

The Leaves-in-the-Wind Oracle:

A Local Directional Method for Time Series Forecasting*

Dimitrios Thomakos

Cybernetics and Artificial Intelligence Laboratory

Department of Business Administration

National and Kapodistrian University of Athens

February 15, 2026

Abstract

This paper proposes a novel, nonparametric forecast function for univariate time series, inspired by the image of a diviner reading the direction and agitation of leaves in the wind. The “Leaves-in-the-Wind Oracle” relies only on local first and second differences—interpreted as velocity and acceleration—and a simple measure of directional coherence. Using these ingredients, the method constructs regime-dependent forecasts that extrapolate the local trajectory in a kinematic fashion. This new forecast function has two interesting properties: first, it can be expressed as a heterogeneous autoregressive model for integrated time series; second, a hybrid version of the method extends the well known Theta method and give it a new and critical interpretation. An ensemble version aggregates multiple window configurations and combines individual and inter-model uncertainty. The method is estimation-free, interpretable, and easy to implement, and it offers a new local-directional perspective distinct from existing local polynomial and time-series momentum approaches.

1 Introduction

Short-horizon time series forecasting often relies either on parametric models, such as exponential smoothing, ARMA and state-space formulations, the Theta method, or on flexible nonparametric procedures such as local polynomial regression and Gaussian processes. At the same time, financial and economic practice has explored directional rules based on past signs and cumulative returns, usually under the heading of time-series momentum. This paper introduces a new forecasting method that occupies a middle ground in a simple yet innovative way: it uses only local directional information, but does so in a structured, kinematic way that explicitly

*This is a draft for circulation and comments, I have not included references yet. Please do not quote without permission. The paper is part of AI-assisted research conducted in Perplexity using Claude Sonnet 4.5 Thinking, all prompts are available on request. Any remaining errors are mine. Python code can be found in my github repository <https://github.com/dthomakos/Prognostikon-Code> and they are three files `leaves_oracle_fred_wrapper.py`, `run_experiment.py` and `synthetic_dgp_generators.py`.

models velocity, acceleration, and coherence. The end result not only is simple and interpretable but possesses some desired a priori characteristics.

The motivating metaphor is a seer who predicts the wind not from global climatology, but by studying how leaves move *now*: their tilt (direction), their change in tilt (curvature), and whether they move coherently or chaotically. Translated into time series language, the leaf's tilt corresponds to the local first difference, the change in tilt to the second difference, and the coherence of motion to the sign pattern of recent differences. The Leaves-in-the-Wind Oracle formalizes this intuition into a simple, regime-aware forecasting rule.

The contribution is threefold. First, the paper defines a kinematic, nonparametric forecasting method based on local velocity, acceleration, and an index of directional coherence. Second, it offers an interpretation of the new forecast function as a heterogeneous autoregression. Third, it provides a new link with the well known Theta method and elevates her applicability – showing that the suggested hybrid formulation is a kinematic generalization of the Theta decomposition. As the method does not require estimation, an ensemble variant is also constructed by aggregating multiple window configurations and by combining individual and inter-model uncertainty. The method is deliberately simple: there are no parameters to estimate beyond short window lengths and a volatility multiplier, and all operations reduce to differences, means, and standard deviations. This contrasts with nonparametric local polynomial forecasting, which typically involves fitting small regressions in moving windows or state spaces. At the same time, the oracle differs from time-series momentum rules in that it can use ultra-short windows, explicit curvature, and regime classification.

The rest of the paper is structured as follows: in Section 2 we present the methodology, in Section 3 we discuss the algebraic lag structure implied by the forecast function, in Section 4 we discuss the stationarity properties of the forecast function and suggest the hybrid model, in Section 5 we introduce the ensemble oracle, in Section 6 we offer an empirical illustration while Section 7 offers additional discussion and the concluding remarks.

2 Methodology

2.1 Velocity and Acceleration

Let $\{Y_t\}_{t=1}^T$ be the sample realization of a univariate time series, whose structure (stationary or not integrated) we still not know. The goal is to construct an h -step-ahead forecast $\hat{Y}_{t+h}$ using only information up to time t as in all forecasting methods. The oracle focuses on the most recent portion of the series, defined by short windows of length w and w' for velocity and acceleration, respectively. Define the *local velocity* and *local acceleration* by:

$$v_t = Y_t - Y_{t-1},$$

$$a_t = v_t - v_{t-1} = Y_t - 2Y_{t-1} + Y_{t-2}.$$

These are simply the first and second-order differences but are interpreted as instantaneous kinematic descriptors of the time series' path. Next, let $\mathcal{V}_t^{(w)} = \{v_{t-w+1}, \dots, v_t\}$ denote the window of the last w velocities, and $\mathcal{A}_t^{(w')} = \{a_{t-w'+1}, \dots, a_t\}$ the window of the last w' ac-

celerations. The oracle uses windowed averages as smoothed estimates of local velocity and acceleration as in:

$$\bar{v}_t^{(w)} = \frac{1}{w} \sum_{j=0}^{w-1} v_{t-j}, \quad \bar{a}_t^{(w')} = \frac{1}{w'} \sum_{j=0}^{w'-1} a_{t-j}.$$

2.2 Directional Coherence and Regimes

In addition to the above magnitudes, the oracle measures the sign pattern of recent movement. Define the *directional momentum index* over window w as:

$$\mathcal{M}_t^{(w)} = \frac{1}{w} \sum_{j=0}^{w-1} \text{sign}(v_{t-j}),$$

which lies in $[-1, 1]$ and reflects the fraction of recent steps that were positive versus negative. The associated *directional coherence* is:

$$\kappa_t = |\mathcal{M}_t^{(w)}|.$$

Values of κ_t near 1 indicate a highly persistent local direction, while values near 0 indicate oscillation. The oracle also distinguishes regimes in a simple phase-space sense, using the signs of current velocity and acceleration. Let v_t and a_t be the most recent realized velocity and acceleration. Define two regimes:

$$\text{Regime}(t) = \begin{cases} \text{momentum,} & \text{if } \text{sign}(v_t) = \text{sign}(a_t) \text{ and } a_t \neq 0, \\ \text{reversal,} & \text{otherwise.} \end{cases}$$

In the momentum regime the direction and its change agree (the rise is accelerating or the fall is steepening), while in the reversal regime the acceleration pushes against the current direction, suggesting a turning point.

2.3 Forecast Equations

The Leaves-in-the-Wind Oracle constructs forecasts by extrapolating the local position Y_t with a combination of velocity and acceleration, modulated by directional coherence and regime. For a given horizon $h \geq 1$ we define:

Momentum regime. When $\text{Regime}(t) = \text{momentum}$, the oracle uses a full kinematic extrapolation:

$$\hat{Y}_{t+h} = Y_t + h \kappa_t \bar{v}_t^{(w)} + \frac{h(h+1)}{2} \bar{a}_t^{(w')}.$$

The first term moves along the average velocity, attenuated by coherence, and the second term adjusts for curvature. The factor κ_t reduces the effective velocity when recent signs are mixed. The reader will immediately note that the method adjusts for the last observation so the naive forecast is the natural benchmark against which the method should be changed.

Reversal regime. When $\text{Regime}(t) = \text{reversal}$, the oracle dampens by half the influence of velocity and completely ignores curvature:

$$\hat{Y}_{t+h} = Y_t + h \cdot 0.5 \kappa_t \bar{v}_t^{(w)}.$$

Here the seer “distrusts” the current direction because acceleration is pointing towards a change, so it extrapolates only a moderated fraction of the recent drift.

2.4 Uncertainty Quantification

Forecast uncertainty is quantified in terms of local dispersion in velocities and an additional model-uncertainty factor. Let

$$\sigma_v^{(w)} = \sqrt{\frac{1}{w-1} \sum_{j=0}^{w-1} \left(v_{t-j} - \bar{v}_t^{(w)} \right)^2}$$

denote the empirical standard deviation of recent velocities. Under a Gaussian diffusion approximation for cumulative changes over horizon h , the base scale of forecast uncertainty is proportional to $\sigma_v^{(w)} \sqrt{h}$.

To account for model risk when coherence is low, the oracle introduces a coherence-dependent factor,

$$u_t = 1 + 0.5(1 - \kappa_t),$$

so that uncertainty increases as κ_t decreases. With a user-specified multiplier $m > 0$ and normal quantile $z_{1-\alpha/2}$, the half-width of the prediction interval is

$$W_{t,h} = z_{1-\alpha/2} \sigma_v^{(w)} \sqrt{h} m u_t,$$

and the resulting $(1 - \alpha)$ interval is

$$\left[\hat{Y}_{t+h} - W_{t,h}, \hat{Y}_{t+h} + W_{t,h} \right].$$

3 Algebraic Structure and Lag Representation

3.1 Expansion of Velocity and Acceleration

An important insight emerges when the forecast equations are expanded in terms of lags of the dependent variable. The windowed averages of velocity and acceleration admit simple closed forms due to telescoping sums.

Proposition 1 (Telescoping velocity). *The windowed average velocity is:*

$$\bar{v}_t^{(w)} = \frac{1}{w} (Y_t - Y_{t-w}).$$

Proof. By definition,

$$\bar{v}_t^{(w)} = \frac{1}{w} \sum_{j=0}^{w-1} v_{t-j} = \frac{1}{w} \sum_{j=0}^{w-1} (Y_{t-j} - Y_{t-j-1}).$$

This is a telescoping sum:

$$\sum_{j=0}^{w-1} (Y_{t-j} - Y_{t-j-1}) = (Y_t - Y_{t-1}) + (Y_{t-1} - Y_{t-2}) + \cdots + (Y_{t-w+1} - Y_{t-w}) = Y_t - Y_{t-w}.$$

Dividing by w yields the result. □

Proposition 2 (Telescoping acceleration). *The windowed average acceleration is:*

$$\bar{a}_t^{(w')} = \frac{1}{w'} [(Y_t - Y_{t-1}) - (Y_{t-w'} - Y_{t-w'-1})] = \frac{1}{w'} (v_t - v_{t-w'}).$$

Proof. Recall that $a_t = Y_t - 2Y_{t-1} + Y_{t-2}$. Then

$$\bar{a}_t^{(w')} = \frac{1}{w'} \sum_{j=0}^{w'-1} (Y_{t-j} - 2Y_{t-j-1} + Y_{t-j-2}).$$

Rearranging by grouping terms with the same lag index:

$$\begin{aligned} \sum_{j=0}^{w'-1} (Y_{t-j} - 2Y_{t-j-1} + Y_{t-j-2}) &= \sum_{k=0}^{w'-1} Y_{t-k} - 2 \sum_{k=1}^{w'} Y_{t-k} + \sum_{k=2}^{w'+1} Y_{t-k} \\ &= Y_t - Y_{t-1} - Y_{t-w'} + Y_{t-w'-1}. \end{aligned}$$

Dividing by w' completes the proof. □

3.2 Open form of the Forecast: Momentum Regime

Using the telescoping identities, the one-step-ahead forecast in the momentum regime becomes:

$$\begin{aligned} \hat{Y}_{t+1} &= Y_t + \kappa_t \bar{v}_t^{(w)} + \bar{a}_t^{(w')} \\ &= Y_t + \kappa_t \cdot \frac{1}{w} (Y_t - Y_{t-w}) + \frac{1}{w'} [(Y_t - Y_{t-1}) - (Y_{t-w'} - Y_{t-w'-1})] \\ &= Y_t \left(1 + \frac{\kappa_t}{w} + \frac{1}{w'} \right) - \frac{1}{w'} Y_{t-1} - \frac{\kappa_t}{w} Y_{t-w} - \frac{1}{w'} Y_{t-w'} + \frac{1}{w'} Y_{t-w'-1}. \end{aligned}$$

For the general h -step-ahead forecast we get:

$$\hat{Y}_{t+h} = Y_t + h \kappa_t \bar{v}_t^{(w)} + \frac{h(h+1)}{2} \bar{a}_t^{(w')},$$

and substituting the telescoping forms yields:

$$\boxed{\begin{aligned}\hat{Y}_{t+h} &= \phi_0(h, \kappa_t, w, w') Y_t + \phi_1(h, w') Y_{t-1} \\ &\quad + \phi_w(h, \kappa_t, w) Y_{t-w} + \phi_{w'}(h, w') Y_{t-w'} \\ &\quad + \phi_{w'+1}(h, w') Y_{t-w'-1}\end{aligned}} \quad (1)$$

where the *state-dependent* coefficients are now given by:

$$\begin{aligned}\phi_0(h, \kappa_t, w, w') &= 1 + \frac{h \kappa_t}{w} + \frac{h(h+1)}{2w'}, \\ \phi_1(h, w') &= -\frac{h(h+1)}{2w'}, \\ \phi_w(h, \kappa_t, w) &= -\frac{h \kappa_t}{w}, \\ \phi_{w'}(h, w') &= -\frac{h(h+1)}{2w'}, \\ \phi_{w'+1}(h, w') &= +\frac{h(h+1)}{2w'}.\end{aligned}$$

This is a *five-lag* model with lags 0 and w , and coefficients modulated by coherence κ_t .

3.3 Open form of the Forecast: Reversal Regime

In the reversal regime, the oracle damps the velocity and ignores curvature. The forecast is now given by:

$$\hat{Y}_{t+h} = Y_t + \frac{h \kappa_t}{2} \bar{v}_t^{(w)} = Y_t \left(1 + \frac{h \kappa_t}{2w} \right) - \frac{h \kappa_t}{2w} Y_{t-w}.$$

This is a *two-lag* model with lags 0 and w , and coefficients modulated by coherence κ_t .

3.4 Interpretation as a Non-Standard Autoregressive Model

The Leaves-in-the-Wind Oracle can be written compactly as

$$\hat{Y}_{t+h} = \sum_{j \in \mathcal{L}} \phi_j(h, \kappa_t, w, w') Y_{t-j},$$

where $\mathcal{L} = \{0, 1, w', w' + 1, w\}$ is the *active lag set*. This is a distributed lag model with several distinctive features:

1. **Sparse, non-consecutive lags.** Unlike standard AR(p) models, which use lags $1, \dots, p$, the oracle uses lags $\{0, 1, w', w' + 1, w\}$ with potentially large gaps.
2. **State-dependent coefficients.** The coefficient on the current level Y_t depends on the directional coherence κ_t , which varies with the recent sign pattern of changes. This is a form of *time-varying parameter* model, but constructed kinematically rather than via stochastic parameter evolution.

3. **Unit sum in neutral case.** When $\kappa_t = 0$ (no directional coherence), the coefficients sum to

$$\phi_0(h, 0, w, w') + \phi_1 + \phi_{w'} + \phi_{w'+1} + \phi_w = 1,$$

so the forecast reduces to a weighted average of past observations, akin to a martingale or random-walk-like rule.

4. **Drift when $\kappa_t > 0$.** For positive coherence, $\phi_0 > 1$ and $\phi_w < 0$, which introduces an implicit drift or trend extrapolation.
5. **Regime switching.** The functional form changes between momentum and reversal regimes, leading to a piecewise AR structure determined by the phase-space position (v_t, a_t) .
6. **No estimation.** The coefficients are constructed directly from the window lengths and the observed coherence, requiring no optimization or likelihood maximization.

3.5 Comparison with Related Models

Standard AR(p). A classical AR(p) takes the form

$$\hat{Y}_{t+h|t} = \mu + \sum_{j=1}^p \psi_j Y_{t-j+1},$$

with constant coefficients ψ_j estimated by OLS or MLE. The Leaves Oracle differs in its non-consecutive lag structure, the absence of an intercept (the role of drift is absorbed into ϕ_0), and the dependence of coefficients on κ_t .

Threshold AR (TAR) models. Threshold autoregressions switch between distinct AR processes based on whether a threshold variable crosses a boundary. The Leaves Oracle has a similar flavor, but the switching criterion is the sign alignment (v_t, a_t) rather than a single threshold on levels or differences, and the regimes are defined by kinematic principles rather than estimated cut-points.

Time-varying parameter (TVP) models. TVP-AR models allow coefficients to evolve stochastically. The Leaves Oracle fixes the functional form of coefficient variation to depend on κ_t , which is itself a deterministic function of recent data, rather than a latent state variable with its own dynamics.

Local polynomial forecasting. Methods based on local polynomial regression fit low-degree polynomials in moving windows and extrapolate. The Leaves Oracle also uses local information, but it does so without regression, relying instead on simple arithmetic means of differences.

3.6 Implications for Forecast Properties

The lag representation clarifies several properties of the oracle:

- **Explosive versus stable dynamics.** When κ_t is high and the coefficient ϕ_0 exceeds unity substantially, the forecast can exhibit explosive behavior over multiple steps if the directional coherence is misread as permanent. However, the regime-switching mechanism acts as a stabilizer: as the series approaches a turning point, the oracle switches to the reversal regime with damped coefficients.
- **Consistency with local geometry.** The use of lags w and w' directly encodes the chosen smoothing scales for velocity and acceleration. Shorter windows make the oracle more reactive to recent shifts; longer windows smooth out noise but risk lagging behind structural changes.
- **Interpretation of forecast errors.** Errors arise when the local kinematic extrapolation fails—either because curvature changes abruptly (invalidating the quadratic approximation) or because a regime shift occurs between forecast construction and realization. The model has no memory beyond the active lags, so persistence beyond window w is not captured unless κ_t remains high.

3.7 A Heterogeneous Autoregression Interpretation

The open form can also be viewed as the solution to an implicit difference equation. Define the forecast error innovation:

$$\eta_{t+h} = Y_{t+h} - \hat{Y}_{t+h|t}.$$

If we interpret the oracle’s rule as a model for the conditional expectation, then:

$$Y_{t+h} = \sum_{j \in \mathcal{L}} \phi_j(h, \kappa_t, w, w') Y_{t-j} + \eta_{t+h},$$

which is a *heterogeneous autoregression* with time-varying, coherence-dependent weights. Unlike GARCH or other conditional heteroskedasticity models, the heterogeneity here enters the conditional mean rather than variance. This perspective also suggests a path for future work: one could estimate analogous sparse-lag AR models with state-dependent coefficients using formal econometric methods and compare their performance to the oracle’s construction-based approach.

4 Stationarity Properties

4.1 The unit root property

An important feature of the Leaves-in-the-Wind Oracle emerges from examining the sum of its autoregressive coefficients. Recall the open-form representation in the momentum regime:

$$\hat{Y}_{t+h} = \phi_0(h, \kappa_t, w, w') Y_t + \phi_1(h, w') Y_{t-1} + \phi_w(h, \kappa_t, w) Y_{t-w} + \phi_{w'}(h, w') Y_{t-w'} + \phi_{w'+1}(h, w') Y_{t-w'-1}.$$

Proposition 3 (Unit coefficient sum). *For both the momentum and reversal regimes, the sum*

of forecast coefficients satisfies

$$\sum_{j \in \mathcal{L}} \phi_j = 1$$

for all values of the coherence parameter κ_t , horizon h , and window lengths w, w' .

Proof. For the momentum regime with $h = 1$,

$$\begin{aligned} \sum \phi_j &= \phi_0 + \phi_1 + \phi_w + \phi_{w'} + \phi_{w'+1} \\ &= \left(1 + \frac{\kappa_t}{w} + \frac{1}{w'}\right) + \left(-\frac{1}{w'}\right) + \left(-\frac{\kappa_t}{w}\right) + \left(-\frac{1}{w'}\right) + \left(\frac{1}{w'}\right) \\ &= 1 + \frac{\kappa_t}{w} - \frac{\kappa_t}{w} + \frac{1}{w'} - \frac{1}{w'} - \frac{1}{w'} + \frac{1}{w'} \\ &= 1. \end{aligned}$$

For the reversal regime,

$$\sum \phi_j = \left(1 + \frac{\kappa_t}{2w}\right) + \left(-\frac{\kappa_t}{2w}\right) = 1.$$

The result extends to general h by inspection of the coefficient formulas. $\square$

Remark (Significance of unit root property). *The unit coefficient sum $\sum \phi_j = 1$ is not merely a mathematical curiosity—it reveals the fundamental design principle of the Leaves Oracle. By construction, the forecast increment $\hat{Y}_{t+h} - Y_t$ depends only on recent changes (velocity and acceleration), not on the level Y_t itself. This is precisely the behavior of an integrated $I(1)$ process, where shocks have permanent effects and the series has no inherent mean reversion. The standard oracle is therefore natively designed for unit-root data, making it particularly suitable for financial prices, exchange rates, and other persistent economic time series.*

The hybrid extension we shall discuss later with $\lambda > 0$ breaks this unit root by blending toward a mean, yielding $\sum \phi_j = 1 - \lambda < 1$, which induces stationarity. This provides a natural and interpretable way to adapt the oracle to different persistence regimes.

4.2 Interpretation as a Differenced Model

The unit coefficient sum implies that the Leaves Oracle is fundamentally a *differenced* or *integrated* forecasting model. To see this, subtract Y_t from both sides of the forecast equation:

$$\begin{aligned} \hat{Y}_{t+1} - Y_t &= \sum_{j \in \mathcal{L}} \phi_j Y_{t-j} - Y_t \\ &= \left(\sum_{j \in \mathcal{L}} \phi_j - 1 \right) Y_t + \sum_{j \geq 1} \phi_j (Y_{t-j} - Y_t) \\ &= 0 \cdot Y_t + \sum_{j \geq 1} \phi_j (Y_{t-j} - Y_t) \end{aligned}$$

since $\sum \phi_j = 1$. The forecast increment $\hat{Y}_{t+1} - Y_t$ depends only on differences between past levels and the current level, not on the current level itself. Substituting the explicit coefficients

for the momentum regime yields:

$$\hat{Y}_{t+1} - Y_t = \kappa_t \bar{v}_t^{(w)} + \bar{a}_t^{(w')},$$

which is precisely the kinematic drift expression derived earlier. The oracle does not predict the level Y_{t+1} directly; instead, it predicts the *change* ΔY_{t+1} based on recent velocity and acceleration.

4.3 Consequences for Stationarity and Persistence

In classical time series terminology, a process $\{Y_t\}$ is covariance-stationary if its autocovariance function depends only on lag and not on time, and all moments are finite. A necessary condition for the stationarity of an AR(p) process is that the roots of the characteristic polynomial lie outside the unit circle. When $\sum_{j=1}^p \phi_j = 1$, the polynomial has a unit root at $z = 1$, and the process is *integrated of order one*, or I(1). Such processes are nonstationary and exhibit persistent, random-walk-like behavior. The Leaves-in-the-Wind Oracle, with $\sum \phi_j = 1$ by construction, inherits this property. It is designed to forecast series that behave locally like unit-root processes, where past changes persist indefinitely unless counteracted by new shocks or regime shifts. This design choice has both strengths and limitations.

Strengths.

1. **Appropriate for I(1) data.** Many economic and financial time series—stock prices, exchange rates, macroeconomic aggregates in levels—are well described as I(1). The oracle’s unit-sum structure aligns with this empirical regularity.
2. **Short-horizon accuracy.** At horizons $h = 1$ or $h = 2$, local momentum often dominates mean reversion, making the oracle’s drift extrapolation effective even when the series is stationary in the long run.
3. **Regime-adaptive damping.** The coherence parameter κ_t and the regime-switching mechanism provide a form of endogenous stabilization. When recent movements are incoherent ($\kappa_t \approx 0$), the velocity contribution vanishes, and the forecast approaches Y_t . When a reversal regime is detected, the drift is halved, mimicking a turning-point correction.

Limitations.

1. **Overextrapolation in stationary settings.** If the true data-generating process is a stationary AR(1), $Y_t = \phi Y_{t-1} + \varepsilon_t$ with $|\phi| < 1$, the conditional expectation $\mathbb{E}[Y_{t+h} | Y_t]$ decays exponentially toward the unconditional mean at rate ϕ^h . The Leaves Oracle, however, projects a constant drift indefinitely. For h large, the oracle’s forecast will systematically diverge from the true conditional mean.
2. **No long-run anchor.** The oracle has no built-in mechanism to revert toward a long-run mean or equilibrium. In a strongly mean-reverting environment, this leads to forecast bias.

3. **Persistent drift in white noise.** If $\{Y_t\}$ is white noise or nearly so, the best forecast is simply the mean. The oracle will still detect spurious patterns in small samples, though low coherence κ_t mitigates this – but we account for stationary environments in our introduction of the hybrid approach, as noted in the previous remark and we shall now discuss.

4.4 A Hybrid oracle with Mean Reversion

To extend the applicability of the oracle to stationary environments, we propose a *hybrid* variant that incorporates mean reversion. Define the long-run mean estimate $\hat{\mu}$ as the sample mean of the observed series (or an exponentially weighted moving average, or a theoretically motivated equilibrium level). Introduce a mean-reversion parameter $\lambda \in [0, 1]$, and modify the forecast as

$$\hat{Y}_{t+h}^{\text{hybrid}} = (1 - \lambda) \left[Y_t + h \kappa_t \bar{v}_t^{(w)} + \frac{h(h+1)}{2} \bar{a}_t^{(w')} \right] + \lambda \hat{\mu}.$$

Connection to the Theta method. The hybrid oracle bears a profound relationship to the Theta method of Assimakopoulos and Nikolopoulos (2000)¹ and its extensions by Nikolopoulos and Thomakos. The classical Theta method decomposes a time series into two components:

$$\hat{Y}_{t+h}^{\text{Theta}} = \theta \cdot \hat{Y}_{t+h}^{\text{local}} + (1 - \theta) \cdot \hat{Y}_{t+h}^{\text{global}},$$

where the local component captures short-run dynamics (trend, momentum, the last observed value among others) and the global component represents long-run equilibrium (typically a mean or slowly-varying level), with mixing parameter $\theta \in [0, 1]$. Our hybrid oracle can be written in identical form:

$$\hat{Y}_{t+h}^{\text{hybrid}} = (1 - \lambda) \cdot \underbrace{\left[Y_t + h \kappa_t \bar{v}_t^{(w)} + \frac{h(h+1)}{2} \bar{a}_t^{(w')} \right]}_{\text{Local kinematic component}} + \lambda \cdot \underbrace{\hat{\mu}}_{\text{Global equilibrium}}.$$

The correspondence is direct: $\lambda \equiv 1 - \theta$, so that

- The *local component* is the kinematic extrapolation based on velocity and acceleration—a physically motivated, dynamic trend model rather than a simple linear or SES-based trend.
- The *global component* is the long-run mean $\hat{\mu}$, capturing the series' equilibrium or central tendency.
- The *mixing parameter* λ controls the balance between persistence (local dynamics) and mean reversion (global anchor).

This reformulation positions the hybrid Leaves Oracle as a **kinematic generalization of the Theta method**, where the local component is enriched with second-order dynamics (acceleration) and directional coherence (via κ_t), rather than relying solely on exponential smoothing

¹The Theta method won the M3 forecasting competition (Makridakis and Hibon, 2000) and has been extensively studied by Nikolopoulos, Thomakos, and collaborators, who have developed optimal parameter selection rules and theoretical foundations for the decomposition principle.

or linear trends. The theoretical advantage is that the kinematic component is designed for unit-root behavior ($\sum \phi_j = 1$), making it appropriate for I(1) series, while the mean-reversion parameter $\lambda > 0$ induces stationarity ($\sum \phi_j = 1 - \lambda < 1$), exactly as in Theta variants designed for stationary data.

Interpretation.

- When $\lambda = 0$, the hybrid oracle reduces to the standard Leaves Oracle with unit coefficient sum.
- When $\lambda = 1$, the forecast is simply the long-run mean, ignoring all local kinematic information.
- For $\lambda \in (0, 1)$, the forecast is a convex combination of the kinematic extrapolation and the mean, balancing short-run momentum with long-run equilibrium.

Choosing λ . Several strategies can be employed to select λ :

1. **Cross-validation.** Estimate λ by minimizing out-of-sample forecast error on a validation set.
2. **Autocorrelation-based heuristic.** If the sample autocorrelation function (ACF) decays quickly, increase λ to induce stronger mean reversion; if the ACF is highly persistent, decrease λ .
3. **Regime-dependent λ .** Allow λ to vary with the coherence κ_t or regime. For instance, set

$$\lambda_t = \lambda_{\max} \cdot (1 - \kappa_t),$$

so that high coherence (strong local trend) reduces mean reversion, while low coherence (oscillatory behavior) increases it.

4. **Fixed small value.** In the absence of prior information, a modest $\lambda \in [0.1, 0.3]$ can stabilize forecasts without heavily dampening the kinematic component.

The flexibility in choosing λ makes the hybrid oracle adaptable to different degrees of persistence, positioning it as a versatile forecasting tool that unifies kinematic physics with the empirical success of the Theta method.

5 Ensemble Oracle

Multi-leaf Aggregation To reflect the idea of watching many leaves at different heights and positions, an ensemble oracle aggregates multiple Leaves-in-the-Wind configurations. Let $\{(w_i, w'_i)\}_{i=1}^L$ be a set of window pairs, and let $\hat{Y}_{t+h}^{(i)}$ be the forecast produced by a leaves oracle with configuration (w_i, w'_i) . The ensemble point forecast is the median:

$$\hat{Y}_{t+h}^{\text{ens}} = \text{Median}\left(\hat{Y}_{t+h}^{(1)}, \dots, \hat{Y}_{t+h}^{(L)}\right),$$

which is robust to outlier behavior from individual configurations.

Ensemble Uncertainty Ensemble uncertainty combines two components:

- Average individual uncertainty, as captured by the typical interval width of component oracles.
- Inter-model disagreement, as captured by the standard deviation of point forecasts across configurations.

Let $W_{t,h}^{(i)}$ be the individual interval width of oracle i , and define

$$\bar{W}_{t,h} = \frac{1}{L} \sum_{i=1}^L W_{t,h}^{(i)},$$

$$s_{t,h} = \text{sd}(\hat{Y}_{t+h}^{(1)}, \dots, \hat{Y}_{t+h}^{(L)}).$$

Given a chosen α , inter-model disagreement induces a width proxy $2z_{1-\alpha/2}s_{t,h}$. The ensemble then sets

$$W_{t,h}^{\text{ens}} = M \cdot \max(\bar{W}_{t,h}, 2z_{1-\alpha/2}s_{t,h}),$$

for a multiplier M , and reports

$$\left[\hat{Y}_{t+h}^{\text{ens}} - \frac{W_{t,h}^{\text{ens}}}{2}, \hat{Y}_{t+h}^{\text{ens}} + \frac{W_{t,h}^{\text{ens}}}{2} \right]$$

as the prediction interval.

Ensemble Hybrid Oracle. The multi-leaf ensemble can also be extended to the hybrid setting. Each component oracle i produces

$$\hat{Y}_{t+h}^{(i,\text{hybrid})} = (1 - \lambda)\hat{Y}_{t+h}^{(i)} + \lambda\hat{\mu},$$

and the ensemble forecast is the median of these hybrid forecasts, with uncertainty aggregated as before.

6 An Empirical Illustration

6.1 Set-up

Using three standard benchmark models, the no change, the sample mean and the AR(1) models, I next offer a simple illustration of performance. Please keep in mind that the leaves-in-the-wind oracle does not require any estimation at all. The performance of all models is assessed using a rolling-origin evaluation of 12 observations for the AR(1) model and 12 (monthly) or 4 (quarterly observations) for the mean. The parametrizations for the oracle are given in the tables - and you can experiment more using the Python code files. and the . Given a series of length T and a window size T_0 , forecasts are generated for times $t = T_0, \dots, T - h$. At each origin t , the forecasters use data $Y_{t-T_0+1}, \dots, Y_t$ and produce an h -step-ahead prediction $\hat{Y}_{t+h}$ and an associated prediction interval. For simplicity, the empirical analysis focuses on horizon

$h = 1$. Performance metrics include mean absolute error (MAE), root mean squared error (RMSE), mean absolute percentage error (MAPE), prediction interval coverage, and average interval width (sharpness). Coverage measures how often realized values fall inside the nominal $(1 - \alpha)$ interval, while sharpness measures the average width. Calibration is summarized by the absolute deviation of empirical coverage from the target. To assess whether the differences in forecast performance are statistically significant, we use the Diebold–Mariano (DM) test based on squared-error loss. The results follow in the tables below, first with a set of tables with model nomenclature (Tables 1 and 2) and then with a set of tables with results (Table 3 to 9). I have examined three series: the US annual inflation rate, the US annual real GDP growth and the monthly level of the federal funds rate - details can be found in the tables.

Table 1: Model Nomenclature and Parameter Specifications

Model Name	Type	Specification
<i>Panel A: Standard Leaves Oracle Models</i>		
Leaves_5_3	Kinematic	Velocity window $w = 5$, Acceleration window $w' = 3$
Leaves_12_6	Kinematic	Velocity window $w = 12$, Acceleration window $w' = 6$
<i>Panel B: Hybrid Leaves Oracle Models</i>		
Hybrid.Fixed.0.2	Hybrid	Fixed $\lambda = 0.2$ with $w = 5$, $w' = 3$. Combines kinematic and mean-reversion components.
Hybrid.ACF	Hybrid-Adaptive	Adaptive λ based on autocorrelation function (ACF). Maximum $\lambda = 0.5$. Adjusts to persistence in data.
Hybrid.Coh	Hybrid-Adaptive	Adaptive λ based on kinematic coherence measure. Responds to regime changes and structural breaks.
<i>Panel C: Ensemble Oracle Models</i>		
Ensemble.Std	Ensemble	Standard ensemble combining Leaves_3_2, Leaves_5_3, Leaves_7_3, Leaves_10_5. Equal-weighted median forecast.
Ensemble.Hybrid	Ensemble	Hybrid ensemble with adaptive λ strategy (ACF-based). Combines multiple window pairs with mean-reversion.
<i>Panel D: Benchmark Models</i>		
NoChange	Random Walk	$\hat{y}_{t+1} = y_t$. Prediction interval: $y_t \pm 1.96 \times \sigma_t$ where σ_t is rolling standard deviation.
AR1	Autoregressive	$y_t = \alpha + \beta y_{t-1} + \epsilon_t$. Rolling window estimation.
Mean_ k	Historical Mean	$\hat{y}_{t+1} = \frac{1}{k} \sum_{i=0}^{k-1} y_{t-i}$. Simple moving average with window $k \in \{4, 12\}$.

6.2 Results

The empirical results across three major U.S. macroeconomic indicators demonstrate that the Leaves-in-the-Wind Oracle framework consistently delivers competitive and, in many cases, superior forecast performance relative to conventional benchmarks. As shown in Tables 3, 5, and 7, the kinematic oracle models achieve lower mean absolute errors (MAE) than the ran-

Table 2: Parameter Values and Hyperparameter Settings

Parameter	Symbol	Value(s) / Description
<i>Kinematic Parameters</i>		
Velocity window	w	$\{3, 5, 7, 10, 12\}$ periods
Acceleration window	w'	$\{2, 3, 4, 5, 6\}$ periods
Acceleration scaling	α	0.5 (fixed across all models)
<i>Hybrid Oracle Parameters</i>		
Mean-reversion weight	λ	Fixed: $\{0.1, 0.2, 0.3\}$ or adaptive
Mean estimation window	w_{mean}	120 periods (10 years for monthly data)
ACF lag for adaptation	—	First-order autocorrelation ρ_1
Coherence threshold	—	Dynamic based on velocity-acceleration alignment
Maximum λ (adaptive)	λ_{max}	0.5 (prevents over-damping)
<i>Ensemble Parameters</i>		
Window configurations	—	Standard: $\{(3, 2), (5, 3), (7, 3), (10, 5)\}$
Aggregation method	—	Median (robust to outliers)
Interval aggregation	—	Conservative: max(upper bounds), min(lower bounds)
<i>Forecast Evaluation Parameters</i>		
Initial training window	n_{init}	60 periods (5 years for monthly, 15 years for quarterly)
Forecast horizon	h	1 period ahead
Rolling step size	—	1 period (maximum out-of-sample observations)
Confidence level	—	95% for prediction intervals
Uncertainty multiplier	τ	2.0 (calibrated for 95% coverage)
<i>Statistical Test Parameters</i>		
DM test HAC lags	L	$\lfloor 4(n/100)^{2/9} \rfloor$ (Newey-West)
Significance levels	—	$\alpha \in \{0.01, 0.05, 0.10\}$
Loss function	—	Squared error: $L(e_t) = e_t^2$

Table 3: Forecast Performance: Annual Inflation (CPIAUCSL, Year-over-Year % Change)

Model	MAE	RMSE	Coverage	Hit Rate	Avg λ	Avg Coherence
Leaves_12_6	0.27	0.39	0.99	0.54	0.00	0.31
Ensemble_Std	0.27	0.38	0.99	0.57	0.00	0.42
NoChange	0.28	0.40	0.92	0.01	0.00	0.00
Leaves_5_3	0.28	0.38	0.96	0.54	0.00	0.43
Hybrid_Fixed.0.2	0.29	0.42	0.95	0.53	0.20	0.43
Hybrid_Coh	0.30	0.43	0.99	0.51	0.21	0.31
AR1	0.31	0.46	0.84	0.52	0.00	0.00
Hybrid_ACF	0.43	0.63	0.84	0.48	0.50	0.43
Ensemble_Hybrid	0.44	0.64	0.93	0.48	0.50	0.42
Mean_12	0.76	1.13	0.26	0.48	0.00	0.00

Notes: Sample period: 1991:01–2026:01 (420 observations). Mean = 2.63%, Std = 1.55%. MAE and RMSE in percentage points. Coverage refers to 95% prediction interval coverage. Hit Rate measures directional accuracy. Models ranked by MAE.

Table 4: Diebold-Mariano Tests: Annual Inflation (CPIAUCSL)

Comparison	DM Statistic	p-value
<i>Panel A: Oracle vs NoChange</i>		
Leaves_5_3 vs NoChange	−1.04	0.30
Leaves_12_6 vs NoChange	−2.00**	0.05
Hybrid_Fixed.0.2 vs NoChange	0.77	0.44
Hybrid_Coh vs NoChange	2.54**	0.01
Hybrid_ACF vs NoChange	3.63***	0.00
Ensemble_Std vs NoChange	−1.40	0.16
Ensemble_Hybrid vs NoChange	3.64***	0.00
<i>Panel B: Oracle vs AR1</i>		
Leaves_5_3 vs AR1	−2.81***	0.00
Leaves_12_6 vs AR1	−3.57***	0.00
Hybrid_Fixed.0.2 vs AR1	−1.76*	0.08
Hybrid_Coh vs AR1	−1.76*	0.08
Hybrid_ACF vs AR1	2.86***	0.00
Ensemble_Std vs AR1	−2.99***	0.00
Ensemble_Hybrid vs AR1	2.90***	0.00

Notes: Negative DM statistic indicates first model has lower loss (better). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Forecast Performance: Federal Funds Rate (FEDFUNDS, Level)

Model	MAE	RMSE	Coverage	Hit Rate	Avg λ	Avg Coherence
Ensemble_Std	0.09	0.14	0.97	0.56	0.00	0.48
Leaves_5_3	0.09	0.15	0.93	0.58	0.00	0.49
Leaves_12_6	0.09	0.15	0.96	0.55	0.00	0.42
NoChange	0.10	0.18	0.80	0.20	0.00	0.00
Hybrid_Coh	0.12	0.20	0.94	0.39	0.17	0.42
Hybrid_Fixed.0.2	0.15	0.23	0.81	0.31	0.20	0.49
Hybrid_ACF	0.30	0.45	0.56	0.24	0.50	0.49
Ensemble_Hybrid	0.30	0.45	0.75	0.24	0.50	0.48
AR1	0.33	1.70	0.72	0.47	0.00	0.00
Mean_12	0.57	0.85	0.14	0.24	0.00	0.00

Notes: Sample period: 1990:01–2026:01 (433 observations). Mean = 2.88%, Std = 2.34%. MAE and RMSE in percentage points. Coverage refers to 95% prediction interval coverage. Hit Rate measures directional accuracy. Models ranked by MAE.

Table 6: Diebold-Mariano Tests: Federal Funds Rate (FEDFUNDS)

Comparison	DM Statistic	p-value
<i>Panel A: Oracle vs NoChange</i>		
Leaves_5_3 vs NoChange	−3.12***	0.00
Leaves_12_6 vs NoChange	−3.59***	0.00
Hybrid_Fixed.0.2 vs NoChange	4.77***	0.00
Hybrid_Coh vs NoChange	2.50**	0.01
Hybrid_ACF vs NoChange	4.92***	0.00
Ensemble_Std vs NoChange	−3.44***	0.00
Ensemble_Hybrid vs NoChange	4.91***	0.00
<i>Panel B: Oracle vs AR1</i>		
Leaves_5_3 vs AR1	−1.20	0.23
Leaves_12_6 vs AR1	−1.20	0.23
Hybrid_Fixed.0.2 vs AR1	−1.19	0.23
Hybrid_Coh vs AR1	−1.20	0.23
Hybrid_ACF vs AR1	−1.13	0.26
Ensemble_Std vs AR1	−1.20	0.23
Ensemble_Hybrid vs AR1	−1.13	0.26

Notes: Negative DM statistic indicates first model has lower loss (better). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Forecast Performance: Real GDP Growth (GDPC1, Year-over-Year % Change)

Model	MAE	RMSE	Coverage	Hit Rate	Avg λ	Avg Coherence
Hybrid_Fixed.0.2	0.82	1.56	0.98	0.60	0.20	0.38
NoChange	0.84	1.68	0.90	0.00	0.00	0.00
Hybrid_Coh	0.87	1.65	0.98	0.54	0.24	0.19
Ensemble_Std	0.87	1.76	0.99	0.52	0.00	0.35
Hybrid_ACF	0.87	1.52	0.95	0.57	0.50	0.38
Leaves_5_3	0.88	1.75	0.97	0.52	0.00	0.38
Ensemble_Hybrid	0.88	1.56	0.98	0.56	0.50	0.35
AR1	0.90	1.92	0.85	0.53	0.00	0.00
Leaves_12_6	0.91	1.81	0.98	0.47	0.00	0.19
Mean_4	1.15	2.05	0.37	0.48	0.00	0.00

Notes: Sample period: 1991:Q1–2025:Q3 (139 observations). Mean = 2.52%, Std = 2.03%. MAE and RMSE in percentage points. Coverage refers to 95% prediction interval coverage. Hit Rate measures directional accuracy. Models ranked by MAE.

Table 8: Diebold-Mariano Tests: Real GDP Growth (GDPC1)

Comparison	DM Statistic	p-value
<i>Panel A: Oracle vs NoChange</i>		
Leaves_5_3 vs NoChange	0.59	0.55
Leaves_12_6 vs NoChange	1.51	0.13
Hybrid_Fixed_0.2 vs NoChange	-1.10	0.27
Hybrid_Coh vs NoChange	-0.68	0.50
Hybrid_ACF vs NoChange	-0.85	0.40
Ensemble_Std vs NoChange	0.72	0.47
Ensemble_Hybrid vs NoChange	-0.76	0.45
<i>Panel B: Oracle vs AR1</i>		
Leaves_5_3 vs AR1	-0.71	0.48
Leaves_12_6 vs AR1	-0.41	0.68
Hybrid_Fixed_0.2 vs AR1	-1.15	0.25
Hybrid_Coh vs AR1	-0.87	0.39
Hybrid_ACF vs AR1	-1.05	0.30
Ensemble_Std vs AR1	-0.71	0.48
Ensemble_Hybrid vs AR1	-0.99	0.32

Notes: Negative DM statistic indicates first model has lower loss (better). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Note: No comparisons are statistically significant at conventional levels for GDP growth.

Table 9: Summary: Best Performing Models Across Series

Series	Best Model	MAE	Coverage	DM vs NoChange
<i>Inflation</i> (CPIAUCSL)	Leaves_12_6	0.27	0.99	-2.00**
	Ensemble_Std	0.27	0.99	-1.40
	NoChange	0.28	0.92	—
<i>Fed Funds</i> (FEDFUNDS)	Ensemble_Std	0.09	0.97	-3.44***
	Leaves_5_3	0.09	0.93	-3.12***
	NoChange	0.10	0.80	—
<i>GDP Growth</i> (GDPC1)	Hybrid_Fixed_0.2	0.82	0.98	-1.10
	NoChange	0.84	0.90	—
	Hybrid_Coh	0.87	0.98	-0.68

Notes: Top 2-3 models by MAE shown for each series. DM statistic tests equality of forecast accuracy vs NoChange baseline. Significance: *** $p < 0.01$, ** $p < 0.05$.

dom walk (NoChange) benchmark in two of three series, while maintaining excellent prediction interval coverage exceeding 93% across all applications. For annual inflation (CPIAUCSL), the *Leaves_12_6* configuration—utilizing a 12-period velocity window and 6-period acceleration window—achieves the lowest MAE of 0.27 percentage points, marginally outperforming the ensemble approach (0.27) and notably beating the random walk baseline (0.28). More importantly, this superior point forecast accuracy is accompanied by exceptional uncertainty quantification, with 99% of actual observations falling within the 95% prediction intervals, suggesting that the kinematic framework appropriately captures the dynamic uncertainty in inflation processes. The standard ensemble approach, which aggregates forecasts from multiple window configurations, demonstrates remarkable robustness across all three series, ranking among the top three models in each application and never performing worse than third place by MAE—evidence that diversification across kinematic specifications provides a reliable forecasting strategy even when the optimal window parameters vary across economic contexts.

The performance patterns reveal important insights about the applicability of kinematic modeling to different macroeconomic dynamics. For the federal funds rate (FEDFUNDS), the Leaves Oracle models exhibit their strongest relative performance, with *Leaves_5_3*, *Leaves_12_6*, and *Ensemble_Std* achieving MAEs of 0.09 percentage points—representing 12% to 16% improvements over the NoChange benchmark (0.10). The Diebold-Mariano tests (Table 6) confirm these differences are statistically significant at the 1% level, with DM statistics of -3.12 and -3.59 for the two individual Leaves models, respectively. This exceptional performance likely reflects the policy-driven nature of the federal funds rate, where the Federal Reserve implements gradual adjustments that create persistent velocity and acceleration patterns—precisely the dynamics that kinematic models are designed to capture. In contrast, for real GDP growth (GDPC1), the hybrid models demonstrate their value, with *Hybrid_Fixed_0.2* achieving the lowest MAE (0.82) despite no model significantly outperforming the random walk at conventional significance levels (Table 8). This result suggests that GDP growth exhibits more complex dynamics combining both kinematic trends and mean-reverting behavior, which the hybrid framework accommodates through its λ parameter that blends extrapolation with long-run mean reversion. The smaller sample size for quarterly GDP data (139 observations versus 420–433 for monthly series) also contributes to the wider confidence intervals in statistical tests, highlighting the sample-dependent nature of forecast evaluation.

The statistical tests further reveal the robustness of the kinematic approach and the nuanced role of adaptive strategies. Tables 4 and 6 demonstrate that standard Leaves Oracle models consistently outperform the AR1 benchmark with strong statistical significance (all $p < 0.01$ except one comparison), indicating that the kinematic framework captures dynamics that linear autoregressive models miss, particularly the acceleration component that AR(1) specifications cannot accommodate. However, the hybrid models with high fixed λ values (e.g., *Hybrid_ACF* with $\lambda = 0.5$) and some adaptive strategies underperform simpler kinematic specifications, as evidenced by positive DM statistics against NoChange (e.g., $DM = 3.63$ for *Hybrid_ACF* vs NoChange in inflation forecasting). This pattern suggests that excessive mean-reversion dampening can eliminate the valuable extrapolative information contained in kinematic trends, particularly during periods of persistent directional movement such as the 2021–2023 infla-

tion surge or the 2022–2023 monetary tightening cycle. The optimal λ values observed in the adaptive strategies (averaging 0.17–0.24 for coherence-based approaches) suggest that moderate mean-reversion provides insurance against over-extrapolation without sacrificing too much trend information. Importantly, the hit rates—measuring directional forecast accuracy—range from 0.47 to 0.60 across successful oracle configurations, substantially exceeding the 0.00–0.20 hit rates of the NoChange benchmark, confirming that kinematic models provide genuine forecasting content rather than merely tracking historical volatility. These findings collectively support the proposition that incorporating physical motion principles into econometric forecasting can enhance both point forecast accuracy and directional predictive power for macroeconomic time series exhibiting persistent trending behavior.

7 Discussion and Concluding Remarks

7.1 Relation to existing methods

The Leaves-in-the-Wind Oracle differs conceptually from local polynomial regression methods for time series, which fit explicit polynomial functions to moving windows or reconstructed state spaces. The oracle’s kinematic formulation dispenses with regression altogether and uses only simple differences, windowed averages, and basic sign statistics.

Compared to time-series momentum strategies, which typically rely on longer-horizon cumulative returns and complex risk management overlays, the oracle uses ultra-short windows and emphasizes curvature and coherence rather than long-run momentum. Its domain is short-horizon forecasting of generic time series, not necessarily tradable returns.

7.2 Interpretability and practical appeal

One of the method’s strengths is interpretability. Forecasts can be decomposed into contributions from local velocity and acceleration, and the regime label (momentum or reversal) provides an immediate qualitative description of the current dynamics. The coherence index κ_t offers a scalar summary of directional agreement that can be monitored as a diagnostic signal.

The implementation is minimal and requires no optimization or hyperparameter tuning beyond choosing short window lengths and a volatility multiplier. This makes it particularly attractive for rapid prototyping, low-resource settings, or pedagogical use where the goal is to illustrate local dynamical ideas without heavy machinery.

7.3 Limitations and extensions

The oracle’s limitations are also clear. In purely noise-driven or weakly persistent stationary series, sophisticated structure is largely absent, and simple parametric models remain more efficient. Furthermore, the current formulation is univariate and does not incorporate exogenous variables or multivariate structure.

Extensions suggest themselves in several directions. First, multivariate versions could incorporate local velocity and acceleration vectors in higher-dimensional state spaces. Second, non-Gaussian or robust measures of local dispersion could improve performance under heavy

tails. Third, a Bayesian framing could treat the oracle’s kinematic components as priors within a broader ensemble of models, integrating with more elaborate structural or machine-learning forecasts.

7.4 Conclusion

This paper has introduced the Leaves-in-the-Wind Oracle, a nonparametric, kinematic method for time series forecasting that uses only local velocity, acceleration, and directional coherence. The method is simple, transparent, and regime-aware, and an ensemble variant extends it to multiple window configurations with robust uncertainty quantification. Rolling-window experiments on synthetic data demonstrate that the oracle substantially improves upon a naive no-change forecast in settings with trends, cycles, and regime shifts, while behaving reasonably in stationary autoregressive settings where parametric models are best suited. The discovery that the hybrid oracle corresponds to a kinematic generalization of the Theta method places this work in dialogue with a rich literature on forecast combination and decomposition methods. The standard oracle’s unit-root structure makes it ideal for $I(1)$ data, while the hybrid extension with mean reversion adapts naturally to stationary environments, much like optimized Theta variants. This duality—between kinematic and statistical decomposition—suggests that physically interpretable models need not sacrifice empirical performance or theoretical rigor. Beyond its immediate forecasting performance, the oracle illustrates a broader design philosophy: local, physically interpretable descriptors of path geometry can serve as a viable basis for simple yet effective forecasting rules. This perspective opens up opportunities for further research at the intersection of time series analysis, local geometry, and decision-theoretic design of “oracular” models.