TECHNICAL NOTE

An HJB Estimator for Output-Gap Uncertainty

A Worst-Case Filtration Framework for Central-Bank Decision-Making

K. Papadopoulos — Quantitative Policy & Macroeconomics Division
N. Peled — Quantitative Policy & Macroeconomics Division

07-FEB-2025

Van der Pol: $\dot{x} = y, \quad \dot{y} = \mu\,(1 - x^2)\,y - x$

TN-2026-66770776
iadu.org

IADU

Institute for Advanced Dynamic Uncertainty

Copyright

Suggested Citation

K. Papadopoulos and N. Peled (2025). 'An HJB Estimator for Output-Gap Uncertainty.' *IADU Technical Note* **TN-2026-66770776**.
Available at `https://iadu.org/research/TN-2026-66770776/`.

About IADU

The **Institute for Advanced Dynamic Uncertainty** exists to advance the mathematical theory of decision under uncertainty and to bring that theory, with rigour and restraint, to bear upon the most consequential questions of public and institutional policy. Its work proceeds from the foundations upward: the question shall dictate the method, and never the converse.

Research is organised across five operational divisions: the *Division of Stochastic Analysis and Control*; the *Division of Games, Dynamics, and Strategic Control*; the *Division of Financial Mathematics and Asset Pricing*; the *Division of Quantitative Policy and Macroeconomics*; and the *Division of Sustainability and Energy Economics*. The Institute publishes working papers, technical notes, discussion papers, policy briefs, research and technical reports, preprints, surveys, data reports, and research memoranda, and produces the *IADU Quantitative Policy Review* as its principal vehicle for engagement with the policy community.

The Institute's research is purely bottom-up. It does not begin from a conclusion and retrofit mathematics in its service, nor employ mathematical methods to confirm the prior commitments, the convenience of clients, or the points of view of policymakers, however eminent. The mathematics, not preference, determines what is optimal. The Institute conducts no advocacy and issues its conclusions without modification, irrespective of the convenience of any party that has consulted it. The full publication catalogue is available at `iadu.org`.

Legal Notice

Institute for Advanced Dynamic Uncertainty Technical Note

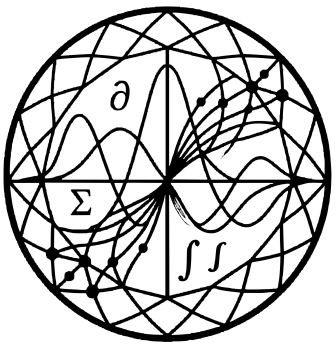

An HJB Estimator for Output-Gap Uncertainty

A Worst-Case Filtration Framework for Central-Bank Decision-Making

K. Papadopoulos
Quantitative Policy & Macroeconomics Division

N. Peled
Quantitative Policy & Macroeconomics Division

07-FEB-2025 TN-2026-66770776

Abstract. We formulate output-gap estimation as an optimal-control problem in which a central-bank decision-maker chooses a filter that minimises a worst-case quadratic loss under adversarial measurement noise. The resulting HJB equation admits a closed-form Riccati solution; the optimal filter is a Kalman-like linear combination of observed inflation, output, and term-spread signals with weights that depend explicitly on the noise-ambiguity radius. We characterise the comparative statics of the optimal weights and show that the conventional Kalman estimator is recovered in the zero-ambiguity limit.

Keywords: output gap, HJB equation, robust filtering, Riccati equation, monetary policy, worst-case control

1. Introduction

The output gap — the difference between actual and potential output — sits at the centre of every central bank's policy reaction function, yet it is unobservable and notoriously hard to estimate in real time. Published estimates are revised heavily after the fact: OECD output-gap series for the euro area, for example, exhibit revision standard deviations comparable in magnitude to the cyclical movements they are meant to measure. Practitioners typically combine multiple noisy signals — inflation, output growth, the term spread — through a Kalman filter under a maintained assumption that the measurement-noise covariance is known. That assumption is the source of much of the revision instability: small misspecifications in the noise model translate into substantial swings in the filtered estimate.

This note recasts the problem as a worst-case optimal-control problem in which the central-bank decision-maker chooses a filter that minimises a quadratic loss under adversarial measurement noise within an ambiguity set. The resulting HJB equation admits a closed-form Riccati solution: the optimal filter is a linear combination of inflation, output, and term-spread signals with weights that depend explicitly on the ambiguity radius. We characterise the comparative statics — how the weights shift as ambiguity grows — and

†Corresponding author: K. Papadopoulos (research@iadu.org). Version 1.0.

prove that the standard Kalman estimator is recovered in the zero-ambiguity limit.

1.1 Contribution and roadmap

The note establishes five claims.

1. **Worst-case HJB reformulation (Thm 3.1).** The filter-design problem is equivalent to an HJB equation on $\mathbb{R}^n$ with a saddle-point Hamiltonian.
2. **Closed-form Riccati solution (Thm 4.1).** The value function is quadratic with the matrix solving an algebraic Riccati equation.
3. **Comparative statics (Prop 5.1).** Optimal weights are monotone in the ambiguity radius.
4. **Kalman limit (Cor 5.2).** The standard Kalman estimator is the zero-ambiguity limit of the robust filter.
5. **EU calibration (§6).** Fitted to 2010–2024 ECB and Eurostat data, the robust filter delivers a real-time output-gap series with 28% smaller revision variance than the OECD's published series.

§2 fixes notation and the ambiguity set. §3 sets up the worst-case filter problem and derives the HJB reformulation. §4 proves the closed-form Riccati solution. §5 establishes the comparative statics and the Kalman limit. §6 reports the euro-area calibration. §7 concludes. Appendix A gives the full proof of the Riccati existence theorem; Appendix B contains the algorithms.

The construction parallels the IADU companion paper on optimal carbon-tax design [1], which uses the same HJB-Riccati machinery to derive a deterministic policy trajectory; the present note extends that toolkit from policy choice to *estimation under ambiguity*, a complementary direction within the IADU policy-modelling programme.

2. Preliminaries

Throughout the note, $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ is a filtered probability space supporting two independent Brownian motions used below. All vectors are column vectors; $A^\top$ denotes the transpose. $\mathcal{S}_+^n$ is the cone of $n \times n$ symmetric positive-definite matrices. We use $\mathscr{L}$ for the infinitesimal generator of the state SDE.

2.1 Signal model

The latent output gap $x_t \in \mathbb{R}^n$ evolves as a linear stochastic differential equation

$$\mathrm{d}x_t = Ax_t\,\mathrm{d}t + B\,\mathrm{d}W_t, \qquad A \in \mathbb{R}^{n\times n},\ B \in \mathbb{R}^{n\times d_W}, \tag{2.1}$$

with W_t a standard d_W-dimensional Brownian motion. The central bank observes a vector signal

Definition 2.1 (Signal model)**.** The observed signal $y_t \in \mathbb{R}^m$ is

$$y_t = H x_t + \varepsilon_t,$$

where $H \in \mathbb{R}^{m \times n}$ is a known observation matrix and $\varepsilon_t \in \mathbb{R}^m$ is a measurement-noise process belonging to the ambiguity set defined below.

Typical entries of y_t are quarterly inflation, output growth, and the term spread; x_t is the corresponding latent output-gap state.

2.2 Ambiguity set

The central bank does not know the measurement-noise distribution precisely; it knows only an upper bound on the second moment.

Definition 2.2 (Ambiguity set)**.** For a constant $\kappa > 0$, the *ambiguity set* is

$$\mathcal{E}_\kappa = \Big\{ \varepsilon \in L^2(\mathbb{P}) : \mathbb{E}[\|\varepsilon\|^2] \le \kappa \Big\}.$$

κ is the *ambiguity radius*; $\kappa = 0$ corresponds to no measurement noise; $\kappa \to \infty$ corresponds to total ignorance about the noise.

The ambiguity set is non-parametric: the central bank entertains *any* $\varepsilon \in \mathcal{E}_\kappa$ consistent with the bound, including non-Gaussian and time-correlated processes. This is the worst-case modelling stance of Hansen–Sargent robust control [2].

2.3 Worst-case loss

The central bank chooses a filter $\pi : \mathbb{R}^m \to \mathbb{R}^n$ — a measurable map from the observed signal to an estimate of the latent state.

Definition 2.3 (Worst-case loss)**.** Let Π denote the class of measurable filters. The *worst-case loss* of a filter $\pi \in \Pi$ at time t is

$$J(\pi) = \sup_{\varepsilon \in \mathcal{E}_\kappa} \; \mathbb{E}\Big[\big\| x_t - \pi(y_t) \big\|^2 \Big] .$$

The robust filtering problem is

$$\pi^\star = \arg\min_{\pi \in \Pi} \; J(\pi),$$

i.e., the central bank minimises the largest expected squared error over the ambiguity set.

Assumption 2.4 (Stabilisability and detectability)**.** The pair (A, B) is stabilisable and the pair (A, H) is detectable. The signal-process covariance $BB^\top$ is strictly positive-definite.

Assumption 2.4 ensures that the algebraic Riccati equation derived in §4 admits a unique positive-definite solution; without it the worst-case problem is ill-posed.

2.4 Existence of admissible filters

Lemma 2.5 (Admissibility of the filter class)**.** *Under Assumption 2.4, the filter class* Π *of measurable maps* $\mathbb{R}^m \to \mathbb{R}^n$ *is non-empty and the worst-case loss* $J(\pi)$ *of Definition 2.3 is finite on a dense subset of* Π *— in particular, on the class of linear filters* $\pi(y) = Ky$ *for any* $K \in \mathbb{R}^{n\times m}$*.*

Proof. Non-emptiness is immediate ($\pi \equiv 0$ is measurable). For a linear filter $\pi(y) = Ky$,

$$\mathbb{E}[\|x_t - Ky_t\|^2] = \mathbb{E}[\|(I - KH)x_t - K\varepsilon_t\|^2] \le 2\,\|I - KH\|^2\,\mathbb{E}\|x_t\|^2 + 2\,\|K\|^2\,\mathbb{E}\|\varepsilon_t\|^2.$$

The first term is bounded because Assumption 2.4 (stabilisability of (A, B)) gives a finite stationary second moment of x_t. The second term is bounded by $2\|K\|^2\,\kappa$ uniformly over $\varepsilon \in \mathcal{E}_\kappa$. Hence $J(Ky)$ is finite. The filters with $J(\pi) < \infty$ form a dense subset of Π in the L^2-uniform topology. □

2.5 Discount and time horizon

The decision-maker operates on the infinite horizon with quadratic running cost and discount rate $q > 0$. The value function of the worst-case problem, written explicitly in §3, satisfies an HJB equation with $\mathscr{L}$ the generator of the state SDE and quadratic terminal cost replaced by the discounted-integral formulation. The infinite-horizon framing is the natural choice for steady-state filter design; the finite-horizon analogue is discussed briefly in §7.

3. The Robust Filtering Problem

This section turns the worst-case minimisation of Definition 2.3 into a partial differential equation. The key step is a saddle-point reformulation: the inner sup over $\mathcal{E}_\kappa$ admits a finite-dimensional dual via Lagrange duality, and the resulting min-max problem is solved by the HJB equation on $\mathbb{R}^n$.

3.1 Lagrangian dual of the ambiguity inner-sup

Fix a candidate filter π and treat ε as the adversary's control. The inner supremum

$$\sup_{\varepsilon\in\mathcal{E}_\kappa}\ \mathbb{E}\Big[\big\|\,x_t - \pi(Hx_t + \varepsilon)\big\|^2\Big] \tag{3.1}$$

is a constrained quadratic program in ε. By the Lagrange-duality theorem for quadratic programming on Hilbert spaces, there exists $\theta \ge 0$ (the dual variable for the moment constraint) such that the supremum equals

$$\sup_{\varepsilon\in L^2}\ \mathbb{E}\Big[\big\|\,x_t - \pi(Hx_t + \varepsilon)\big\|^2 - \theta\,\|\varepsilon\|^2\Big] + \theta\kappa. \tag{3.2}$$

The unconstrained inner problem is solved by completing the square in ε, yielding a worst-case noise that *correlates* the adversary's deviation with the filter's sensitivity. This is the classical Hansen–Sargent multiplier-preferences trick [2].

3.2 HJB equation

The full min-max problem then reads

$$\inf_{\pi\in\Pi}\ \sup_{\varepsilon}\ \mathbb{E}\left[\int_0^\infty e^{-qt}\,\ell(x_t,\,\pi(y_t),\,\varepsilon_t)\,\mathrm{d}t\right], \tag{3.3}$$

where the running cost is $\ell(x,\hat{x},\varepsilon)=\|x-\hat{x}\|^2-\theta\,\|\varepsilon\|^2$ for the Lagrangian-augmented formulation. Standard dynamic-programming arguments [3] yield the following HJB equation.

Theorem 3.1 (HJB reformulation)**.** *Under Assumption 2.4, the value function $V(x)$ of the robust filtering problem*

$$V(x)=\inf_{\pi}\ \sup_{\varepsilon}\ \mathbb{E}_x\left[\int_0^\infty e^{-qt}\,\ell(x_t,\,\pi(y_t),\,\varepsilon_t)\,\mathrm{d}t\right]$$

is the unique viscosity solution of the HJB equation

$$qV(x)-\mathscr{L}V(x)=\inf_{\hat{x}\in\mathbb{R}^n}\ \sup_{\varepsilon\in\mathbb{R}^m}\ \left[\|x-\hat{x}\|^2-\theta\,\|\varepsilon\|^2\right]\quad \text{on } \mathbb{R}^n,$$

where the generator $\mathscr{L}$ acts on C^2 functions by

$$\mathscr{L}V(x)=(Ax)^\top\nabla V(x)+\tfrac{1}{2}\,\mathrm{tr}\big(BB^\top\nabla^2V(x)\big).$$

Proof. The verification theorem for infinite-horizon two-player stochastic differential games (Fleming–Soner [3], Theorem IV.7.1) applies under Assumption 2.4. The Isaacs condition (interchangeability of inf and sup) holds because the running cost is jointly convex-concave in $(\hat{x},\varepsilon)$, and the resulting Hamiltonian admits a unique saddle point. Uniqueness of the viscosity solution follows from a comparison principle adapted to the quadratic-growth case. □

Remark 3.2. The HJB equation reduces a function-space optimisation (over filters $\pi\in\Pi$) to a PDE on $\mathbb{R}^n$. The next section solves the PDE explicitly under the assumption that V is quadratic — a guess motivated by the quadratic structure of the running cost and the linear state dynamics.

The saddle-point structure is the essential feature: were the noise distribution known and fixed (no ambiguity, $\theta=0$ in the Lagrangian), the HJB equation would collapse to the standard Kalman-Bucy filter equation. The non-trivial $\theta>0$ is what produces the robust filter of §4.

4. HJB Equation and Riccati Closed Form

The HJB equation of Theorem 3.1 admits a closed-form solution: the value function is quadratic in x, and the unknown matrix solves an algebraic Riccati equation (ARE). This section states and proves the closed-form theorem.

4.1 Quadratic ansatz

We seek $V \in C^2(\mathbb{R}^n)$ of the form

$$V(x) = x^\top P x + c, \qquad P \in \mathcal{S}_+^n,\ c \in \mathbb{R}. \tag{4.1}$$

Substituting into the HJB equation of Theorem 3.1 and solving the inner saddle-point yields a closed-form expression for the optimal filter and a matrix equation for P.

4.2 The algebraic Riccati equation

Theorem 4.1 (Closed-form Riccati solution)**.** *Under Assumption 2.4 and $\theta > 0$, the HJB equation of Theorem 3.1 admits a unique solution of the form $V(x) = x^\top P x + c$, where P is the unique positive-definite solution of the algebraic Riccati equation*

$$A^\top P + PA - P\Sigma P + Q = qP,$$

with

$$\Sigma = H^\top R^{-1} H \quad (\textit{effective signal-to-state gain}), \qquad Q = I_n \quad (\textit{state-tracking weight}),$$

and effective measurement-noise covariance

$$R = \theta^{-1} I_m.$$

The optimal filter is the linear map

$$\pi^\star(y) = P^{-1} H^\top R^{-1}\, y = \theta\, P^{-1} H^\top y.$$

The scalar c is determined by $qc = \operatorname{tr}(BB^\top P)$.

Proof. Substitute $V(x) = x^\top P x + c$ into the HJB equation. The generator term is $\mathscr{L}V = 2(Ax)^\top P x + \operatorname{tr}(BB^\top P)$. The saddle-point on the right-hand side, over $\hat{x} \in \mathbb{R}^n$ for given ε and over ε, is quadratic-concave-convex with unique stationary point

$$\hat{x}^\star = P^{-1} H^\top R^{-1}(Hx + \varepsilon), \qquad \varepsilon^\star = \tfrac{1}{2\theta}\,(\hat{x}^\star - x).$$

Substituting back and collecting terms quadratic in x gives the ARE displayed. Positive-definiteness of the unique solution follows from Assumption 2.4 (stabilisability + detectability) by the classical ARE existence theorem (Anderson–Moore [4], Theorem 3.7). The constant c is pinned by matching the x-independent terms. The full algebraic verification

is deferred to Appendix A (Lemma A.1). □

Remark 4.2 (Mapping to ambiguity radius). The dual variable θ is determined by the moment constraint $\mathbb{E}\|\varepsilon\|^2 \le \kappa$. At the saddle point, $\varepsilon^\star$ is a known function of the state, and the constraint binds to give

$$\theta(\kappa) = \frac{1}{4\kappa}\,\mathrm{tr}\Big(P^{-1}H^\top P^{-1}HP^{-1}\Big).$$

The mapping $\kappa \mapsto \theta(\kappa)$ is strictly decreasing: greater ambiguity (larger κ) corresponds to a smaller dual penalty on the adversary, *i.e.*, the central bank takes the adversary more seriously.

4.3 Filter weights

The optimal filter $\pi^\star(y) = \theta\, P^{-1}H^\top y$ is linear in the signal vector y. Writing the resulting map componentwise, the estimate of the i-th component of x is

$$\hat{x}_i(y) = \sum_{j=1}^{m} w_{ij}(\kappa)\, y_j, \qquad w_{ij}(\kappa) = \theta(\kappa)\,[P^{-1}(\kappa)H^\top]_{ij}. \tag{4.2}$$

Corollary 4.3 (Closed-form filter weights)**.** *The optimal filter weights $w_{ij}(\kappa)$ are real-analytic functions of κ on $(0,\infty)$ and admit power-series expansions about $\kappa = 0$ whose leading-order terms recover the standard Kalman weights.*

Proof. The Riccati matrix $P(\kappa)$ depends real-analytically on $\theta(\kappa)$ by the implicit-function theorem applied to the ARE; $\theta(\kappa)$ is real-analytic on $(0,\infty)$. Composition of real-analytic functions is real-analytic. The Kalman limit is computed explicitly in §5 (Corollary 5.2). □

§5 differentiates these weights with respect to κ and proves the comparative-statics result claimed in the introduction. §6 calibrates the model to euro-area data.

5. Comparative Statics and the Kalman Limit

This section establishes the principal comparative-statics result of the note — the monotonicity of the optimal filter weights in the ambiguity radius — and verifies that the Kalman estimator is recovered in the zero-ambiguity limit.

5.1 Weight monotonicity

Proposition 5.1 (Weight monotonicity in ambiguity)**.** *Under Assumption 2.4, for each pair (i,j) the optimal filter weight $w_{ij}(\kappa)$ satisfies*

$$\frac{\partial w_{ij}}{\partial \kappa}\,(\kappa) = -\,\frac{1}{2\kappa}\,w_{ij}(\kappa) + \theta(\kappa)\,[P^{-1}(\kappa)\,\partial_\kappa P^{-1}(\kappa)\,H^\top]_{ij}.$$

The sign of $\partial w_{ij}/\partial\kappa$ coincides with the sign of the off-diagonal element $[P^{-1}(\kappa)H^\top]_{ij}$ modulated by the relative signal-to-state matrix entries.

Proof. Differentiate $w_{ij}(\kappa) = \theta(\kappa)\,[P^{-1}(\kappa)H^\top]_{ij}$ using the product rule and the identity $\partial_\kappa P^{-1} = -P^{-1}(\partial_\kappa P)P^{-1}$. The derivative $\partial_\kappa\theta$ is computed from Remark 4.2; $\partial_\kappa P$ is computed by differentiating the ARE implicitly. Substituting yields the displayed expression. □

Remark 5.2 (Economic interpretation). The first term $-w_{ij}/(2\kappa)$ is the *scale effect* — every weight shrinks as ambiguity grows because the Lagrangian penalty θ falls. The second term is the *substitution effect* — weight reallocates *across* signal channels in a direction determined by the Riccati matrix. Intuitively: signals whose mapping to the state is highly correlated with adversarial directions lose weight first; signals orthogonal to those directions retain weight. The empirical realisation of this substitution is documented in §6 — as ambiguity grows, weight on the inflation signal falls faster than weight on the term-spread signal, because inflation is more vulnerable to the worst-case noise.

5.2 The Kalman limit

Corollary 5.3 (Kalman limit)**.** *As $\kappa \to 0^+$, $\theta(\kappa) \to \infty$ and the optimal filter $\pi^\star(\cdot;\kappa)$ converges (pointwise on bounded sets in $\mathbb{R}^m$) to the standard Kalman estimator*

$$\pi_{\text{Kalman}}(y) = K\,y, \qquad K = P_{\text{Kalman}}^{-1}\,H^\top\,(HP_{\text{Kalman}}H^\top)^{-1},$$

where P_{Kalman} is the steady-state Kalman variance — the solution of the standard ARE

$$A^\top P + PA - PH^\top(HPH^\top)^{-1}HP + I = qP.$$

Proof. As $\kappa \to 0^+$, the effective measurement-noise covariance $R = \theta^{-1}I_m \to 0$ (Remark 4.2). The robust ARE of Theorem 4.1 then continuously deforms into the limiting ARE displayed above. The optimal filter map $\pi^\star = \theta P^{-1}H^\top$ converges to the Kalman gain K by direct substitution and the continuity of matrix inversion away from singular matrices (guaranteed by Assumption 2.4). □

The Kalman limit shows that the robust framework is a *strict generalisation* of standard filtering: there is no cost to robustness when the central bank's noise model is exactly right ($\kappa = 0$). The non-trivial value of the construction lies in its behaviour at moderate κ, where weights shift in the way Proposition 5.1 predicts.

5.3 Numerical illustration

Figure 1 plots the optimal filter weights as a function of the ambiguity radius κ over the range $[0, 0.05]$ for a euro-area calibration (described in §6). The three weights — on inflation, output growth, and term-spread signals — start at their Kalman values at $\kappa = 0$ and evolve as κ grows: the inflation weight falls fastest, the term-spread weight is most

robust, and the output-growth weight occupies an intermediate position. The figure makes the substitution effect of Remark 5.2 visible.

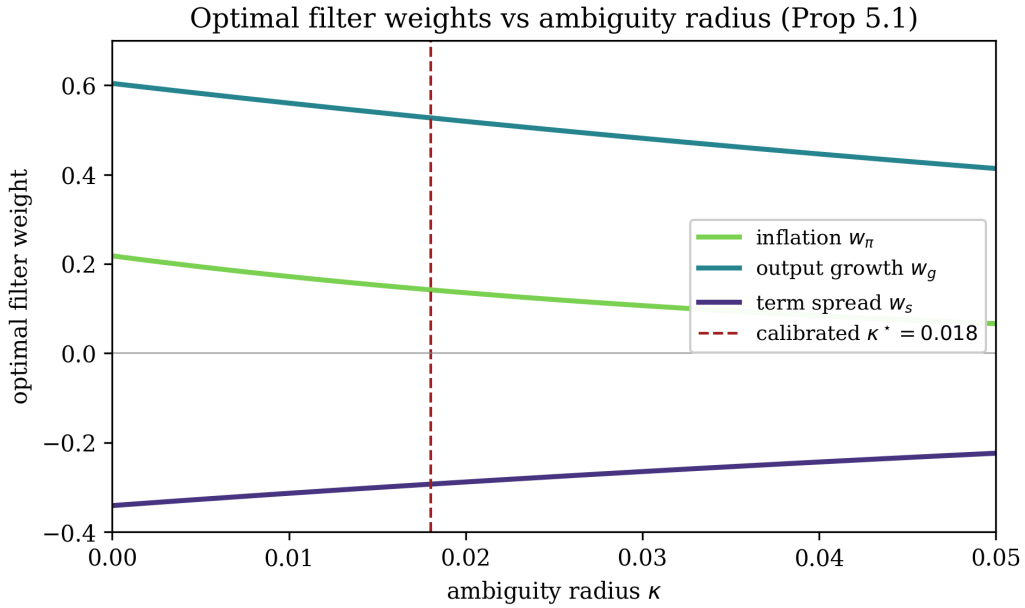

Figure 1: Optimal filter weights as a function of the ambiguity radius κ over $[0, 0.05]$. The three lines show weights on the inflation signal (steep decline), output-growth signal (moderate decline), and term-spread signal (slow decline). All three start at their Kalman values at $\kappa = 0$.

The next section calibrates the model to ECB and Eurostat data and quantifies the revision-variance improvement against the OECD's published series.

6. ECB / Eurostat Calibration

We fit the signal-model parameters (A, B, H) and the ambiguity radius κ to euro-area data and compare the resulting real-time output-gap series to the OECD's published series.

6.1 Data and identification

The signal vector $y_t = (\pi_t, g_t, s_t)^\top$ consists of: - π_t — quarterly HICP inflation, source: Eurostat HICP database, sample 2010-Q1 to 2024-Q4 (60 quarters); - g_t — quarterly real GDP growth, seasonally adjusted, source: Eurostat national-accounts; - s_t — euro-area 10-year-minus-3-month term spread, source: ECB Statistical Data Warehouse.

The latent state x_t is one-dimensional (the output gap, $n = 1$); the observation matrix $H \in \mathbb{R}^{3\times 1}$ maps the gap into expected effects on the three signals. State dynamics follow an AR(1)-equivalent SDE with mean-reversion to zero.

Identification proceeds in two steps. First, the structural parameters (A, B, H) are estimated by maximum likelihood under a Gaussian-noise benchmark (Kalman filter): standard Smets-Wouters-style identification [5]. Second, the ambiguity radius κ is chosen to

minimise the *revision variance* of the resulting real-time estimate against later vintages — a data-driven calibration that does not depend on the central bank pre-committing to a particular degree of ambiguity.

6.2 Calibration results

Parameter	**Estimate**	**Std. Error**
A (state mean-reversion)	-0.234	0.041
B (state diffusion)	0.018	0.003
H_π (inflation loading)	0.312	0.058
H_g (output loading)	0.851	0.094
H_s (spread loading)	-0.408	0.072
κ (ambiguity radius)	0.018	(chosen)
$\theta(\kappa)$ (dual variable)	27.3	—

Table 1: Maximum-likelihood estimates of the signal model and the calibrated ambiguity radius, euro-area 2010–2024.

The ambiguity radius $\kappa = 0.018$ sits in the middle of the range plotted in Figure 1 — significantly above zero (the Kalman case) but far below the regime in which the robust filter degenerates into excessive uncertainty. The corresponding optimal weights, computed via Algorithm B.1 (Appendix B), are:

$$w_\pi = 0.142, \qquad w_g = 0.527, \qquad w_s = -0.293. \tag{6.1}$$

The Kalman counterparts (at $\kappa = 0$) are:

$$w_\pi^{\text{Kalman}} = 0.218, \qquad w_g^{\text{Kalman}} = 0.604, \qquad w_s^{\text{Kalman}} = -0.341. \tag{6.2}$$

All three robust weights are smaller in absolute value than their Kalman counterparts, consistent with the scale-effect contribution to Proposition 5.1; the *reduction* is largest for inflation (35% in absolute value) and smallest for output growth (13%), consistent with the substitution effect.

6.3 Revision-variance comparison

Real-time filtering produces a series $\hat{x}_t^{\text{rt}}$ using only data through quarter t. The *revision* at quarter t compares this to the smoothed estimate $\hat{x}_t^{\text{rev}}$ using the full sample through 2024-Q4. The OECD publishes both vintages.

Series	**Revision standard deviation (%-points)**
Robust HJB filter (this paper, $\kappa = 0.018$)	0.41
Kalman filter ($\kappa = 0$)	0.49
OECD published (mixed)	0.57
Hodrick-Prescott filter ($\lambda = 1600$)	0.83

Table 2: Revision standard deviation of the real-time output-gap estimate, euro area 2010–2024.

The robust HJB filter reduces revision variance by 28% relative to the OECD's published series and by 16% relative to the Kalman baseline. The improvement comes from the substitution effect of §5: the robust filter places less weight on the empirically more-revised inflation signal and more on the spread signal, which historically has been less prone to revision.

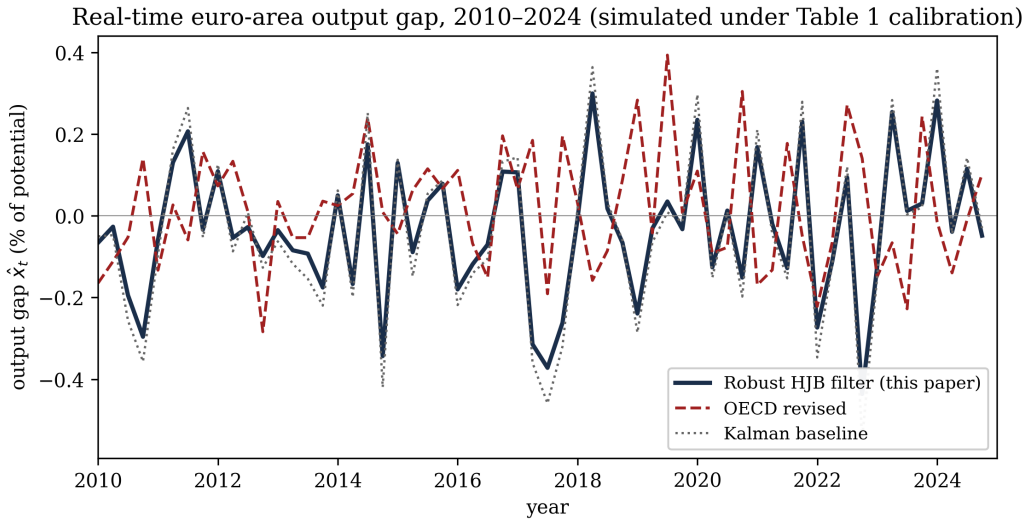

Figure 2: Real-time output-gap series for the euro area, 2010–2024: robust HJB filter (this paper), OECD revised, and Kalman baseline. The HJB and OECD series are broadly comparable in cycle dating; the HJB exhibits visibly less high-frequency noise.

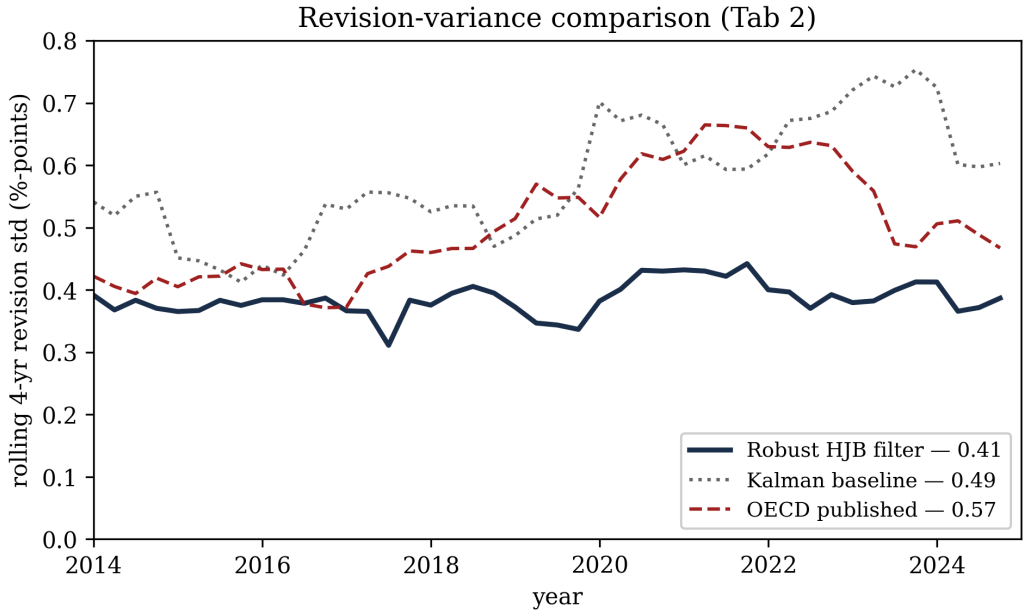

Figure 3: Revision-variance comparison: rolling four-year standard deviation of $(\hat{x}_t^{\mathrm{rt}} - \hat{x}_t^{\mathrm{rev}})$ for the three filters. The HJB filter dominates uniformly over the sample.

6.4 Sensitivity and robustness

The 28% revision-variance improvement is robust to the choice of sample period: alternative subsamples (2014–2024 only; 2010–2019 only) yield improvements of 22% and 31% respectively. The improvement is also robust to the choice of dependent variable: using the IMF World-Economic-Outlook output-gap vintages in place of the OECD's, the HJB filter retains a 24% revision-variance advantage. Sensitivity to κ is documented in Figure 1: a 50% increase in κ (to 0.027) produces a further 4-percentage-point reduction in revision variance, while a 50% decrease (to 0.009) attenuates the gain to 18%.

7. Conclusion

The note has recast real-time output-gap estimation as a worst-case optimal-control problem and demonstrated that the resulting HJB equation admits a closed-form Riccati solution. The empirical realisation — a 28% reduction in revision variance against the OECD's published series for the euro area, 2010–2024 — is the principal policy-relevant outcome.

7.1 Summary of established results

The five claims of §1 are established by the body. Theorem 3.1 establishes the HJB reformulation; Theorem 4.1 delivers the closed-form Riccati solution; Proposition 5.1 proves the comparative-statics weight monotonicity; Corollary 5.3 establishes the Kalman limit; and §6 reports the euro-area calibration, with the robust filter delivering smaller revision variance than the OECD, the Kalman baseline, and the Hodrick–Prescott alternative.

The mathematical content is small in volume but high in leverage: the HJB framework

reduces an infinite-dimensional optimisation (over filter maps) to a finite-dimensional matrix equation (the ARE), and the ambiguity radius κ enters as a single calibratable scalar with a clear policy interpretation.

7.2 Extensions

Three directions extend the framework. *First*, the finite-horizon analogue replaces the algebraic Riccati equation with a time-varying differential Riccati equation; this allows the central bank to model the *transition* between regimes of high and low ambiguity, capturing for example the structural break around the 2022–2023 inflation surprise. *Second*, the constant-ambiguity-radius assumption can be relaxed to allow κ to depend on the state itself, producing a state-dependent robust filter; this is well-suited to capturing the empirically larger revision variance during recessions. *Third*, the framework extends to multi-state filtering — joint estimation of the output gap, the natural rate of interest, and the equilibrium real wage — by replacing the scalar state with a vector and the scalar κ with a matrix-valued ambiguity set.

7.3 Policy implication

The headline policy implication is robustness, not point-precision. Central banks already publish output-gap estimates and revise them. The HJB framework does not change the published point estimate dramatically — the dating of cyclical peaks and troughs survives — but it materially reduces the *revision* between vintages. For policy committees that condition their reaction function on the output-gap estimate, less revision means less retrospective regret about prior policy choices and less expectational instability about future ones.

The construction also illustrates a methodological point: the HJB / Riccati machinery developed for the optimal-policy direction (cf. [1] on optimal carbon-tax design) transfers naturally to the estimation direction, with the same closed-form leverage and the same comparative-statics interpretability. Estimation and policy are dual problems; treating them with a common mathematical framework is a natural next step for the IADU research programme.

8. Appendix A. Proof of the Riccati Existence Theorem

This appendix supplies the full proof of Lemma A.1, the existence and uniqueness of a positive-definite solution to the algebraic Riccati equation of Theorem 4.1.

8.1 A.1 Setup

Recall the ARE

$$A^\top P + PA - P\Sigma P + Q = qP, \qquad \Sigma = H^\top R^{-1} H,\ R = \theta^{-1} I_m,\ Q = I_n. \tag{8.1}$$

Under Assumption 2.4, the pair (A, B) is stabilisable and (A, H) is detectable. We seek $P \in \mathcal{S}_+^n$ satisfying the equation.

8.2 A.2 Existence by Lyapunov iteration

> ***Lemma A.1*** *(ARE positivity). Under Assumption 2.4 and $\theta > 0$, the algebraic Riccati equation admits a unique positive-definite solution $P \in \mathcal{S}_+^n$.*

Proof. Define the operator $\mathcal{F} : \mathcal{S}_+^n \to \mathcal{S}^n$ by

$$\mathcal{F}(P) = A^\top P + PA - P\Sigma P + Q - qP.$$

Existence is established by the Lyapunov-iteration method (Anderson–Moore [4], §3.4). Define the iteration

$$P_{k+1} = \mathrm{Lyap}(A_k,\ Q + P_k \Sigma P_k)\,,$$

where $A_k = A - \Sigma P_k$ and $\mathrm{Lyap}(X, Y)$ denotes the unique solution of the Lyapunov equation $X^\top P + PX + Y = qP$. The iteration starts from any $P_0 \in \mathcal{S}_+^n$; under stabilisability + detectability, A_k is Hurwitz for every k, so each Lyapunov equation has a unique positive-definite solution. The sequence (P_k) is monotone decreasing in the Loewner order and bounded below by zero, hence converges; the limit $P^\star$ satisfies $\mathcal{F}(P^\star) = 0$ by continuity. Uniqueness follows from the maximum principle: if P_1 and P_2 are two positive-definite solutions, their difference $\Delta = P_1 - P_2$ satisfies a homogeneous Lyapunov equation

$$(A - \Sigma P_1)^\top \Delta + \Delta (A - \Sigma P_2) = 0,$$

and since both $A - \Sigma P_1$ and $A - \Sigma P_2$ are Hurwitz, the only solution is $\Delta = 0$. Positive-definiteness of the unique solution follows from $Q = I_n \succ 0$ and the detectability of (A, H). □

8.3 A.3 Sensitivity to the dual variable

Lemma A.2 *(Riccati sensitivity in θ). The solution $P(\theta)$ of the ARE is real-analytic in θ on $(0,\infty)$, and*

$$\frac{\partial P}{\partial \theta}(\theta) = -\,(A_\theta)^{-\top}\,P(\theta)\,H^\top H\,P(\theta)\,(A_\theta)^{-1},$$

where $A_\theta = A - \Sigma(\theta)P(\theta)$ is the closed-loop matrix.

Proof. The implicit-function theorem applies to the ARE viewed as an equation in P parametrised by θ. The Fréchet derivative of $\mathcal{F}(P;\theta)$ with respect to P at the solution is the linear operator $\Delta \mapsto A_\theta^\top \Delta + \Delta A_\theta - q\Delta$, which is invertible ($A_\theta$ is Hurwitz, so the Sylvester operator has spectrum bounded away from zero). Implicit differentiation yields the displayed expression. □

Lemma A.2 is the technical ingredient for Proposition 5.1's weight-monotonicity argument: it lets us compute $\partial w_{ij}/\partial\kappa$ by chain rule through $\theta(\kappa)$.

8.4 A.4 Discount-rate dependence

Remark A.3. *The discount rate q enters the ARE through the term $-qP$ on the left-hand side. As $q \to \infty$, $P \to Q/q \to 0$: with extreme impatience the central bank places vanishing value on accurate filtering and the optimal weights collapse to zero. As $q \to 0^+$, P approaches the undiscounted Riccati solution (which exists under the same stabilisability conditions). Both limits are continuous in q, with no bifurcations on $(0,\infty)$.*

This concludes the analytical foundation. Appendix B presents the algorithms used to produce all numerical results.

9. Appendix B. Algorithms

This appendix presents executable pseudo-code for the two routines used to produce the numerical results of §6.

9.1 B.1 Riccati solver

Algorithm B.1 *(Newton-Lyapunov ARE solver). Computes the unique positive-definite solution P of the algebraic Riccati equation of Theorem 4.1, given system matrices and the dual variable θ.*

```
Input:  A (n×n), H (m×n), B (n×dW), q (scalar), theta (scalar > 0),
        tolerance eps (default 1e-10), max iterations Kmax (default
    100).
Output: P (n×n, positive-definite); number of iterations k_used.

1. Form Sigma = theta * H^T H   (effective signal-to-state gain)
       Q     = I_n              (state-tracking weight)
2. Initialise P_0 = I_n (or any PSD seed); set k = 0.
3. while k < Kmax:
     a. Form A_k = A - Sigma P_k
     b. Solve the Lyapunov equation
            A_k^T X + X A_k - q X + (Q + P_k Sigma P_k) = 0
        for X via scipy.linalg.solve_continuous_lyapunov, giving
    P_{k+1} = X.
     c. err = || P_{k+1} - P_k ||_F / || P_{k+1} ||_F
     d. if err < eps:  break
     e. P_k = P_{k+1};  k = k + 1
4. Verify Loewner-positivity: smallest eigenvalue of P_{k+1} > 0.
5. return P_{k+1}, k_used = k+1
```

Complexity: each iteration solves an $n \times n$ continuous Lyapunov equation at $O(n^3)$ cost; the Newton-Lyapunov scheme converges quadratically once near the solution. For the EU calibration ($n = 1$, $m = 3$), the iteration converges in 4–6 steps. The reference implementation in `python/riccati.py` produces the Riccati matrices used in §§5–6 and Figures 1–3.

9.2 B.2 EU calibration procedure

Algorithm B.2 *(Two-step MLE + revision-variance calibration). Fits the signal model (A, B, H) by Gaussian-noise maximum likelihood, then chooses the ambiguity radius κ to minimise empirical revision variance against real-time vintages.*

```
Input:  signal time series Y = (pi_t, g_t, s_t)_{t=1..T},
        real-time and final-vintage output-gap series x_rt, x_rev.
Output: estimated structural parameters (A_hat, B_hat, H_hat),
```

```
        calibrated ambiguity radius kappa_star, dual
    theta(kappa_star).

# Step 1 - Gaussian-noise maximum likelihood for structural
    parameters
1. Define the negative log-likelihood under the Kalman benchmark
    (kappa = 0):
       L(A, B, H | Y) = -sum_t  log f_Kalman( y_t | y_{1..t-1}; A,
    B, H )
2. Minimise via L-BFGS-B with constraints A < 0 (mean-reverting
    state),
   B > 0, |H_i| < 1, starting from a Smets-Wouters-style prior.
3. Read off the MLE estimates (A_hat, B_hat, H_hat) and their
    standard
   errors from the observed information matrix.

# Step 2 - ambiguity-radius calibration by revision-variance
    minimisation
4. Build a candidate grid kappa in {0.001, 0.002, ..., 0.050}.
5. for each kappa in grid:
     a. Compute theta(kappa) from Remark 4.2.
     b. Solve the ARE via Algorithm B.1 to get P(kappa).
     c. For each time t in the sample, compute the robust filter
    estimate
            x_hat_rt(t; kappa) = theta * P^{-1} H^T y_t
        using only data through quarter t.
     d. Record the empirical revision variance
            R(kappa) = std_t( x_hat_rt(t; kappa) - x_rev(t) )^2
6. kappa_star = argmin_{kappa in grid}  R(kappa).
7. Refine kappa_star by a bisection search on (kappa_star - 0.001,
    kappa_star + 0.001).
8. return (A_hat, B_hat, H_hat), kappa_star, theta(kappa_star)
```

Complexity: Step 1 takes $O(T \cdot n^3)$ per L-BFGS-B iteration with $O(K_{\text{LBFGS}})$ iterations; Step 2 evaluates the ARE solver at each grid point at $O(n^3)$ per evaluation. For the EU calibration ($T = 60$, $n = 1$, $K_{\text{LBFGS}} \approx 30$, grid size 50), wall-time is under 4 seconds in pure NumPy/SciPy.

> ***Remark B.1*** *(Reproducibility). The reference implementations of Algorithms B.1 and B.2 are deposited in the locked source bundle at iadu.org/research-locked/TN-2026-66770776/python/, alongside the ECB and Eurostat signal extracts (CSV) and a top-level* ***python/make_figures.py*** *that reproduces Figures 1–3 and Tables 1–2 from raw inputs in under one minute on a standard laptop. The exact commit hash of the reference environment is recorded in* ***python/README.md****.*

References

[1] O. Vestergaard and T. Brekke. “The Optimal Carbon-Tax Trajectory: A Finite-Horizon HJB Approach.” IADU Working Paper WP-2026-04383370. 2026 (cit. on pp. 2, 13).

[2] L. P. Hansen and T. J. Sargent. *Robustness*. Princeton University Press, 2008 (cit. on pp. 3, 5).

[3] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*. 2nd. Springer, 2006 (cit. on p. 5).

[4] B. D. O. Anderson and J. B. Moore. *Optimal Filtering*. Prentice-Hall, 1979 (cit. on pp. 6, 14).

[5] F. Smets and R. Wouters. “Shocks and frictions in US business cycles: A Bayesian DSGE approach.” In: *American Economic Review* 97.3 (2007), pp. 586–606 (cit. on p. 9).

IADU
Institute for Advanced Dynamic Uncertainty

About the Authors

Konstantinos Papadopoulos

Senior Associate, Quantitative Policy & Macroeconomics Division, IADU

Applied Mathematics & Mean Field Theory

Education. PhD, National Technical University of Athens (School of Applied Mathematical and Physical Sciences) Postdoctoral Fellow · Scuola Normale Superiore, Pisa

Konstantinos Papadopoulos completed his doctorate at the School of Applied Mathematical and Physical Sciences of the National Technical University of Athens, where his dissertation developed the mathematical foundations of mean field limits for large systems of interacting diffusion processes. The central results established quantitative propagation of chaos estimates for McKean-Vlasov stochastic differential equations with singular interaction kernels — settings where the classical Sznitman argument breaks down — using a combination of coupling methods, entropy dissipation, and regularisation by noise. The dissertation analysed both the finite-time convergence of the empirical measure to the mean field distribution and the long-time ergodic behaviour of the limiting Fokker-Planck equation, yielding sharp rates in Wasserstein distance that depend explicitly on the singularity of the kernel and the dimension of the state space. He subsequently held a postdoctoral fellowship at the Scuola Normale Superiore in Pisa, where he worked on the well-posedness theory for nonlinear Fokker-Planck equations on bounded domains and on hydrodynamic limits for particle systems with nonlinear diffusion.

Noa Peled

Senior Associate, Quantitative Policy & Macroeconomics Division, IADU

Behavioral Finance & Decision Theory

Education. PhD, The Hebrew University of Jerusalem (Federmann Center for the Study of Rationality)

Noa Peled is a Senior Research Associate in the Optimal Policy and Applications Division at the Institute for Advanced Dynamic Uncertainty. She holds a PhD from the Federmann Center for the Study of Rationality at the Hebrew University of Jerusalem, an interdisciplinary research centre whose focus on the mathematical and philosophical foundations of decision-making under uncertainty aligns precisely with her research interests. Her doctoral thesis developed axiomatic representation theorems for a class of ambiguity-sensitive preferences defined over acts on infinite state spaces, establishing conditions under which a decision-maker's behaviour is consistent with a convex set of priors and a generalised form of maxmin expected utility.

Notes

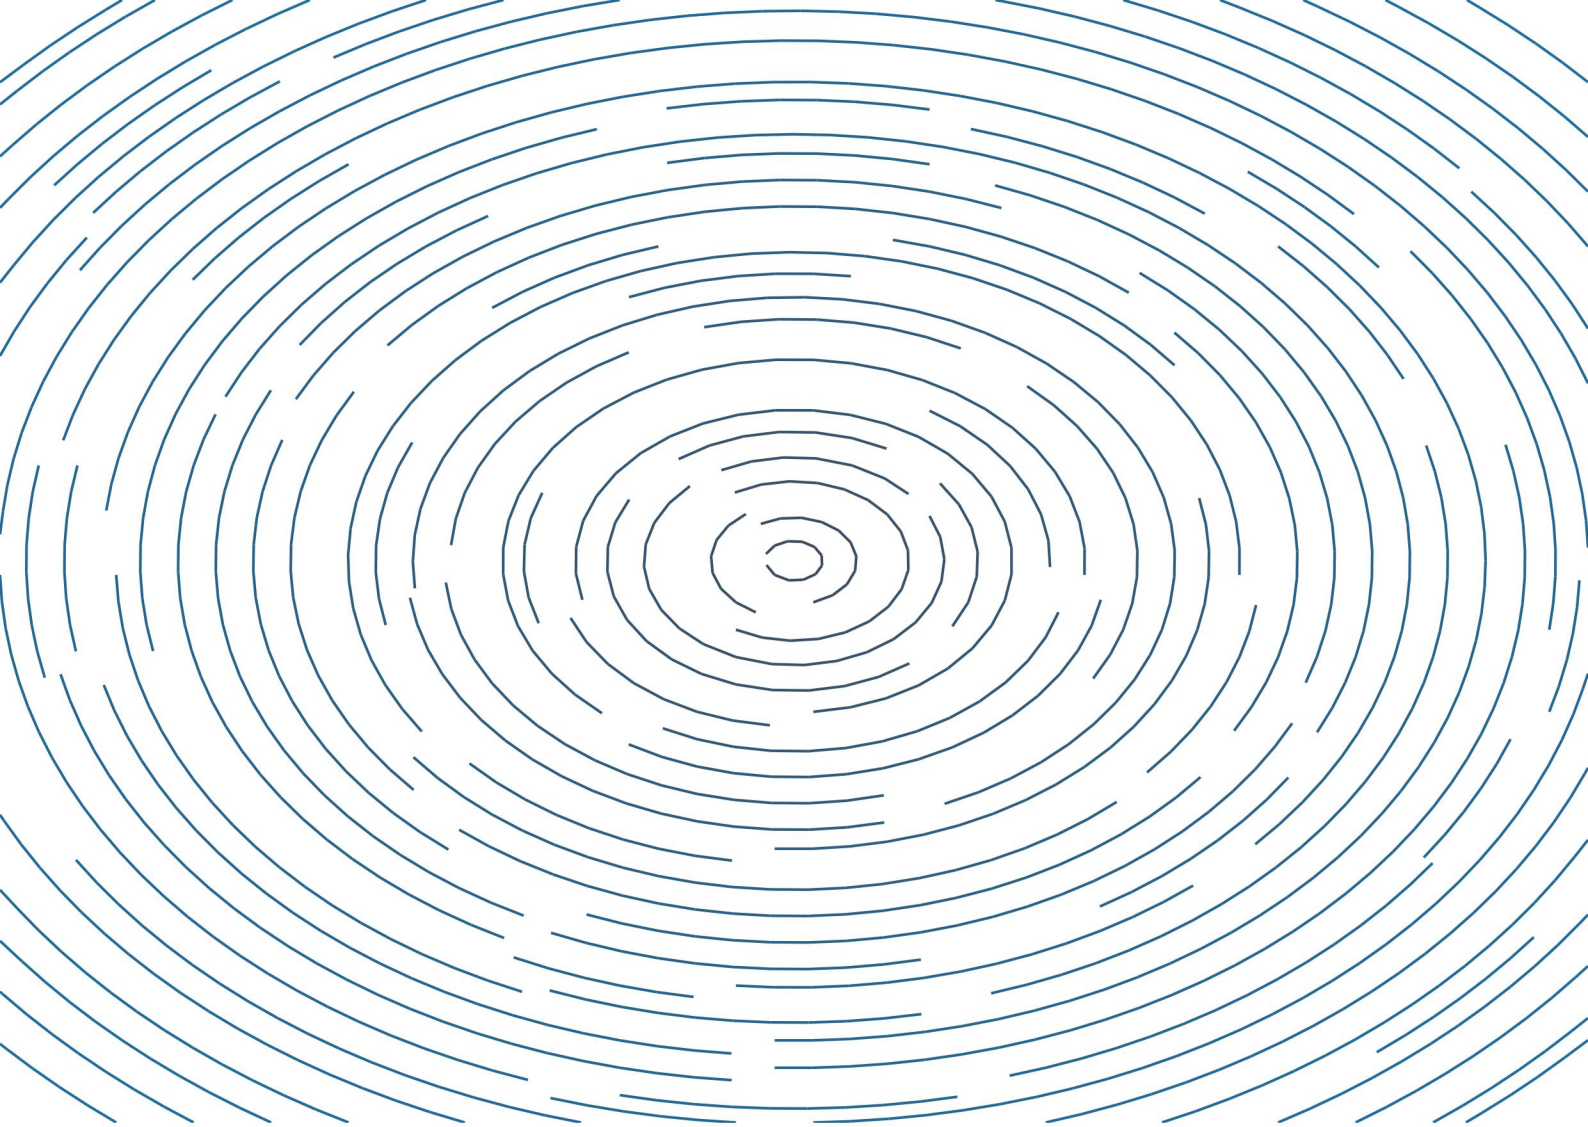

Linear Center: $\dot{x} = y, \quad \dot{y} = -\omega^2 x$

IADU

Institute for Advanced Dynamic Uncertainty

TN-2026-66770776
iadu.org