, comparison with the theoretical optimum, and a fixed-regime threshold-condition test. Across this record, I found strong support for INSTa in delayed-estimation comparison, abrupt delay-jump response, constrained actuation, finite-window sufficiency, robustness-bound validation, theorem-level validation, direct stability validation, and corrected CRLB agreement. At the same time, I retained the main limitation in full view: Smith predictor remained stronger in the corrected matched-model fixed-delay benchmark.

INTRODUCTION

I developed INSTa from an operational premise: once observation delay is admitted explicitly, the object available at decision time is no longer a directly usable present state. What is available is delayed information. I therefore treat the control problem as one of reconstructing an operational state from a finite near-past window and projecting that reconstruction to the horizon at which action must take effect.

In this paper I do not present the full archival record of every intermediate run. I present the benchmark evidence that I regard as decisive for the INSTa formulation. My purpose is not to argue that standard INSTa dominates every classical baseline. My purpose is to show where the formulation is genuinely strong, where it remains only competitive, and where stronger classical structure persists. I organize the evidence in three stages. First, I report primary baseline tests against strong comparators in delayed estimation, delay compensation, and predictive control. Second, I report the broader empirical suite, because a delayed-observation framework should not be judged from a single convenient setting. Third, I report eight theorem-linked validation tests of the governing INSTa formulation: equation sanity, finite-window sufficiency, deviation-state robustness with a general robustness-bound interpretation, Theorem-1 validation, direct closed-loop stability validation, corrected CRLB validation, comparison with the theoretical optimum, and a fixed-regime threshold-condition test.

GOVERNING EQUATIONS AND TESTED FORM

I used the discrete-time formulation throughout the validated numerical record. I wrote the delayed observation model as: $W_k^{(L)} \to \hat{x}^{op}_{k+\ell}|W_k \to u_k$

- Delayed observation

$$y_k = Cx_{k-d} + v_k$$

- Definition of the information window

$$W_K^{(L)} = \{y_k, y_{k-1}, \ldots, y_{k-L+1}, u_{k-d-L+1}, \ldots, u_{k-1}\}$$

- Operational estimate from the window

$$\hat{x}^{op}_{k+\ell}|W_k = P_{L,\ell}(W_k^{(L)})$$

- Then the control

$$u_k = -K\,\hat{x}^{op}_{k+\ell}|W_k$$

These equations define the standard INSTa form on which I concentrate in this paper. When I needed to show whether a gain came from the standard formulation or from an observer-corrected extension, I reported the observer-corrected result explicitly and treated it only as a boundary marker around the claim.

I therefore interpret every test in this paper through one question: when I reconstruct an operational estimate from delayed information and project it forward, does that structure remain useful against strong comparators and under difficult delayed-information regimes?

PRIMARY BASELINE TESTS

I begin with primary baseline tests because they place INSTa immediately against comparators that matter technically: delayed-estimation baselines, a corrected delay-compensation baseline, and explicit predictive-control baselines.

Test	Setting	Baseline result	INSTa result	Interpretation
Kalman-like delayed estimation	MSD; $\tau = 0.25$, $\Delta = 0.06$, $\sigma = 0$	SSE = 0.820327	SSE = 0.128582	Strong positive INSTa result
Corrected matched model Smith predictor	MSD; $\tau = 0.25$, $\Delta = 0.06$, $\sigma = 0$	Steady-state error = 0.558929 IAE =7.971349	Steady-state error = 0.664755 IAE =8.440066	Negative result for INSTa
Predictive baselines	Early LQR proxy and later true MPC	LQR proxy unstable; true MPC IAE = 14.23-14.45	INSTa IAE = 1.41-3.89 across the retained sweep	The predictive baselines did not dominate in the tested forms

Against the Kalman-like delayed-estimation baseline, I obtained one of the clearest positive results in the entire project. Under the validated benchmark setting, INSTa reduced the steady-state error from 0.820327 to 0.128582.

I also retained a clear negative result. In the corrected matched-model fixed-delay benchmark, Smith predictor outperformed standard INSTa in steady-state error, RMSE, IAE, and oscillation envelope. I preserved that result because I do not regard a delayed-observation framework as credible unless its limitations remain visible in the main record.

Finally, I tested predictive-control baselines. The early LQR-based proxy became unstable and was not a credible long-term comparator in that form. I then ran a true MPC-style receding-horizon implementation. It remained bounded, but it did not outperform the strongest delayed-window methods in the retained implementation.

EXTENDED EMPIRICAL BENCHMARKS

I next tested INSTa across the broader empirical suite that defined the mature phase of the project. I report the INSTa numbers directly and I mention the observer-corrected variant only where its presence changes the interpretation of the standard result.

Benchmark	INSTa result	Comparator context	Interpretation
Time-varying delay + parameter drift + mixed noise	Average IAE = 2.255; Average RMSE = 0.204 all runs stable	Present = 2.189; Smith = 2.357; Kalman-like = 2.206	Mixed; INSTa remained competitive but not best
Switching system + abrupt delay jumps	Average IAE = 4.650; Average RMSE= 0.302	Smith = 4.964; Present = 5.733; Kalman-like = 5.736	Strong positive standard INSTa result
Actuator saturation + rate limit	Average IAE = 6.784; Average RMSE= 0.396	Smith = 6.886; Present = 6.913; Kalman-like = 6.910	Positive by a small but repeatable margin
Dropout / missing observations	Average IAE = 5.131; Average RMSE= 0.336	Smith = 4.992; Present = 5.706 ; Kalman-like = 5.818	INSTa remained strong; observer-corrected variant was best

I treat the mixed-noise and parameter-drift benchmark as a mixed result. All methods remained stable. INSTa remained competitive, but it was not the best performer on the main error metric.

I treat the switching benchmark as one of the strongest positive INSTa results in the project. In that test, INSTa outperformed Smith, Present, and the Kalman-like baseline. The observer-corrected variant was only marginally lower.
I treat the actuator-constrained benchmark as a second positive INSTa result. The gain was small, but it repeated under realistic saturation and rate-limit conditions.

I retain the missing-observation benchmark because it is conceptually aligned with the delayed-information logic of INSTa. INSTa remained strong there, although the observer-corrected form was stronger.

EQUATION SANITY

Before drawing theorem-level conclusions, I checked the internal consistency of the governing formulation itself. I verified dimensions, time indices, causality, boundedness, and edge-case behaviour. Across the reported checks, the equation set remained operationally coherent.

REPRESENTATIVE EQUATION-SANITY RESULTS

Check	Recorded outcome	Status
Dimensions and symbol consistency	All tested cases passed; dimensions = True	Passed
Causality and time-index consistency	causality = True	Passed
boundedness under the reported checks	bounded = True	Passed
Edge-case screening	All tested cases passed	Passed

I therefore treat the equation-sanity stage as passed: the formulation remained dimensionally consistent, causally valid, and operationally well-posed under the reported checks.

FINITE-WINDOW SUFFICIENCY

I then tested the finite-window sufficiency claim directly. The central question was whether the reconstruction gap decreases systematically as the near-past window length L increases, and whether the decay is well captured by an exponential fit.

REPRESENTATIVE FINITE-WINDOW SUFFICIENCY RESULTS

Criterion	Recorded outcome	Status
Monotonic decrease of the reconstruction gap with increasing L	7/7 downward steps	Passed strongly
Exponential fit parameter rho	0.790032, 0.786220, 0.841310, 0.853162	Supportive
Best log-fit quality	log-fit R^2 = 0.9861, 0.9818	Supportive

I therefore regard the finite-window sufficiency result as strong: the gap decreased systematically with L, and the reported exponential fits remained numerically tight.

DEVIATION-STATE ROBUSTNESS / GENERAL ROBUSTNESS BOUND

I also tested whether the aggregate perturbation scale Xi tracks practical degradation in the deviation state, and whether the resulting behaviour is compatible with the general robustness-bound interpretation used in the project. In this part of the project, I treated the robustness statement as conditional and bounded rather than global.

$$A_k = \bar{A} + \Delta A_k, \qquad B_k = \bar{B} + \Delta B_k$$

$$x_{k+1} = A_k x_k + B_k u_k + w_k, \qquad y_k = C x_{k-d_k} + v_k$$

$$\|z_k - \hat{z}_k\| \le \varepsilon_L + c_d \delta_d + c_j j_d + c_\theta \delta_\theta + c_v \bar{v} + c_w \bar{w}$$

$$\Xi = a_L \varepsilon_L + a_d \delta_d + a_j j_d + a_\theta \delta_\theta + a_v \bar{v} + a_w \bar{w}$$

$$\|s_k\| \le c\rho^k \|s_0\| + \gamma \Xi$$

REPRESENTATIVE ROBUSTNESS-BOUND RESULTS

Criterion	Recorded outcome	Status
Xi monotonicity across ordered scenarios	5/5	Passed strongly
Fit quality between the deviation-state bound and observed degradation	R^2 = 0.9961 and 0.9965	Passed strongly
Reported bound coefficient	γ = 0.304551	Supportive
Coverage of the reported ordered scenarios	Full coverage	Passed

I therefore treat the robustness-bound result as strong: once the deviation state was used, the aggregate perturbation scale tracked degradation almost linearly across the ordered scenarios.

THEOREM-1 VALIDATION

I also ran the initial theorem-level validation in which I checked whether the practical condition held outside trivial regimes. In the reported grid, rank(O_L) remained equal to 2, and the practical condition was reported as True in every listed run. What matters most here is the pattern: increasing L strongly reduced the reconstruction burden and the tail behaviour, especially at larger delays.

REPRESENTATIVE IMPROVEMENT WITH INCREASING WINDOW LENGTH

Delay d	L = 2: ε_RMS / tail_max / final norm	L = 6: ε_RMS / tail_max / final norm	Practical?
1	0.4233 / 0.6866 / 0.373970	0.1569 / 0.1061 / 0.059058	True in both cases
2	1.6117 / 2.6654 / 2.665354	0.1850 / 0.1036 / 0.103569	True in both cases
4	7.9269 /26.1967/26.196686	0.2192 / 0.1667 / 0.026106	True in both cases
6	3.4981 /11.3052/10.181042	0.2299 / 0.2086 / 0.049593	True in both cases
8	12.7809/37.7428/19.627460	0.2197 / 0.2005 / 0.200509	True in both cases

I therefore interpret this test as strong support: the theorem-level practical condition did not only survive easy cases, but remained numerically meaningful in harder delayed regimes.

DIRECT CLOSED-LOOP STABILITY VALIDATION

I ran the exact discrete-time stability check for the core INSTa closed-loop form across d = 1, 2, 4, 6, 8 and ℓ = 0, 1, 2, 3. In every reported case, the simulation remained stable and the spectral radius stayed strictly below one. The key point is not only that $\rho\left(A_{(cl)}\right)$< 1 held, but that the time-domain simulations were consistent with that spectral result.

REPRESENTATIVE NUMERICAL SUMMARY

Horizon ℓ	Spectral radius	Final norm across d	Max norm across d	Interpretation
0	0.965142	0.000818 to 0.000888	0.8984 to 0.9406	Stable in all tested delays

1	0.963213	0.000687 to 0.000731	0.8851 to 0.9281	Stable in all tested delays
2	0.961367	0.000542 to 0.000562	0.8742 to 0.9167	Stable in all tested delays
3	0.959607	0.000391 to 0.000394	0.8656 to 0.9067	Stable in all tested delays

From these runs, I conclude that the corrected discrete-time INSTa closed-loop formulation is numerically stable in the tested grid, and that the stability evidence is direct rather than merely qualitative.

CRLB VALIDATION

The CRLB test remains one of the strongest theorem-linked results in the project. I first checked model consistency and obtained residuals essentially equal to zero. I then compared the sample covariance with the corrected CRLB across the reported grid. The key numerical fact is that the ratio cov/CRLB stayed extremely close to 1.

REPRESENTATIVE CRLB ROWS

Case	Trace (CRLB)	trace(sample covariance)	cov/CRLB	Interpretation
d=1,L=2,ℓ=0 σ=0.010	0.076566	0.076213	0.995398	Very close agreement
d=1,L=3,ℓ=0 σ=0.010	0.018970	0.019089	1.006254	Very close agreement
d=1,L=4,ℓ=0 σ=0.010	0.007527	0.007697	1.022600	Still close to unity
d=1,L=5,ℓ=0 σ=0.010	0.003737	0.003740	1.000740	Near-identity match
d=1,L=6,ℓ=0 σ=0.010	0.002123	0.002144	1.010150	Near-identity match

Across the reported rows, cov/CRLB remained approximately in the band 0.995 to 1.023. I regard that as strong evidence that the corrected delayed-window formulation is statistically consistent.

COMPARISON WITH THE THEORETICAL OPTIMUM

I also compared INSTa with the exact Bayesian optimum, that is, the LMMSE/MMSE benchmark for the same delayed estimation problem. I did not find that INSTa became exactly optimal. What I found instead was more important for the scientific claim: as the window length L increased, the gap shrank sharply, and INSTa became near-optimal in the informative-window regime.

REPRESENTATIVE INSTA / OPTIMUM RATIOS

Representative case	MSE OPT_EMP	MSE INSTa_ EMP	INSTa / OPT	Interpretation
d=1,L=2,ℓ=0 σ=0.020	0.186303	0.307031	1.648016	Clearly above optimum
d=1,L=4,ℓ=0 σ=0.020	0.027957	0.029929	1.070525	Gap already much smaller
d=1,L=6,ℓ=0 σ=0.005	0.000542	0.000543	1.001659	Near-optimal
d=4,L=6,ℓ=3 σ=0.020	0.007892	0.007974	1.010390	Near-optimal
d=8,L=6,ℓ=3 σ=0.020	0.007653	0.007820	1.021784	Still close to optimum

This is the correct interpretation I draw from these results: standard INSTa is not theoretically optimal in general, but the distance from the optimum can become very small when the near-past window is sufficiently informative.

FIXED-REGIME THRESHOLD-CONDITION TEST

In its earliest compact form, I expressed the threshold idea as:

$$\varepsilon_L < \varepsilon_{\mathrm{crit}}(\ell)$$

I include the strongest threshold result, namely the fixed-regime version rather than the mixed-regime version.

$$\varepsilon_L < \varepsilon_{\mathrm{crit}}(\ell;\, d,\, \sigma)$$

Finally, I do not claim that a single universal threshold was recovered across all regimes. I claim something more precise: within fixed regimes, the threshold interpretation was numerically meaningful, particularly in the low-and moderate-noise cases.

REPRESENTATIVE FIXED-REGIME THRESHOLD RESULTS

Regime	EPS_crit	Balanced accuracy	Accuracy	Interpretation
d=1, σ=0.001,ℓ=0	0.003283	0.8333	0.8750	Strong fixed-regime support
d=1, σ=0.010,ℓ=4	0.039908	0.9167	0.8750	Very strong fixed -regime support
d=2, σ=0.010,ℓ=2	0.039853	0.9167	0.8750	Very strong fixed--regime support
d=4, σ=0.010,ℓ=0	0.041216	0.9167	0.8750	Very strong fixed-regime support

d=4, σ=0.020,ℓ=0	0.062371	0.5000	1.0000	All runs accepted; threshold non informative

I therefore do not present the threshold as a universal scalar law. I present it in that the fixed-regime threshold condition is supported mainly in the low-and moderate-noise regimes and remains the scientifically defensible version of the result.

CONCLUSION

I conclude that INSTa should be read as an operational framework for decision and control under delayed observation. Its central move is simple: I replace direct action on a supposedly accessible present state with action based on a finite near-past information window and a forward operational estimate. In the validated benchmark record, I found decisive positive evidence for INSTa in delayed-estimation comparison, abrupt delay-jump response, and constrained actuation. I also found that the governing formulation remained numerically strong in equation sanity, finite-window sufficiency, deviation-state robustness with a robustness-bound interpretation, theorem-level validation, direct closed-loop stability validation, corrected CRLB agreement, near-optimality under informative windows, and conditionally supported fixed-regime threshold behaviour. Rather, I make a bounded and operational claim: INSTa is a coherent delayed-observation framework whose strongest support appears in delayed-estimation comparison, abrupt delay-jump response, constrained actuation, and the theory-linked numerical tests reported here. I therefore regard the present result as operational, strong, and delimited rather than universal.

DECLARATIONS

- *Data accessibility/This article is theoretical and computational in nature. Supporting scripts and summary outputs accompany the submission as supplementary material.*
- *Author contributions/Fares Alabdali. I conceived the study, developed the theoretical framework, designed and ran the simulations, analysed the results, and I wrote the manuscript.*
- *Acknowledgements/ None.*
- *Competing interests/ I declares no competing interests.*
- *Funding/This research received no specific grant from any funding agency in the public, commercial, or not for-profit sectors.*

References

- Fares Alabdali. 2026. Instant Time (INSTa). Research Gate preprint. Available at: https://www.researchgate.net/publication/399854095_Instant_Time_INSTa
- *Smith, O. J. M. 1957 Closer Control of Loops with Dead Time. Chemical Engineering Progress 53(5), 217–219.*
- Fares Alabdali. 2026. Instant Time (INSTa), Part II. Research Gate preprint Available at: https://www.researchgate.net/publication/404232666_Insta_Time_INSTa_Part_II
- *Kalman, R. E. 1960 A New Approach to Linear Filtering and Prediction Problems. Journal of Basic Engineering 82(1), 35–45.*
- *Bar-Shalom, Y. 2000 Update with out-of-sequence measurements in tracking: exact solution. Proceedings of SPIE 4048.*
- *Normey-Rico, J. E. & Camacho, E. F. 2007 Control of Dead-time Processes. Springer.*
- *Rawlings, J. B., Mayne, D. Q. & Diehl, M. 2022 Model Predictive Control: Theory, Computation, and Design, 2nd edn. Nob Hill Publishing.*
- *Kay, S. M. 1993 Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory. Prentice Hall.*
- *Van Trees, H. L. 2001 Detection, Estimation, and Linear Modulation Theory. Wiley.*