

One Number Sets the Limit of Forecasting

— Each Extra Day Costs **1.5874** Times the Initial Accuracy — Observe **10** Times More Precisely and You Gain
Only **4.98** Days — [Paper 289]

Yuuki YAMAGISHI (山岸 勇輝)
Yamagishi Complexity Torus Lab, Japan
August 2026

— *From thinking to syncing.* —

Abstract

That a weather forecast cannot reach beyond a certain horizon is due neither to missing equations nor to slow computers. This paper asks what sets the limit—the answer is one number, the error doubling time τ_d . No new mathematical theorem and no new law is claimed. First, the price per day is a fixed factor. With $\tau_d = 1.5$ days, each extra day costs **1.5874** times the initial accuracy (Section 2). Second, this is the core of the paper. Going from 14 to 21 days costs **25.40** times; to 30 days, **1625.5** times; to 60 days, **1.70×10^9** times (Section 2). Third, read backwards, the return is logarithmic. Observing **10** times more precisely gains only **4.98** days (Section 3). Fourth, even 10^9 times gains only **44.85** days (Section 3). Fifth, the familiar “about two weeks” comes from here. If the initial error is 10^{-3} of saturation, the forecastable span is **14.95** days (Section 4). Sixth, the separator is τ_d itself. At $\tau_d = 1.0$ day the same extension costs **65536** times; at 2.5 days only **84.4**—everything rides on one number (Section 5).

Scope of this paper (scope note): No new mathematical theorem and no new law is claimed——exponential error growth, the doubling time, the limit of predictability, and the value $\tau_d \approx 1.5$ days are all standard. We do not build meteorology——all we use is one exponential and its inverse. We do not discuss chaos——the Lorenz equations, attractors, and bifurcations are not treated at all. We do not discuss numerical weather prediction——grid resolution, parameterisation, and data assimilation are not treated. We do not say the error grows exactly exponentially—— $e^{\lambda t}$ holds only while the error is small, and growth stops near saturation. The computations here are confined to the linear-growth regime. We assert no value for τ_d ——1.5 days is a representative value widely used in the literature, and it moves from about 1 to 2.5 days with season, region, and variable. Section 5 shows the size of that dependence itself. We do not say there is a single exponent——the real atmosphere has different growth rates at different scales, with smaller eddies growing faster. A single τ_d is a crude approximation. We do not deny that forecasts improve——forecasts have in fact grown longer. What this paper says is only that the growth is logarithmic, not that improvement is pointless. Relation to earlier papers: Paper 253 showed that time is one-dimensional because prediction demands it, not because a law says so——this paper turns how far that demand can be met into a number. Paper 195 separated “stable” into six words——that paper is a classification of stability; this one is a time scale of predictability, the same hyperbolicity as material with a different question. Paper 190 measured “rare” on a logarithmic scale——the return here is likewise logarithmic. Paper 266 showed that the premise of the sampling theorem is never met——“knowing the initial state exactly” here is likewise a premise never met, the same figure. What is added is writing the price per day as the fixed factor 1.5874, computing the accuracy needed for $14 \rightarrow 21 \rightarrow 30 \rightarrow 60$ days as $25.40 / 1625.5 / 1.70 \times 10^9$, writing backwards that 10 times the observation gains only 4.98 days, and sweeping τ_d from 1.0 to 2.5 to show the answer moving from 65536 to 84.4.

1. Notation, and One Number

Setting: A small error ε_0 in the initial state grows as

$$\varepsilon(t) = \varepsilon_0 e^{\lambda t} = \varepsilon_0 2^{t/\tau_d}$$

Here τ_d is the time for the error to double, taken as about 1.5 days for the atmosphere.

Check 1.1:

Quantity	Value
Doubling time τ_d	1.5 days
$\lambda = \ln 2 / \tau_d$	0.462098 per day

Question: To extend a forecast, what must be improved, and by how much?

Observation 1.2

The span T over which a forecast holds runs until the error reaches a useless size $\varepsilon_{\max}$. Solving $\varepsilon_0 2^{T/\tau_d} = \varepsilon_{\max}$,

$$T = \tau_d \log_2 \left(\frac{\varepsilon_{\max}}{\varepsilon_0} \right)$$

——the initial accuracy enters only inside a logarithm.

2. The Core: 1.5874 Times per Day

Check 2.1: The factor needed to extend T by one day is $2^{1/\tau_d}$.

Quantity	Value
Accuracy factor per day, $2^{1/1.5}$	1.5874

Observation 2.2

1.5874 per day, and the factor is the same wherever the day is added. The day at 14 and the day at 60 cost the same 1.5874——but they compound.

Check 2.3: We compound from a base of 14 days.

Extension	Doublings $(T - 14)/\tau_d$	Accuracy factor required
14 → 16 days	1.3333	2.5198
14 → 18 days	2.6667	6.3496
14 → 21 days	4.6667	25.398
14 → 24 days	6.6667	101.59
14 → 28 days	9.3333	645.08
14 → 30 days	10.6667	1625.5
14 → 45 days	20.6667	1.6645×10^6
14 → 60 days	30.6667	1.7045×10^9
14 → 90 days	50.6667	1.7873×10^{15}

Observation 2.4

This is the core of the paper. One more week costs **25.4** times, two weeks **645**, six weeks **1.7 billion**—the extension adds while the price multiplies.

Observation 2.5

14 → 16 days costs a mere 2.52. The same “two more days” from 88 to 90 requires a starting accuracy already 10^{15} times better—the factor is constant; the starting point is not.

3. Read Backwards: 10 Times the Observation, 4.98 Days

Setting: In practice the question runs the other way—make the observing network N times more precise, and how many days does the forecast gain? The answer is $\tau_d \log_2 N$.

Check 3.1:

Accuracy factor N	Days gained, $\tau_d \log_2 N$
2	1.5000
10	4.9829
100	9.9658
1000	14.9487
10^6	29.8974
10^9	44.8460

Observation 3.2

Observing 10 times more precisely gains only **4.98** days. Halving the error buys exactly $\tau_d = 1.5$ days—which is the definition itself.

Observation 3.3

Even 10^9 times gains **44.85** days. Put an instrument in every cubic metre on Earth and you still do not reach two months—because the return is logarithmic.

Observation 3.4

Paper 190 wrote that “rare” is measured on a logarithmic scale. Here the relation between resource and result is itself logarithmic—the input multiplies while the output only adds.

Observation 3.5

A point to write honestly—forecasts have in fact grown longer. What this paper describes is the shape of that growth, not that improvement is pointless—five days of improvement is the fruit of the enormous labour of a **10-fold** gain in accuracy.

4. “About Two Weeks” Comes from Here

Setting: The forecastable span is set by how small the initial error ε_0 is against the saturation error $\varepsilon_{\max}$.

Check 4.1:

$\varepsilon_0/\varepsilon_{\max}$	Forecastable span $\tau_d \log_2(\varepsilon_{\max}/\varepsilon_0)$
10^{-2}	9.9658 days
10^{-3}	14.9487 days
10^{-4}	19.9316 days
10^{-6}	29.8974 days

Observation 4.2

If the initial error is 10^{-3} of saturation, the forecastable span is **14.95** days. The familiar “the limit of forecasting is about two weeks” comes from this single line—that two-week figure is fixed by just two things: τ_d and the initial accuracy.

Observation 4.3

What acts here is not the complexity of the equations. Solve the Navier--Stokes equations perfectly and you still exhaust the forecast in 14.95 days at 10^{-3} initial accuracy—the limit lies on the input side, not the solver side.

Observation 4.4

Paper 266 showed that the premise of the sampling theorem is never met. “Knowing the initial state exactly” is likewise a premise never met—the theorem correct and the premise out of reach, the same figure.

5. Separator: Everything Rides on τ_d

Observation 5.1

Separator——what must be moved for the answer to change?Only τ_d .

Check 5.2: We sweep τ_d and look at the answer to the same question.

τ_d [days]	Accuracy factor for 14 $\rightarrow$ 30 days	Days gained by 10 $\times$ observation
1.0	6.5536×10^4	3.3219
1.5	1.6255×10^3	4.9829
2.0	2.5600×10^2	6.6439
2.5	8.4449×10^1	8.3048

Observation 5.3

Moving τ_d from 1.0 to 2.5 takes the required accuracy from **65536** to **84.4**, a factor of 776.The whole enterprise of forecasting rides on this one number——so measuring τ_d matters as much as extending the forecast.

Observation 5.4

A point to write honestly——the real atmosphere has no single τ_d .Growth rates differ by scale, with smaller eddies growing faster——the single τ_d here is a crude approximation, and what can be shown is the structure, not the value.

Observation 5.5

Paper 195 separated “stable” into six words.That τ_d is finite at all is due to hyperbolic instability, but this paper does not enter that classification——that paper separates by what is perturbed; this one only turns the speed of growth into a number.

Observation 5.6

Paper 253 showed that time is one-dimensional because of the demand that prediction be possible, not because a law says so. This paper is its continuation—that demand is met only out to 14.95 days.

6. Roll Call and Closing

1. $\lambda = 0.462098$ per day (Check 1.1): From $\tau_d = 1.5$ days.
2. 1.5874 per day (Check 2.1): The same factor wherever added.
3. 14 → 21 days costs 25.398 (Check 2.3): The price of one week.
4. 14 → 30 days costs 1625.5 (Check 2.3): Two and a half weeks.
5. 14 → 60 days costs 1.70×10^9 (Check 2.3): 1.7 billion for six weeks.
6. $10\times$ observation gains 4.9829 days (Check 3.1): The return is logarithmic.
7. $10^9\times$ gains 44.846 days (Check 3.1): Not even two months.
8. 10^{-3} gives 14.9487 days (Check 4.1): The source of “about two weeks.”
9. $\tau_d = 1.0$ costs 65536 (Check 5.2): 2.5 costs 84.4.
10. A factor of 776 (Observation 5.3): All of it on one number.

Closing: What sets the limit of forecasting is neither the equations nor the computers, but one number, the error doubling time τ_d . At $\tau_d = 1.5$ days, each extra day costs 1.5874 times the initial accuracy—the factor is the same wherever the day is added, but the extension adds while the price multiplies, so one week costs 25.4, two weeks 645, six weeks 1.7 billion. Read backwards, observing 10 times more precisely gains only 4.98 days, and even 10^9 times gains 44.85. The familiar “about two weeks” comes from this one line—14.95 days at an initial error of 10^{-3} of saturation. One thing separates them— τ_d itself. At 1.0 day the same extension costs 65536; at 2.5 days, 84.4. A factor of 776 arises from a single number. So the work of extending forecasts and the work of measuring τ_d carry the same weight.

References

1. Lorenz, E. N. (1963). Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 20(2).
2. Lorenz, E. N. (1969). The Predictability of a Flow Which Possesses Many Scales of Motion. *Tellus*, 21(3).
3. Palmer, T. N. and Hagedorn, R. (eds.) (2006). *Predictability of Weather and Climate*. Cambridge: Cambridge University Press.
4. Zhang, F. et al. (2019). What Is the Predictability Limit of Midlatitude Weather? *Journal of the Atmospheric Sciences*, 76(4).

5. Bauer, P., Thorpe, A. and Brunet, G. (2015). The Quiet Revolution of Numerical Weather Prediction. *Nature*, 525(7567).
6. Yamagishi, Y. (2026). “Rare” Is Measured on a Logarithmic Scale [Paper 190]. Zenodo. 10.5281/zenodo.22106182
7. Yamagishi, Y. (2026). “Stable” Is Six Separate Words, and What Separates Them Is What You Perturb [Paper 195]. Zenodo. 10.5281/zenodo.22107170
8. Yamagishi, Y. (2026). Time Is One-Dimensional Because Prediction Demands It, Not Because a Law Says So [Paper 253]. Zenodo. 10.5281/zenodo.22161926
9. Yamagishi, Y. (2026). The Sampling Theorem Is Correct and Its Premise Is Never Met [Paper 266]. Zenodo. 10.5281/zenodo.22169765

— *From thinking to syncing.* —

Yuuki YAMAGISHI (山岸 勇輝)

© 2026 Yamagishi Complexity Torus Lab, Japan · Licensed under Creative Commons Attribution 4.0 International (CC BY 4.0)