Real-Time Recovery of Structural Shocks from High-Frequency Data*†

Réka Rátfai‡

October 12, 2026

Abstract

How should a macroeconomic model whose grid is quarterly read data that arrive daily? Aggregating the fast series to the slow grid destroys the information a real-time observer needs most—the within-window movements that are the structural shocks themselves—and refining the grid cannot restore it under the standard representation of flows. This paper develops a continuous-time, mixed-frequency filtering framework that reads each release on its own calendar and delivers both objects of real-time interest online: the latent state and the contemporaneous structural shocks hitting the economy. Identification rests on a measurement primitive: a reported flow is the window increment of a cumulative observable and carries the state's Brownian increment—the pass-through—read instantaneously by a correlated-noise Kalman–Bucy filter. Four closed-form factors and a recovery-correlation law separate what finer sampling repairs from what only the correct representation can; Monte Carlo evidence in two New Keynesian economies verifies each factor. Daily observation of a single financial stock cuts the tracking error of a never-directly-observed state nearly thirtyfold and rescues the weakly identified shock from 0.2-0.5 to 0.85–0.94, while densification under the wrong primitive makes filtering worse, not better. The proposed misspecification tests maintain the transparency of the underlying structural constraints. On US data a model-free test shows official flows are averages, while asset prices carry the shocks.

JEL classification: C32; C53; E32; E37

Keywords: macro policy, macro-finance, correlated noise structure, rate-form representation, real-time structural shock recovery, Kalman–Bucy filter, temporal aggregation, mixed stock-flow data, nowcasting, mixed-frequency data, preference shocks, high-frequency monetary policy shocks, learning

1 Introduction

Real-time assessment of economic conditions is among the most consequential problems central banks, finance ministries, and financial institutions face—and it has become a

*I am grateful to Olaf Posch for insightful comments and suggestions. All remaining errors are my own.

†I thank Gerrit Meyerheim for a detailed written comment that corrected the scope of the gain-equivalence result and refined the parameter-drift argument and the statement of the measurement primitive.

‡Universität Hamburg. Email: reka.ratfai@uni-hamburg.de

problem of timing mismatch. The economy evolves continuously; the financial data are available daily or even intraday; the prices are released monthly, in form of inflation; national accounts arrive quarterly. Macroeconometric practice, built on a single discrete grid, must either aggregate the fast series down to the lowest frequency or force the model onto a grid finer than most of its data—both of which come at a price. Aggregation destroys the within-window movments a real-time observer needs the most, and a refined grid leaves a largely unobserved panel behind. The nowcasting programmes of central banks manage this mismatch with reduced-form mixed-frequency machinery (Giannone et al., 2008; Bok et al., 2018; Bànbura et al., 2013; Richardson et al., 2021) for their object of interest: the estimate of the current state of the economy. However, a policymaker navigating in real time needs two objects including the an estimate of what fundamental shock is hitting the economy now. The former task is a filtering exercise of latent variables, while the second is an identification problem for structural shocks. Nevertheless, this second object is beyond the reach of any forecast: a structural shock is a martingale increment, unpredictable from past data by construction, so it can only be read from contemporaneous observations or reconstructed *ex post*. This paper develops the filtering framework that delivers both objects online, as each release arrives, on whatever calendar it arrives.

The exercise is most naturally modeled in continuous time, a framework supported by a growing body of economic theory literature: heterogeneous-agent models specify wealth and income dynamics as diffusions (Achdou et al., 2022; Kaplan et al., 2018), the macro-finance literature builds equilibrium dynamics directly on stochastic differential equations (Brunnermeier and Sannikov, 2014), and the robustness programme of (Hansen and Sargent, 2011) treats hidden continuous-time dynamics as the object agents and econometricians alike confront; the specific laboratory of this paper is the continuous-time New Keynesian framework of (Posch, 2020). A continuous-time economy has no grid, so it accommodates any release calendar—the mismatch disappears on the model side. What remains is the measurement side. There, the continuous-time econometric tradition (Bergstrom, 1966, 1983, 1990; Phillips, 1978; Harvey, 1985; Harvey and Stock, 1993; Zadrozny, 1988) long ago established how to estimate such models from discrete data, and (Christensen et al., 2024) prove that exact discretisation makes parameter estimation invariant to the sampling frequency. The question this paper asks is which part of that invariance survives in the real-time filtering problem. Filter output improves with finer sampling under any specification; the object that does or does not carry over is the filter's distance from the attainable optimum. Under the correct flow representation that distance vanishes with the sampling interval. Under the standard representation it converges to a fixed penalty factor that no finer sampling removes, and the divergence has a single source: the measurement of flows. A measured flow is the window increment of a cumulative observable, and what that increment contains is a measurement assumption. The rate-form primitive maintained here places three pieces in it: a drift term carrying the state's level, the state's own Brownian increment (the pass-through), and measurement error. Under that primitive the structural shock is in the datum itself. The integral-form treatment, standard since (Bergstrom, 1986) and (Harvey and Stock, 1993) and exact in (Christensen et al., 2024), absorbs that increment into an augmented state and gives up its contemporaneous appearance; the rate-form representation developed here keeps it, by means of a correlated-noise filtering structure, standard in filtering theory since (Kalman-Bucy, 1962; Anderson and Moore, 1979) and employed in economics where correlated innovations are an assumed feature of asset-price dynamics under partial information

(Buccheri, 2025). In contrast, in this paper the correlation is not assumed but derived—it is the pass-through of the flow measurement—and utilized for the real-time recovery of structural shocks. Which representation a given release warrants is a property of its publication chain, decided release by release: the framework is an assignment rule with instrumentation, not an advocacy of any of the representations.

The paper makes four contributions. The first is the framework itself: a rate-form, correlated-noise representation of mixed stock–flow systems, in which three filters—the continuous-time correlated-noise Kalman–Bucy filter, the exact discrete filter of the integral form, and the naive Euler scheme—occupy one taxonomy, ordered by how much of the pass-through each retains and at what temporal resolution. The general matrix machinery, proofs, and implementation criteria are collected in the companion technical note (Rátfai, 2026). The second contribution prices every comparison within that taxonomy in closed form: four factors and one law, each attached to a single mechanism, each depending on structural persistence and the sampling interval only through their product. Finer sampling removes the temporal-aggregation factor; no sampling frequency removes the pass-through penalty; the structural-penalty pair prices the Euler shortcut under each primitive; the level-resolution factor establishes the stock channel's near-immunity; and the recovery-correlation law makes the shock-recovery ranking a corollary of the tracking ranking.

The third contribution takes the framework to the mixed frequency by Monte Carlo, in two New Keynesian economies—a single-shock feedback-rule model in which every factor is verified, and a two-shock partial-adjustment model in which the observables play the role of separating increments rather than reinforce a single one. Daily observation of the policy rate alone cuts the tracking error of the never-directly-observed demand state by a factor of twenty-eight and rescues the weakly identified shock of each configuration. Identification arrives unevenly within the quarter, as a wave whose closed-form per-position ceilings the filter attains. And densifying the observation grid under the wrong primitive amplifies the misspecification instead of repairing it, so the take-away is to correct the representation before refining the grid. The fourth contribution is a collection of stated misspecification tests—moment restrictions that assign the primitive, innovation and recovered-shock diagnostics, parameter stability across sampling frequencies, and a two-sided validation—that keeps the full structural commitment testable, release by release. A by-product of the observation-count geometry runs through all the results: proximity to stochastic singularity, detrimental for likelihood-based estimation, is favourable for the online recovery studied here.

The empirical section contributes with developing and applying tests to assign the correct flow representation to real-time data series capturing GDP and inflation. The study in Section 8 finds that US official flows, indeed, do not retain much structural content, however, adding other asset price data can recover the incremental information about the structural shocks and strengthens identification substantially, which is a result of the study. The *ex ante* assessment of flow-type variables are essential to avoid distortions and recover the highest information content that is physically attainable from an economic system in mixed-frequency. Tests can be deployed and are conclusive entirely model-free (with no economic information added), but testing with model restrictions is also displayed for completeness.

The analytical results are established for a scalar latent state driven by a single structural shock, with parameters treated as known—the configuration in which the mechanism is cleanest and the factors are exact. The two-shock economy then carries

the taxonomy where shock separation becomes substantial. This paper interconnects with companion papers, but the division is strict: the general matrix machinery and all proofs are in the companion technical note; the full deployment on observed US data—and the estimation extension it defers—is in (Parra-Alvarez et al., 2026). The present paper establishes the algebraic and simulation results on which both latter rest.

1.1 Related Literature

This paper speaks first to the literature that takes the mismatch of observation calendars as its point of departure. On the reduced-form side, the nowcasting agenda reads mixed-frequency releases through factor and state-space models to estimate the current state of the economy (Giannone et al., 2008; Bànbura et al., 2013; Bok et al., 2018). On the structural side, (Posch et al, 2016) document how a continuous-time equilibrium model can be estimated on daily financial data combined with lower-frequency macroeconomic series, and find that the daily rate strengthens identification of the structural parameters; (Foroni and Marcellino, 2016) show that mixed-frequency methods mitigate the identification problems created by temporal aggregation in structural VARs. Behind both stands the continuous-time tradition of estimation from mixed stock and flow samples (Bergstrom, 1986; Zadrozny, 1988; Harvey and Stock, 1993; Chambers, 2009). The present paper complements this agenda from the filtering and representation side: the objects are the filtered state and the contemporaneous structural shock, delivered online, release by release, on the data's own calendar.

The central motivation of the paper is that the calendar mismatch is not innocuous which is also the subject of the temporal-aggregation literature. In the vocabulary of (Christiano and Eichenbaum, 1987), the resulting specification error "temporal aggregation bias," and (Marcet, 1991) shows that aggregation systematically inflates the leading coefficients of the moving-average representation. The continuous-time tradition localized the remedy in the representation of measurement: a flow must be treated as the integral of an underlying rate, not as a point sample of it (Bergstrom, 1984; Chambers, 2009)—the distinction that prices the naive Euler filter throughout this paper. A recent wave carries the mechanism to structural shocks. (Jacobson et al., 2023) argue that the apparent "Fed information effect" is in part an artefact of temporal aggregation: pairing high-frequency monetary surprises with a daily inflation measure removes the anomalous response, and measurement error in the daily proxies does not overturn the conclusion. (Snudden, 2024) shows that over 70 percent of structural innovations recovered from temporally aggregated data are predictable and mistimed, with the distortion growing in the persistence of the underlying process. These contributions diagnose the mechanism, empirically and by simulation; Section 4 gives it closed form, in four factors and a law that depend on persistence and sampling only through their product. And the representation question of this paper sits one level below the classical one: not integral versus point sample, but two representations of the integral itself—the integral form and the rate form—which differ in exactly one term, the pass-through of the state's Brownian increment into the datum.

The paper's second problem places it in the literature on recovering structural shocks, where the established routes are retrospective. In the SVAR tradition, shocks are read from reduced-form residuals under identifying restrictions imposed after estimation (Ramey, 2016), and are recoverable from the history of the observables only when they are invertible (Fernández-Villaverde et al., 2007). In nowcasting practice, the update at each release is read as news about the current estimate rather than as a structural shock (Bànbura

et al., 2013). Closest to the present framework is (Christensen et al., 2024), a recent advancement of the exact-discretisation programme for estimating continuous-time models with flow data: their route runs from exact discretisation through parameter estimation to ex-post identification of the structural shocks, and it establishes a result this paper relies on throughout—exact discretisation makes parameter estimation invariant to the sampling frequency, an invariance Section 6 converts into a misspecification test. The framework here adds the step their paper does not: real-time filtering of the same model class, with the structural shocks projected out at the same time. That addition is carried by the rate-form representation of flows, which keeps the contemporaneous shock in the datum itself, and it returns the binary invertibility condition of the retrospective route as a graded, computable ceiling on attainable recovery.

A latent continuous-time process, observed discretely and with error, is also the setting in high-frequency financial econometrics. The frictions studied there are the ones of timing. In (Barndorff-Nielsen et al., 2008, 2011), realised kernels absorb microstructure noise and non-synchronous trading; the covariation of asynchronously observed prices is estimated by quasi-likelihood (Shephard and Xiu, 2017) or recast as a missing-data problem as in (Corsi et al., 2015). The finding of (Aït-Sahalia et al., 2010) is the one the flow results of this paper most resemble: when the measurement structure and asynchronity is left unmodelled, ever-finer sampling degrades covariance estimates rather than improving them, and modelling it explicitly restores the value of high-frequency data. Section 4 contains the macroeconomic counterpart of that finding, with one difference in the diagnosis. For flows the damage is done by the representation of the measurement, and it is largest exactly where the releases are most precise.

On the other hand, in this paper the financial series is an observable, and the estimand is the macroeconomic fundamental behind it. That ordering has an equilibrium foundation: asset prices are forward-looking functionals of the consumption process (Lucas, 1978; Breeden, 1979), and a small persistent component of consumption dynamics carries much of their variation (Bansal and Yaron, 2004). Time aggregation confounds the estimation of exactly this component (Bansal and Yaron, 2004). Mixed-frequency data recover it: in (Schorfheide et al., 2018), the evidence for the persistent component strengthens once asset returns enter the estimation. Thereby, daily financial observables are informative about the fundamental shocks this framework projects out, demand and monetary alike, and how informative is what the simulations of Section 5 measure.

2 Continuous-Time State-Space Framework

This section states the representation underlying the analysis, in the compact form the factor derivations of Section 4 consume; the general matrix treatment—arbitrary state dimension, shock count, and stock/flow partition—together with all proofs and implementation criteria, is collected in the companion technical note (Rátfai, 2026). The central object is a distinction usually left implicit: two economically different answers to the question of what a reported flow *is*—the integral-form primitive of the exact-discretisation tradition and the rate-form primitive maintained here—which differ in exactly one term, the *pass-through* of the state's own Brownian increment into the observation increment. The framework filters the latent state and the contemporaneous structural shock *simultaneously*: the same representation that tracks the state places the shock inside the flow innovations, where a fixed projection reads it out in real time

(Section 3).

Lastly, regarding the observation-count geometry,the feedback-rule economy is the canonical stochastically singular configuration ($n_y > n_\varepsilon$), since one shock is pinned down by three observables. Without measurement error the observables would obey exact linear relations, and identification of DSGE parameters in this neighbourhood is treated with corresponding care in the literature (Komunjer and Ng, 2011; Canova and Sala, 2009). Real-time filtering given a calibration is a different exercise. There the same proximity helps, because jointly precise observables pin the one shared increment ever more tightly, and Section 4 prices the gain in closed form. The two-shock economy then leaves the boundary, and precision stops being the binding constraint. What matters instead is whether the observables can tell the two increments apart, a property of the loading geometry, and the simulations answer that question under dense sampling of the policy rate.

2.1 The Model and the Two Primitives

The latent state $x_t \in \mathbb{R}$ is the demeaned demand shock $x_t = d_t - 1$, where d_t follows an Ornstein–Uhlenbeck process—the demand-shock specification of the continuous-time New Keynesian model of (Posch, 2020)—

$$dd_t = -\rho_d(d_t - 1)\,dt + \sigma_d\,dW_t^x, \qquad \rho_d > 0,\ \sigma_d > 0, \tag{1}$$

so that $dx_t = Ax_t\,dt + \sigma_d\,dW_t^x$ with $A = -\rho_d$. Three variables are linked to the state through the structural loadings of the New Keynesian model—output gap, inflation, policy rate—collected in $C = [\,\bar{y}_d,\ \pi_d,\ \phi_\pi\pi_d\,]'$, and the contemporaneous signal vector is the level

$$y_t = C\,x_t, \tag{2}$$

an algebraic equilibrium condition in the sense of (Christensen et al., 2024). The level $y_t = Cx_t$ is the model's *equilibrium intensity*: the policy-function image of the state, which for a cumulative observable plays the role of an instantaneous accrual rate. The econometrician observes functionals of this level recorded by the measurement process; what a reported *flow* is turns out to be the decision on which everything downstream depends.

A reported flow—quarterly GDP, quarterly inflation—is the amount accumulated over a reporting window $(t_{k-1}, t_k]$ of length h, published as a rate. Two primitives answer what that amount is, and the choice is a property of the release, made before any filter is written down.[1] *Primitive I (integral form)*, the flow treatment of the exact-discretisation tradition (Bergstrom, 1986; Harvey and Stock, 1993; Zadrozny, 1988; Christensen et al., 2024), defines the flow as the interval integral of the equilibrium intensity,

$$y_{j,k}^{F,\mathrm{I}} \;=\; \frac{1}{h}\int_{t_{k-1}}^{t_k} C_j\,x_u\,du \;+\; \text{release error}; \tag{3}$$

the intensity is deterministic given the state, no diffusion loading of the measurement

[1]The choice is also, in principle, testable: under Primitive II the reported rate carries excess variance $C_j^2\sigma_d^2/h$ and an innovation covariance structure that Primitive I excludes (display (42)), so reported flow data discriminate between the two. The tests are stated in Section 6 and executed on observed releases in (Parra-Alvarez et al., 2026).

is specified, and the shock reaches the datum only through the path of x. *Primitive II (rate form; this paper)* specifies the flow as the window increment of a *cumulative observable*—cumulated income, cumulated price change—and asserts, as a property of the measurement, that this cumulative is not the time integral of the equilibrium intensity: that integral has finite variation and carries no contemporaneous increment. The primitive assigns instead the observation increment the intensity as its drift and the state's Brownian increment, at the same loading, as its diffusion[2]:

$$dy_{j,t}^{F,\mathrm{II}} \;=\; \underbrace{C_j\, x_t\, dt}_{\text{signal drift}} \;+\; \underbrace{C_j\, \sigma_d\, dW_t^x}_{\text{pass-through}}, \qquad j \in \{y, \pi\}, \tag{4}$$

with the published figure the accumulated increment, reported as a rate and carrying one measurement draw per release,

$$y_{j,k}^{F,\mathrm{II}} \;=\; \frac{1}{h}\int_{t_{k-1}}^{t_k} dy_{j,u}^{F,\mathrm{II}} \;+\; \sqrt{h}\,\eta_j\,\varepsilon_{j,k}, \qquad \varepsilon_{j,k} \sim \mathcal{N}(0,1). \tag{5}$$

Remark 2.1 (The pass-through scale)**.** *The general form of* (4) *carries a free scale: $dy_{j,t}^F = C_j x_t\, dt + \tau\, C_j \sigma_d\, dW_t^x$ with $\tau \geq 0$, where $\tau = 0$ is Primitive I and this paper maintains $\tau = 1$ in the units of the calibration—full pass-through of the contemporaneous increment. The scale enters the excess innovation variance as $\tau^2 C_j^2 \sigma_d^2/h$ and the flow–flow cross-moment as $\tau^2 C_y C_\pi \sigma_d^2/h$, the moments the tests of Section 6.1 measure, so the data identify τ and the tests check for the restriction $\tau = 1$ exactly as they assess between the primitives.*

Primitive II is maintained throughout the paper, so the superscript is dropped hereafter. Carrying out the integration exactly along the Ornstein–Uhlenbeck path, with $\Gamma_h \equiv A^{-1}(e^{Ah} - 1)$ the flow operator and $x_{k-1} \equiv x_{t_{k-1}}$ the state at the window's opening, the published rate decomposes exactly into four named pieces:

$$y_{j,k}^F = \underbrace{\tfrac{C_j}{h}\,\Gamma_h\, x_{k-1}}_{\text{carry-over}} \;+\; \underbrace{\tfrac{C_j}{h}\, u_k}_{\text{path news}} \;+\; \underbrace{\tfrac{C_j}{h}\,\sigma_d\,\Delta W_k}_{\text{pass-through}} \;+\; \underbrace{\sqrt{h}\,\eta_j\,\varepsilon_{j,k}}_{\text{release error}}, \tag{6}$$

where $u_k = \int_0^h \Gamma(h-s)\,\sigma_d\, dW_s^x$ is the within-window drift innovation and ΔW_k the window's structural increment.[3] The carry-over is forecastable from conditions at the window's opening and common to both primitives; the path news is the only stochastic flow content

[2]That the pass-through cannot be derived from the path of the intensity was made precise in a comment by Gerrit Meyerheim; the rate form is a maintained and tested measurement primitive, not a consequence of continuity.

[3]Why the path news is a weighted integral while the pass-through is a plain increment: both are Itô integrals of the same Brownian motion over the same window, $\int_0^h K(s)\, dW_s^x$, and they differ only in the kernel. The pass-through has $K \equiv 1$. It originates in the diffusion term of (4): each shock enters the cumulative observable once, at the instant it occurs, with unit weight, and is not revisited. The path news has $K(s) = \Gamma(h-s)$. It originates in the drift term: a shock at time s moves the state, whose new level then feeds the accrual at every remaining instant $u \in (s, h]$, discounted by mean reversion, summing to $\int_s^h e^{A(u-s)} du = \Gamma(h-s)$—the shock's remaining-window exposure through the level. Since $\Gamma(\tau) \approx \tau$ for small τ, the indirect weight is bounded by the time left in the window (variance $\sigma_d^2 h^3/3$ against the direct channel's $\sigma_d^2 h$), and a shock arriving near the window's close has $\Gamma(h-s) \approx 0$ yet full pass-through weight: late-window shocks are present in the datum through the direct channel alone. Economically, the pass-through records the shock's own transactions as part of the window's total; the path news records the income the repositioned economy generates over the remainder of the window.

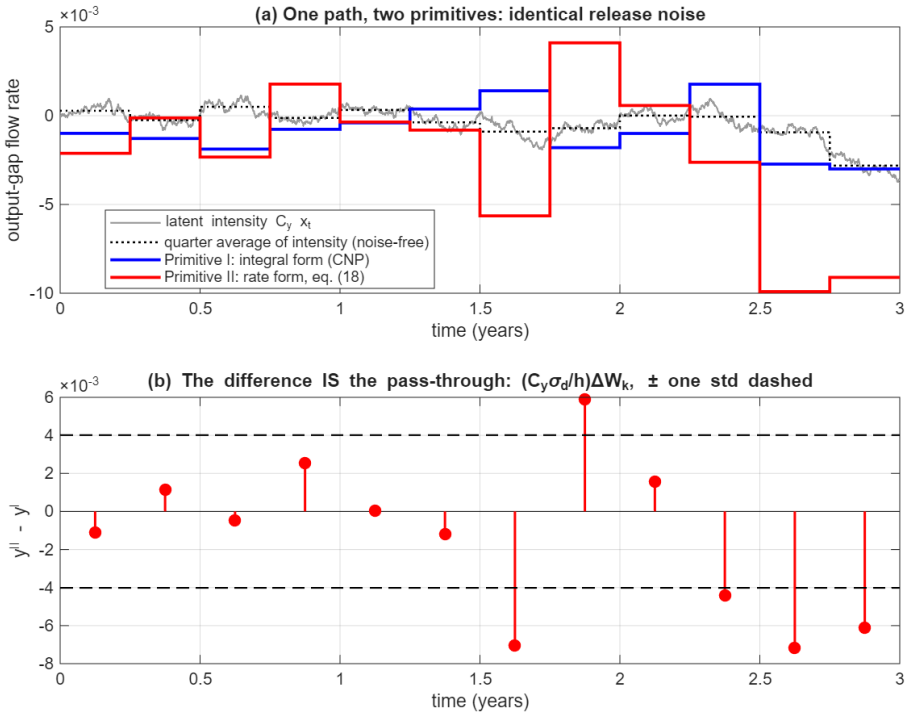


Figure 1: One path, two flow primitives. Panel (a): the latent intensity $C_y x_t$ (grey), its noise-free quarterly average (dotted), and the quarterly reported rate under Primitive I (integral form, blue) and Primitive II (rate form, red), constructed from the same realisation with identical per-release noise draws ($\rho = 0.01$). The anatomy of the visual gaps, in rate units: the gap between Primitive I and the dotted average is pure release noise, of standard deviation $\sqrt{h}\,\eta_f = 0.0015$; the gap between Primitive II and Primitive I is the pass-through, of standard deviation $C_y\sigma_d/\sqrt{h} = 0.0040$—against a signal $C_y x_t$ of stationary standard deviation $C_y\sigma_d/\sqrt{2\rho_d} = 0.0022$ and path news of standard deviation $C_y\sigma_d\sqrt{h/3} = 0.0006$. The pass-through exceeds the path news by the loading-free factor $\sqrt{3}/h \approx 7$: most of a rate-form flow's surprise is the contemporaneous increment itself, which is why the red series does not track the intensity average and is not meant to. Panel (b): the difference of the two primitives, which equals the pass-through $(C_y\sigma_d/h)\Delta W_k$ exactly and is invariant to the release noise; dashed lines mark ± one standard deviation, $C_y\sigma_d/\sqrt{h}$. Under Primitive II the quarter's structural increment is present in the datum itself.

under Primitive I, represented exactly by the Van Loan machinery; the pass-through is present under Primitive II only—the quarter's structural shocks imprinted on the datum as they arrive, the source of the cross-covariance below, and the channel that makes the contemporaneous shock recoverable online. Figure 1 displays the two primitives on a single simulated path with identical release-noise draws: their difference is exactly the pass-through term—under Primitive II, the within-quarter structural increment is *in* the published number.

2.2 Measurement Convention and System Matrices

Measurement error attaches to the *published figure*, not to the accrual: each release—flow rate or stock reading alike—carries one independent draw of standard deviation $\sqrt{h}\,\eta_j$,

$$y_k^S \;=\; C_s\, x_{t_k} \;+\; \sqrt{h}\,\eta_s\,\varepsilon_k^s, \qquad C_s = \phi_\pi \pi_d, \tag{7}$$

with all ε's i.i.d. standard normal, mutually independent, and independent of W^x—the convention native to discrete measurement equations in the exact-discretisation literature (Christensen et al., 2024), treating the error as a property of the release rather than of the accrual, symmetrically across channels. Collecting the flow loadings in $C^f = [\,\bar{y}_d,\ \pi_d\,]'$ and writing $B = \sigma_d$, the continuous block driven by the single scalar Brownian motion is

$$dx_t = A\,x_t\,dt \;+\; B\,dW_t^x, \tag{8}$$

$$dy_t^F = C^f x_t\,dt \;+\; C^f B\,dW_t^x, \qquad C^f B = [\,\bar{y}_d\sigma_d,\ \pi_d\sigma_d\,]', \tag{9}$$

a two-channel observation increment whose diffusion is the rank-one vector $C^f B$: both flows load on the one structural shock. The point-in-time stock has no increment equation—it is a sampled level, not an accumulation. In collected ABCD form—C the observation loading, D the observation *shock*-loading on the same Brownian motion as B, partitioned conformably with C into its flow rows D^f and stock row D^s, and $\Sigma_{v,h}$ the covariance of the stacked per-release measurement errors of (5) and (7)—

$$A = -\rho_d, \quad B = \sigma_d, \quad C = \begin{bmatrix} C^f \\ C_s \end{bmatrix}, \quad D = \begin{bmatrix} D^f \\ D^s \end{bmatrix} = \begin{bmatrix} C^f B \\ 0 \end{bmatrix}, \quad \Sigma_{v,h} = h\,\mathrm{diag}(\eta_y^2, \eta_\pi^2, \eta_s^2). \tag{10}$$

The flow block $D^f = C^f B$ carries the pass-through; the stock row is $D^s = 0$ because a sampled level has no accrual through which to inherit an increment. Primitive I is exactly the restriction $D^f = 0$: the same A, B, C, and measurement variances, with the observation shock-loading removed. For the flow channels the accumulation is generated by the augmented pair

$$d\begin{bmatrix} x_t \\ y_t^F \end{bmatrix} = \underbrace{\begin{bmatrix} A & 0_{1\times 2} \\ C^f & 0_{2\times 2} \end{bmatrix}}_{\mathcal{A}} \begin{bmatrix} x_t \\ y_t^F \end{bmatrix} dt + \underbrace{\begin{bmatrix} B \\ C^f B \end{bmatrix}}_{\mathcal{B}} dW_t^x, \tag{11}$$

whose exact discretisation delivers the discrete representations of Section 3; the integral-form benchmark corresponds to the same $\mathcal{A}$ with the accumulator rows of $\mathcal{B}$—the last two—set to zero. Because the state diffusion and the flow-observation diffusion act on the same Brownian motion, the state innovation and the flow-observation innovation are cross-correlated:

$$\Sigma_{xy} \;\equiv\; \frac{1}{dt}\,\mathrm{cov}\big(dx_t,\ dy_t^F\big) \;=\; B\,(C^f B)' \;=\; \sigma_d^2\,(C^f)', \tag{12}$$

channel by channel $\Sigma_{xy,j} = C_j \sigma_d^2$. The cross-covariance arises entirely from the pass-through; the release errors contribute nothing to it. It is what renders the contemporaneous state innovation partially observable, and it operates on the *gain side*: observability of (A, C)—which holds here through any non-zero loading and settles that the state *level* is identified—is blind to Σ_{xy}, whereas the optimal filter's gain acquires the additive

term $\Sigma_{xy}R^{-1}$—R denoting the variance rate of the total flow-observation disturbance, pass-through plus release noise, defined in (16)—and so reads information about the current shock directly from the flow data (Section 3). A representation that sets $\Sigma_{xy} = 0$—Primitive I—discards that channel by construction. The general observability, fundamentalness, and well-posedness criteria are collected in the companion note.

2.3 Stocks versus Flows: Two Shock Channels

The observables divide by how the data record the signal over the reporting interval: a stock is a point-in-time reading of the level at t_k (here the policy rate); a flow is the accumulation of the increments (4). The asymmetry is finer than presence versus absence. A stock carries no separate within-interval shock term, but its sampled level contains the state innovation $\eta_k^x = \int_0^h e^{A(h-s)}\sigma_d\, dW_s^x$, with variance $\Sigma_{\eta^x} = \sigma_d^2\,(e^{2Ah} - 1)/(2A)$, and for a scalar Ornstein–Uhlenbeck state the innovation is nearly the increment itself:

$$\operatorname{corr}(\eta_k^x,\ \Delta W_k) \;=\; \frac{\sigma_d \Gamma_h}{\sqrt{\Sigma_{\eta^x}\, h}} \;=\; 0.9995 \quad (h = \tfrac{1}{4},\ \text{calibration of Section 5.2}). \tag{13}$$

The stock thus conveys the contemporaneous shock *through the level*—conditionally, insofar as the level is precisely observed—whereas the flow carries it *unconditionally*, in the datum itself. These are the two complementary channels of the online recovery of Section 3: the flow channel dominates when flows are the precise observables, the level channel when the stock is, and the operator that uses both spans the entire precision range.

3 The Three Filters in Brief

Three filters are compared throughout: all address the same online problem—tracking the latent state as mixed stock-flow observations arrive—and they differ only in their *measurement model* for the observed flow, that is, in how each represents the pass-through increment $C_j\sigma_d\, dW^x$ of (4) over the sampling interval. This section states each filter's measurement model and the objects the factor derivations of Section 4 consume—the innovation structures and the recovery operators. Full recursions, the two-step sampling algorithm, existence conditions, proofs, and implementation criteria are collected in the companion technical note; the exact discrete filter of the rate-form process itself, EDT$^+$, serves as a verification instrument and is constructed in Appendix C.

3.1 Online Real-Time Tracking: the Qualifying Criterion

Nowcasting is the estimation of the current state of a target variable from information available up to the present instant, updated as each new observation arrives (Bànbura et al., 2013). The comparisons of this paper are run against the stronger standard of *online real-time tracking*, defined formally in the companion note (Rátfai, 2026): a filter qualifies if it delivers, contemporaneously and using only data observed through t_k, estimates of both

$$\hat{x}_{t_k} \quad \text{(the latent state)} \qquad \text{and} \qquad \hat{\varepsilon}_{t_k} \quad \text{(the contemporaneous structural innovation)}, \tag{14}$$

the second by a fixed linear map $\hat{\varepsilon}_{t_k} = \mathcal{M}\,\nu_{t_k}$ on the stacked contemporaneous innovation $\nu_{t_k} = [\nu^f_{t_k};\ \nu^s_{t_k}]$—a constant operator computed once from the steady innovation moments (given explicitly in (24)), using no future data and no backward pass, where nowcasting practice reads the update at each release as news about the current estimate rather than as a structural shock (Bànbura et al., 2013). The recovery operators are derived once (Section 3.3) and applied identically to every filter's own innovations, so the recovery comparison isolates the representations exactly as the tracking comparison does.

3.2 Three Measurement Models

Method 1: continuous-time correlated-noise Kalman–Bucy filter (CT). The rate-form filter takes (8)–(9) at face value: the flow disturbance shares the structural increment with the state, so the gain carries the cross-covariance term in addition to the usual one,

$$d\hat{x}_t = A\hat{x}_t\,dt + K_t\,(d\tilde{y}_t - C\hat{x}_t\,dt)\,, \qquad K_t = (P_tC' + \Sigma_{xy})\,R^{-1}, \tag{15}$$

Here $\tilde{y}_t$ stacks the cumulative observation processes, K_t is the filter gain, P_t is the filtering error variance $\mathbb{E}(x_t - \hat{x}_t)^2$, and R is the variance rate[4] of the total observation-side disturbance the filter consumes—pass-through plus release noise, whose flow block R^f is displayed below. The error variance solves the correlated-noise Riccati equation under the reduced (output-injection) form $\bar{A} = A - \Sigma_{xy}R^{-1}C$, $\bar{Q} = \Omega - \Sigma_{xy}R^{-1}\Sigma'_{xy}$, with $\Omega = BB' = \sigma_d^2$ the process-noise variance rate (Kalman-Bucy, 1962; Anderson and Moore, 1979). At discrete sampling the filter runs in two steps per interval: the flow channels are assimilated continuously across the window through the flow-only gain $K^f_\infty = (P^f_\infty(C^f)' + \Sigma^f_{xy})(R^f)^{-1}$—$P^f_\infty$ and K^f_∞ the steady-state error variance and gain of the flow-only subproblem, Σ^f_{xy} the flow rows of Σ_{xy}—and the point-in-time stock is incorporated by an exact discrete update at its observation instant—no state augmentation, no accumulators, and no rate-smearing of the stock (the two-step algorithm is stated in full in the companion note). The total flow-disturbance variance rate is the rank-one pass-through outer product plus the release noise at its equivalent continuous variance rate,

$$R^f = (C^f\sigma_d)(C^f\sigma_d)' + \mathrm{diag}\big((h\eta_f)^2\big), \tag{16}$$

whose off-diagonal element $C_yC_\pi\sigma_d^2$ is as structural as the diagonal: both flows carry the same shock. Because the stock enters as an exact discrete update, the recursion reduces, as the stock becomes the dominant precise channel, to the exact discrete Kalman update of a directly observed state—the mechanism of the Anderson–Moore convergence documented in Section 5.3.

Method 2: exact discrete filter of the integral form (EDT). The exact discretisation of Primitive I (Bergstrom, 1986; Harvey and Stock, 1993; Zadrozny, 1988; Christensen et al., 2024) augments the state with one accumulator per flow, built from the lagged state

[4]Throughout, *intensity* is reserved for first-moment rates—the equilibrium intensity $y_t = Cx_t$ at which a cumulative observable accrues. Second-moment rates—variance per unit time of a white-noise disturbance, its spectral density—are called *variance rates*; the filtering literature often calls these intensities as well, and Ω, R, $\bar{Q}$ below are variance rates in that sense.

through Γ_h in rate units,

$$\mathbf{z}_k = \mathbf{A}_z \mathbf{z}_{k-1} + \mathbf{u}_k, \qquad \mathbf{A}_z = \begin{bmatrix} e^{Ah} & 0 & 0 \\ C_y\,\Gamma_h/h & 0 & 0 \\ C_\pi\,\Gamma_h/h & 0 & 0 \end{bmatrix}, \qquad \mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \tag{17}$$

$$\mathbf{Q} = \begin{bmatrix} \Sigma_{\eta^x} & \Sigma_{\eta^x\eta^f} C_y & \Sigma_{\eta^x\eta^f} C_\pi \\ C_y \Sigma_{\eta^x\eta^f} & C_y \Sigma_{\eta^f} C_y & C_y \Sigma_{\eta^f} C_\pi \\ C_\pi \Sigma_{\eta^x\eta^f} & C_\pi \Sigma_{\eta^f} C_y & C_\pi \Sigma_{\eta^f} C_\pi \end{bmatrix}, \tag{18}$$

$$\mathbf{Y}_k = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ C_s & 0 & 0 \end{bmatrix} \mathbf{z}_k + \mathbf{v}_k, \tag{19}$$

Here $\mathbf{z}_k$ stacks the state with the two flow accumulators in rate units; $\mathbf{Q}$ is the covariance of the augmented innovation $\mathbf{u}_k$—the state innovation and the accumulators' path news, the only stochastic content the integral form carries from the state—assembled from the Van Loan blocks ($\Sigma_{\eta^x}, \Sigma_{\eta^f}, \Sigma_{\eta^x\eta^f}$), the state-innovation variance, path-news variance, and their cross-covariance in rate units (Van Loan, 1978); $\mathbf{Y}_k$ stacks the two reported flow rates and the stock reading; and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathrm{diag}(h\eta_y^2, h\eta_\pi^2, h\eta_s^2))$ collects the per-release measurement errors of Section 2.2—identical to the other filters', so all comparisons isolate the representation of the flow, not the noise assumptions. Accumulators are reset after each observation. Within its own measurement model (3) this filter is exact and complete—it is included as the correct filter of the standard model, not as a strawman. Its signature relative to rate-form data: the blocks integrate the process noise only, so the pass-through is absent from $\mathbf{Q}$, and the retained state-flow cross-covariance in rate units is the fraction $\Phi(h)\,\sigma_d^2$ of the instantaneous Σ_{xy} per unit loading—vanishing as $h \to 0$, so fine sampling drives this filter to the *uncorrelated*-noise filter, not to Method 1.

Method 3: naive Euler filter. The common practitioner shortcut replaces the exact discretisation with first-order approximations, $A_{\text{Euler}} = I + Ah$ and $\Sigma_{\eta,\text{Euler}} = \Omega h$, and—most consequentially—ignores the stock-flow distinction, treating every observation as a point-in-time reading of the level. Its three errors are an $O(h^2)$ transition error, a gain-miscalibrating covariance error, and a flow misspecification that omits the temporal-aggregation operator entirely; the first two shrink as the base frequency rises, the third does not, because it is structural misspecification of the observation equation rather than a discretisation error.

Table 1: Comparison of Filtering Methods

Property	**CT (rate form)**	**EDT (integral form)**	**Euler**
State dimension	n	$n + n_{\text{flow}}$	n
Transition	Exact (e^{Ah})	Exact (e^{Ah})	Approx. ($I + Ah$)
Innovation covariance	Van Loan	Van Loan (Ω only)	Ωh
Flow handling	Continuous ODE	Rate accumulator	Ignored
Stock handling	Discrete jump	Augmented vector	Point-in-time
Within-interval path	Dynamic (ODE)	Boundary-anchored	None
Pass-through of (4)	In gain, instantaneous	Omitted from $\mathbf{Q}$	Omitted
Cross-cov. retained	Σ_{xy} (full)	$\Phi(h)\,\sigma_d^2$ (vanishing fraction)	Absent

The three methods occupy three positions on one axis—how much of the pass-through

each retains, and at what temporal resolution: instantaneously in the gain (CT), as the vanishing fraction $\Phi(h)$ through the integrated state (EDT), or not at all (Euler), whose endpoint treatment of the flow does, however, accidentally reintroduce a mis-scaled contemporaneous channel—the *endpoint surrogate*, as it is referred to below—which is why Euler can outperform EDT in the flow-precise corner of Section 5.3 while remaining misspecified everywhere. The CT filter attains the minimum-MSE ceiling of the release information set, as validated against EDT$^+$ (Appendix C), whose own deficit is the negligible resolution effect $\Phi(h)$.

3.3 Online Structural-Shock Recovery: Flow-Only and Stacked Operators

The second object of (14) is the window increment ΔW_k. Because it postdates all data through t_{k-1}, its minimum-variance estimate conditions on the contemporaneous innovation alone,

$$\widehat{\Delta W}_k \;=\; \mathbb{E}\big[\Delta W_k \mid \mathcal{Y}_k\big] \;=\; c'\, S_\nu^{-1}\, \nu_k, \qquad \hat{\varepsilon}_{t_k} = \widehat{\Delta W}_k/\sqrt{h}. \tag{20}$$

where $\mathcal{Y}_k$ is the information set through t_k, $c \equiv \mathrm{cov}(\nu_k, \Delta W_k)$ is the covariance of the stacked innovation with the target increment, and $S_\nu \equiv \mathrm{var}(\nu_k)$ is the innovation covariance matrix—both constant at steady state and given in closed form below. At steady state, with posterior error variance P and window-start estimation error $\tilde{e}_{k-1} = x_{k-1} - \hat{x}_{k-1}$, the innovations of the exact representation are

$$\nu^f_{j,k} = \frac{C_j}{h}\big(\Gamma_h\, \tilde{e}_{k-1} + u_k + \sigma_d \Delta W_k\big) + \sqrt{h}\,\eta_f\,\varepsilon_{j,k}, \quad \nu^s_k = C_s\big(e^{Ah}\,\tilde{e}_{k-1} + \eta^x_k\big) + \sqrt{h}\,\eta_s\,\varepsilon^s_k, \tag{21}$$

with the target covariance c and the blocks of $S_\nu = \left[\begin{smallmatrix} S_{ff} & S_{fs} \\ S'_{fs} & S_{ss} \end{smallmatrix}\right]$ in closed form in the window moments,

$$c_j = \frac{C_j \sigma_d}{h}\Big[\frac{\Gamma_h - h}{A} + h\Big], \qquad c_s = C_s\,\sigma_d\,\Gamma_h \;\neq\; 0, \tag{22}$$

$$\begin{aligned} S_{ff}[j,l] &= \frac{C_j C_l}{h^2}\big(\Gamma_h^2 P + v^+\big) + \delta_{jl}\, h\eta_f^2, \\ S_{fs}[j] &= \frac{C_j C_s}{h}\big(\Gamma_h e^{Ah} P + x^+\big), \qquad S_{ss} = C_s^2\big(e^{2Ah} P + \Sigma_{\eta^x}\big) + h\eta_s^2, \end{aligned} \tag{23}$$

where $v^+ = \Sigma_{\eta^f} + \sigma_d^2 h + 2\sigma_d^2 A^{-1}(\Gamma_h - h)$ and $x^+ = \Sigma_{\eta^x \eta^f} + \sigma_d^2 \Gamma_h$ are the exact second moments of the rate-form flow disturbance. The flow entries of c carry the pass-through plus the path news; the stock entry is non-zero by the level collinearity (13). Two online estimators follow, one per information set:

$$\mathcal{M}^f \nu^f_k \;\equiv\; \frac{(c^f)'\, S_{ff}^{-1}\, \nu^f_k}{\sqrt{h}} \quad \text{(flow-only)}, \qquad\qquad \mathcal{M}\,\nu_k \;\equiv\; \frac{c'\, S_\nu^{-1}\, \nu_k}{\sqrt{h}} \quad \text{(stacked)}, \tag{24}$$

each a constant operator on the contemporaneous innovation, the stacked map being the operator $\mathcal{M}$ of the online criterion (14). Both are the optimal linear reads of the increment from their information sets—the maximisers of the optimal-read lemma of

(Rátfai, 2026)—and their attained correlations, evaluated at the exact filter's steady P, are the ceilings of Appendix B (Proposition 4.5 states the correlation for general fixed maps)—the flow-only operator generalising the classical correlated-noise recovery of the latent Wiener increment from observed signal increments (Nüsken et al., 2019), the stacked operator adding the level channel of Section 2.3. Every filter applies the *same* operators to its *own* innovations. This contrasts with (Christensen et al., 2024), who back out structural shocks *ex post* from reduced-form residuals through a generalised inverse with $O_P(1)$ approximation error (their Proposition 4). The operators (24) add the online counterpart to the ex-post recovery of their Proposition 4: exact conditional expectations for their information sets, evaluated as each release arrives, with residuals equal to the information limit of the window—the attainable ceilings stated in Appendix B.

3.4 Proximity to Stochastic Singularity

A single structural shock observed through multiple channels is the canonical stochastically singular configuration ($n_y > n_\varepsilon$): without measurement error the observables would obey exact linear relations (Komunjer and Ng, 2011; Canova and Sala, 2009)—the configuration (Christensen et al., 2024) exclude by assumption. Here well-posedness is intrinsic to the filter: R^f in (16) is invertible if and only if the release noise is strictly positive, and the effective process-noise variance rate of the correlated-noise reduction has the closed form

$$\bar{Q}^f = \frac{(h\eta_y)^2(h\eta_\pi)^2\,\sigma_d^2}{\sigma_d^2\big(C_\pi^2(h\eta_y)^2 + C_y^2(h\eta_\pi)^2\big) + (h\eta_y)^2(h\eta_\pi)^2} \;>\; 0 \quad \text{for all } \eta_y, \eta_\pi > 0, \qquad \bar{Q}^f \xrightarrow{\eta_f \to 0} 0, \tag{25}$$

so the singular limit is approached, never crossed: strictly positive release noise is a defining condition of the filter, not a regularisation imposed on it. (The general result—the reduction is positive semidefinite in any dimension, with (25) the scalar case of a Woodbury representation—is proved in the companion note.) The proximity is detrimental for likelihood-based parameter estimation (Komunjer and Ng, 2011; Canova and Sala, 2009) but favourable for the complementary exercise conducted here, filtering and recovery *given* the parameters: as the flows become jointly precise they pin the shared increment ever more tightly, the attainable recovery correlation rises monotonically toward its noise-free limit, and the operator attains it at every point of the grid (Section 5.3). With several shocks the channels must *separate* increments rather than reinforce one, a question of the loading geometry (Appendix A) taken up by the partial-adjustment model of Section 5.4.

4 Temporal Aggregation and Filter Selection: Four Factors and a Recovery Law

The treatment of flow variables is where the three filters most sharply diverge, and it places this paper's contribution to the temporal-aggregation literature. That continuous-time formulation followed by exact discretisation avoids the aggregation bias incurred by treating flows as point sampled is a well-established principle (Bergstrom, 1984; Chambers, 2009); the relative accuracy of exact discretisation over a naive Euler–Maruyama scheme has been demonstrated by Monte Carlo in the DSGE setting by (Christensen et al., 2024), and the general question of whether the continuous-time projection can be approximated by sampling more finely was posed by (Marcet, 1991). What the literature has established

qualitatively, or by simulation, this section makes precise. I derive four closed-form scalars, each attached to one mechanism of the filter comparison, and one law that converts any filter's tracking accuracy into its shock-recovery accuracy. The *temporal-aggregation factor* $\Phi(h)$ and the *pass-through penalty* $\Pi(\rho)$ decompose the gap between the continuous-time rate-form filter (CT, Method 1) and the integral-form exact filter (EDT, Method 2) into a removable resolution component—the gap between CT and the rate-form exact discrete filter EDT$^+$ of Appendix C, sized by Φ—and an irreducible representation component—the omitted pass-through, priced in the fine-sampling limit by Π; the *structural-penalty pair* $\Psi_{\mathrm{I}}(h), \Psi_{\mathrm{II}}(h)$ prices the Euler shortcut under each of the two measurement primitives of Section 2.1; the *level-resolution factor* $\Lambda(h)$ is the stock arm's counterpart; and the *recovery-correlation law* maps tracking error into recovery accuracy through a single strictly monotone scalar function. All derivations are collected in Appendix D; this section states each object, its origin, and its meaning. Throughout, h is expressed in the annual units of the calibration (Section 5.2), the units in which the rate-form primitive (4) is specified.

4.1 The Temporal-Aggregation Factor: the Removable Component

This subsection prices temporal *resolution*, and one ratio serves two roles. The continuous-time rate-form filter (CT, Method 1) uses the shared-shock cross-covariance at instantaneous resolution—the object of Section 2.2,

$$\Sigma_{xy,j} = C_j\, \sigma_d^2, \tag{26}$$

the contemporaneous covariance per unit time between the state innovation and the flow disturbance. A discrete filter, by contrast, can carry cross-covariance only as a window moment, and the two discrete filters carry different ones. The integral-form EDT (Method 2) carries the drift-channel block $\Sigma_{\eta^x\eta^f}$ of (18)—the instantaneous object propagated to the window endpoint and integrated across the window—and nothing else; the rate-form exact discrete filter EDT$^+$ (Appendix C) carries, in addition, the boundary-anchored pass-through, its cross block being $\Sigma_{\eta^x\eta^f} + \sigma_d^2\,\Gamma_h$. The ratio of the drift-channel block to the accumulated instantaneous object therefore prices two comparisons at once: it is the *entire* cross-covariance content of the integral-form EDT, and it is the order of the resolution deficit that separates EDT$^+$ from CT. Temporal aggregation, in the sense of (Bergstrom, 1984; Weiss, 1984; Marcellino, 1999), is precisely this replacement of instantaneous moments by moments of window functionals—endpoint samples for stocks, window averages for flows—and Φ is that replacement applied to the cross-covariance Σ_{xy}: it depends on the aggregation scheme (A, h) alone, on no loading, no noise parameter, and no filter.

Proposition 4.1 (Temporal-aggregation factor)**.** *The state-flow block of the integral-form exact discrete representation (EDT) is*

$$\Sigma_{\eta^x\eta^f} \;=\; \sigma_d^2 \int_0^h e^{A(h-u)}\,\Gamma(h-u)\,du, \qquad \Gamma(\tau) = A^{-1}\big(e^{A\tau} - I\big)\,, \tag{27}$$

reproducing the Van Loan block to machine precision. The temporal-aggregation factor*, the ratio of the integrated cross-covariance to the instantaneous cross-covariance accumulated*

over the interval, has the closed form

$$\Phi(h) \;\equiv\; \frac{\Sigma_{\eta^x\eta^f}}{\sigma_d^2\,h} \;=\; \frac{1}{Ah}\left[\frac{e^{2Ah}-1}{2A} - \Gamma_h\right] \;=\; \tfrac{1}{2}\,h + \tfrac{1}{2}\,Ah^2 + O(h^3), \tag{28}$$

is increasing in h on the calibrated range, and satisfies

$$\Phi(h) \;\longrightarrow\; 0 \qquad \text{as } h \to 0. \tag{29}$$

Derivation: Appendix D. The consequence of the limit differs by representation. For the rate-form exact filter (EDT$^+$, Appendix C), which carries the pass-through, the retained cross-covariance is $\sigma_d^2\Gamma_h/h = \sigma_d^2(1+O(h))$; its deficit relative to the continuous-time filter is of order $\Phi(h)$, and the representation collapses onto the continuous-time filter as $h \to 0$. For the integral-form EDT, $\Phi(h)$ *is* the entire retained fraction of Σ_{xy}: as $h \to 0$ it retains none of it, converges to the uncorrelated-noise filter, and its gap against the continuous-time filter persists at a fixed, precision-dependent value—the object of Proposition 4.2 below. Conversely, $\Phi(h)$ rises as sampling coarsens (from 0.048 at $h = 0.1$ to 0.113 at the quarterly $h = 0.25$ to 0.333 at $h = 1$): the coarser the interval, the more the discrete representation's cross-covariance is smeared across it and the further it departs from the contemporaneous quantity the continuous-time filter exploits. Crucially, $\Phi(h)$ is a *resolution* effect—the residual a practitioner removes by sampling more finely. Fine sampling removes how the cross-covariance is smeared; it does not restore an omitted pass-through.

4.2 The Pass-Through Penalty: the Irreducible Component

The factor $\Phi(h)$ measures the removable part of the separation between the continuous-time filter and the integral-form EDT; the irreducible part has its own closed form. Consider the flow block in isolation, with the observation process held fixed—measurement variance rates at their quarterly-calibrated values—and only the assimilation interval h refined.

Proposition 4.2 (Pass-through penalty)**.** *Let the flow block* (8)–(9) *be observed at fixed variance rate $R_c = rI_2$, $r = (\bar{h}\,\eta_f)^2$, $\bar{h} = \frac{1}{4}$, and write $R_{\text{tot}} = \sigma_d^2\,C^f(C^f)' + R_c$ for the total flow-disturbance variance rate. (i) The steady-state error variance of the exact (correlated-noise) filter is*

$$\begin{aligned} P_* &= \frac{\bar{A}_* + \sqrt{\bar{A}_*^2 + a_*\bar{Q}_*}}{a_*}, \qquad \bar{A}_* = A - \Sigma_{xy}' R_{\text{tot}}^{-1} C^f, \\ \bar{Q}_* &= \sigma_d^2 - \Sigma_{xy}' R_{\text{tot}}^{-1}\Sigma_{xy}, \qquad a_* = (C^f)' R_{\text{tot}}^{-1} C^f. \end{aligned} \tag{30}$$

(ii) As $h \to 0$ the integral-form filter converges to the uncorrelated-noise Kalman–Bucy filter with gain $K_0 = P_0\,(C^f)' R_c^{-1}$, where $P_0 = \left(A + \sqrt{A^2 + a_0\sigma_d^2}\right)/a_0$ and $a_0 = (C^f)' R_c^{-1} C^f$—its modelled cross-covariance is the vanishing fraction $\Phi(h)\,\Sigma_{xy}$—and its steady-state error under the true rate-form dynamics solves a Lyapunov equation with solution

$$P_{\text{mis}} = \frac{\sigma_d^2\,(1 - a_0P_0)^2 + a_0P_0^2}{2\,(a_0P_0 - A)}. \tag{31}$$

(iii) The pass-through penalty

$$\Pi(\rho) \;\equiv\; \lim_{h\to 0} \frac{\mathrm{MSE}_{\mathrm{EDT}}}{\mathrm{MSE}_{\mathrm{CT}}} \;=\; \frac{P_{\mathrm{mis}}}{P_*} \tag{32}$$

is free of h and satisfies $\Pi(\rho) \geq 1$, with equality if and only if

$$(C^f)'\, R_c^{-1}\, C^f \;=\; \frac{2\rho_d}{\sigma_d^2}, \tag{33}$$

the condition under which the mismatched gain coincides with the optimal gain. Under the calibration, (33) *has the unique solution $\rho^* = \sigma_d^2 \,\|C^f\|^2 / \big(2\rho_d\,(\bar{h}\,\eta_{base})^2\big) = 0.411$.*

Derivation: Appendix D, where the fine-sampling convergence of the actual filters to (32) is also verified numerically. The proposition is the algebraic form of the section's headline asymmetry: for the *same* filter comparison, the resolution component vanishes ($\Phi \to 0$) while the pass-through component persists (Π fixed). Table 2 evaluates Π across the precision grid. Three readings of the numbers. First, Π explodes precisely where the paper's simulations place the continuous-time advantage—the flow-precise corner—and approaches one where the flow noise drowns the pass-through. Second, Π touches one *exactly* at ρ^*. The two gains are collinear for *any* positive-definite release-noise covariance, not only the symmetric design used here, so the filter is sensitive to a single scalar effective gain throughout, and at ρ^* the integral-form filter's two errors—an understated flow-disturbance variance and an omitted cross term—cancel exactly rather than approximately. At that one precision ratio the misspecified filter is optimal; on either side of it the penalty reopens. Third, the convergence in (32) is one-sided in an instructive way: at finite h the integral-form filter still retains the fraction $\Phi(h)$ of the cross-covariance, so in the flow-precise corner its finite-h MSE lies *below* its own fine-sampling limit, and refining the interval degrades it toward P_{mis}—finer sampling strictly worsens the integral-form filter exactly where the pass-through matters most. The proposition concerns the flow block in isolation; in the full design the stock channel is common to both filters and compresses the finite-h ratios toward one as ρ grows, which is the pattern the simulations of Section 5.3 display.

Table 2: Pass-through penalty $\Pi(\rho)$ across the precision grid (flow block, variance rates at quarterly-calibrated values). The penalty equals one exactly at $\rho^* = 0.411$.

ρ	0.01	0.1	0.5	1	2	10	100
$\Pi(\rho)$	176.9	2.526	1.004	1.037	1.051	1.026	1.003

4.3 The Structural-Penalty Pair: Continuous Time versus Naive Euler

The naive Euler filter does not merely smear the cross-covariance across the interval—it discards the flow integral entirely, treating each flow as a point-in-time reading of the state. The cost of that shortcut must be priced against an object, and Section 2.1 established that there are two candidate objects, one per measurement primitive. Under Primitive I

the shock-driven flow-innovation variance over the interval is, in rate units,

$$\Sigma^f_{\mathrm{I}}(h) = \frac{C_y^2\sigma_d^2}{h^2}\int_0^h \Gamma(h-u)^2\,du = \tfrac{1}{3}\,C_y^2\sigma_d^2\,h + O(h^2), \tag{34}$$

where the leading coefficient $\frac{1}{3}$ is the signature of flow aggregation: a flow is the *integral* of the state over the interval. Under Primitive II the same object carries in addition the pass-through and its covariance with the path news, $\Sigma^f_{\mathrm{II}}(h) = C_y^2\,v^+/h^2 = C_y^2\sigma_d^2\,h^{-1}\big(1+O(h)\big)$, with v^+ the exact second moment of (23). The Euler filter assigns the flow the bare state-innovation variance,

$$\Sigma^f_{\mathrm{Euler}}(h) = C_y^2\sigma_d^2\,h, \tag{35}$$

with no within-interval structure of either kind.

Proposition 4.3 (Structural-penalty pair)**.** *Define the structural penalty under each primitive as the fractional error Euler commits in the shock-driven flow-innovation variance,* $\Psi_{\mathrm{p}}(h) \equiv \big(\Sigma^f_{\mathrm{Euler}} - \Sigma^f_{\mathrm{p}}\big)/\Sigma^f_{\mathrm{p}}$ *for* $\mathrm{p} \in \{\mathrm{I},\mathrm{II}\}$*. Then*

$$\Psi_{\mathrm{I}}(h) = 2 - \tfrac{9}{4}\,Ah + \tfrac{51}{80}\,(Ah)^2 + O\big((Ah)^3\big), \qquad \Psi_{\mathrm{I}}(h) \;\longrightarrow\; 2 \quad \textit{as } h \to 0, \tag{36}$$

$$\Psi_{\mathrm{II}}(h) = -1 + h^2 - h^3 + O(h^4), \qquad \Psi_{\mathrm{II}}(h) \;\longrightarrow\; -1 \quad \textit{as } h \to 0, \tag{37}$$

with persistence entering Ψ_{II} *only at fourth order.*

Derivation: Appendix D. Neither penalty vanishes under fine sampling, in sharp contrast to $\Phi(h) \to 0$—but the two limits are not the same number, nor even the same sign, and that is the content of the pair. Against the integral-form object, Euler *overstates* the flow-innovation variance by a factor of three in the limit (h versus $h/3$): this is the tradition's number, the cost of misreading an integral as a point reading. Against the rate-form object—the paper's data-generating process—the sign flips: the pass-through contributes variance of order σ_d^2/h in rate units, which dwarfs Euler's assigned $\sigma_d^2 h$, so Euler *understates* the true flow-disturbance variance almost totally. At the quarterly interval the Euler filter models 4.96% of the rate-form flow-disturbance variance; as $h \to 0$ the modelled fraction, $h^2(1+O(h))$, vanishes. The structural penalty is therefore primitive-dependent, and not even directionally primitive-free: the same shortcut is a threefold overstatement of one object and a near-total omission of the other. What is primitive-free is the conclusion: one cannot Euler-discretise one's way out of the flow problem by sampling more often. This is the precise, closed-form version of the practitioner-facing warning implicit in the Monte Carlo comparisons of (Christensen et al., 2024): the cost of the naive shortcut for flow variables in a continuous-time economic model is structural, not a matter of resolution. Across the sampling grid the three h-dependent factors read as in Table 3; at intervals beyond a year the two Ψ_{II} error sources trade places and the penalty changes sign, crossing zero at $h^* = 1.72$ (Appendix D)—a regime far coarser than any calibration of interest, recorded for completeness.

4.4 Dependence on Persistence

The factors Φ and Ψ_{I} depend on the structural dynamics and the sampling interval only through their product, Ah—the continuous-time mean-reversion rate times the interval,

Table 3: Temporal-aggregation factors across the sampling interval h (annual units, calibration of Section 5.2). Ψ_{II} changes sign at $h^* = 1.72$.

h	0.05	0.1	0.25	0.5	1.0	2.0
$\Phi(h)$ (resolution, CT vs. EDT)	0.025	0.048	0.113	0.203	0.333	0.457
$\Psi_{\mathrm{I}}(h)$ (Euler penalty, integral form)	2.048	2.096	2.244	2.503	3.066	4.386
$\Psi_{\mathrm{II}}(h)$ (Euler penalty, rate form)	−0.998	−0.991	−0.950	−0.837	−0.528	0.219

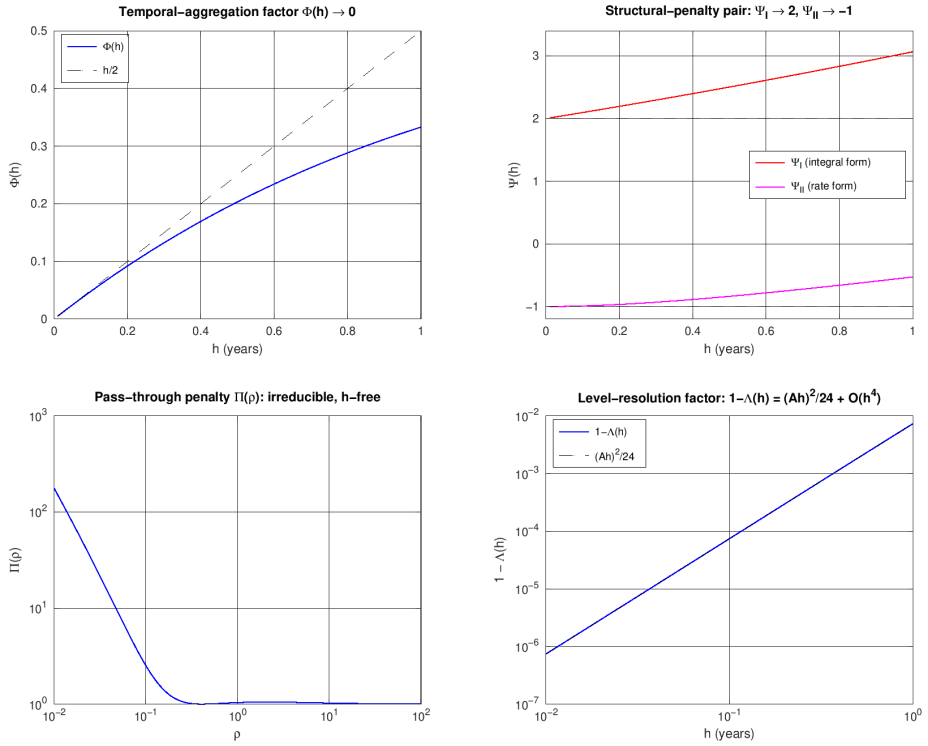


Figure 2: The four closed-form factors under the calibration of Table 4. Top left: the temporal-aggregation factor $\Phi(h)$ with its leading term $h/2$. Top right: the structural-penalty pair—Ψ_{I} (integral form, limit 2) and Ψ_{II} (rate form, limit -1). Bottom left: the pass-through penalty $\Pi(\rho)$, the h-free limit of the integral-form filter's excess MSE on the flow block, equal to one exactly at $\rho^* = 0.411$. Bottom right: the level-resolution factor, $1 - \Lambda(h)$ against its leading term $(Ah)^2/24$ (log–log).

the natural persistence-scaled variable. Expanding the exact expressions,

$$\Phi(h) = \tfrac{1}{2}h + \tfrac{1}{2}Ah^2 + O(h^3), \qquad \Psi_{\mathrm{I}}(h) = 2 - \tfrac{9}{4}\,Ah + \tfrac{51}{80}\,(Ah)^2 + O\big((Ah)^3\big), \tag{38}$$

so persistence enters Ψ_{I} at first order. Since $A = -\rho_d < 0$, a more persistent process (slower mean reversion, $A \to 0^-$) drives Ψ_{I} toward its floor of two, while faster mean reversion inflates it: at the quarterly interval, Ψ_{I} rises from 2.03 when $\rho_d = 0.05$ to 4.95 when $\rho_d = 4$. The mechanism is geometric: when the state mean-reverts slowly, its within-interval path is nearly flat, so the integral that defines the flow and the endpoint that Euler substitutes for it scarcely differ; when mean reversion is fast, the path curves sharply within the interval and the integral-versus-endpoint discrepancy widens. The

rate-form penalty behaves differently: by Proposition 4.3, persistence enters Ψ_{II} only at fourth order. Its size is set not by the within-interval curvature of the path but by the scale mismatch between the pass-through, of order σ_d^2/h in rate units, and Euler's assigned $\sigma_d^2 h$—a pure sampling-dimension effect that no persistence value moderates.

This persistence dependence is the analytical counterpart of an empirical regularity documented across the recent temporal-aggregation literature: (Snudden, 2024) finds that the predictability and mistiming of structural innovations recovered from aggregated data grow with the persistence of the underlying process, and the long-run-risk literature reports that aggregation obscures precisely the persistent component of consumption growth (**?**). These contributions measure the distortion as the predictability of an ex-post reduced-form residual; the factors here measure it as the variance of the within-interval flow innovation, and both depend on persistence through the same product Ah. The two are complementary diagnostics of one within-interval mechanism rather than competing measures, and the closed forms in (38) make the shared persistence channel explicit.

The opposite monotonicities are then no contradiction, for three reasons. First, the floors are not zero: rising persistence strips away only the mean-reversion surcharge—the within-window curvature that separates the integral from the endpoint—and leaves the persistence-independent core exposed, $\Psi_{\mathrm{I}} \geq 2$ and $\Psi_{\mathrm{II}} \approx -1$; the distortion is driven to its irreducible level, not to zero. Second, the two statistics price different moments of the same aggregation, and one parameter moves them in opposite directions: the factors measure the within-window *shape* distortion, which shrinks as slow mean reversion flattens the path, while predictability measures cross-window *memory* contamination, which grows with persistence because the carry-over variance $\Gamma_h^2\,\sigma_d^2/(2\rho_d)$ expands and yesterday's shocks survive into today's window almost undamped ($e^{Ah} \to 1$)—persistence relocates the distortion from the window's second moment to the predictable content across windows; it does not shrink it. Third, the persistence-growing component is specific to ex-post recovery from the integral form, which reads the increment through the drift kernel $\Gamma(h-s)$ and so misattributes carry-over content whose weight rises with persistence; the online rate-form recovery of Section 3.3 reads the increment through the pass-through, with unit kernel and contemporaneously, so its residual is orthogonal to the past at any persistence—the mistiming channel that (Snudden, 2024) documents is closed by the representation, and what the factors price is the component that persistence cannot cure.

4.5 The Level-Resolution Factor

The three objects so far live on the flow side. The stock arm of the recovery—the level channel of Section 2.3—has its own factor.

Remark 4.4 (Level-resolution factor)**.** *Define*

$$\Lambda(h) \;\equiv\; \mathrm{corr}\big(\eta_k^x,\; \Delta W_k\big) \;=\; \frac{\sigma_d\,\Gamma_h}{\sqrt{\Sigma_{\eta^x}\,h}} \;=\; 1 - \frac{(Ah)^2}{24} + O(h^4), \tag{39}$$

the correlation, over the k-th observation window $(t_{k-1}, t_k]$, between the state innovation embedded in the sampled level and the same window's structural increment. The level channel's loss to sampling is the gap $1 - \Lambda(h) = (Ah)^2/24 + O(h^4)$: second order in the persistence-scaled interval, against the first-order $O(h)$ resolution effects of the flow channel; the $O(h^4)$ remainder records that the cubic coefficient vanishes identically (Appendix D). The order gap has a one-line reason: correlation is scale-free, so the leading effect of mean

reversion—a common shrinkage of covariance and standard deviation—cancels in the ratio, and only the uneven discounting of shocks across the window survives, a curvature effect of second order.

Derivation: Appendix D. $\Lambda(h)$ is the algebraic content of the collinearity (13) and of the stock entry $c_s = C_s\sigma_d\Gamma_h$ in the recovery operator (22): it is the reason a precisely measured stock reads the contemporaneous increment through the level, and the reason that channel gains *quadratically* from higher sampling frequency—$1-\Lambda$ falls ninefold from quarterly to monthly under the calibration, and is 10^{-5}-sized at a daily interval. The asymmetry $\{1-\Lambda = O(h^2)\}$ versus $\{\Phi = O(h)\}$ is the formal statement of why frequent stocks are, in this model, almost direct reads of the shock path.

4.6 The Recovery-Correlation Law

The four factors describe tracking; one law ties the recovery of the structural shock to the same objects. Consider any filter whose window-start error $\tilde{e}_{k-1}$ has steady variance P_F and is independent of the window's shocks—true for every construction in this paper—and whose innovations are formed by the exact prediction operators of (21).

Proposition 4.5 (Recovery-correlation law)**.** *For any fixed map $m \in \mathbb{R}^3$ applied to the stacked innovation,*

$$\operatorname{corr}\big(m'\nu_k,\ \Delta W_k\big) \;=\; \frac{m'c}{\sqrt{m'S(P_F)\,m\cdot h}},$$
$$S(P_F) \;=\; S_0 + P_F\,w\,w', \qquad w = \Big[\tfrac{C_y\Gamma_h}{h},\ \tfrac{C_\pi\Gamma_h}{h},\ C_s\,e^{Ah}\Big]', \tag{40}$$

with c as in (22) *and $S_0 = S(0)$ from* (23)*. The correlation is maximised at $m \propto S(P_F)^{-1}c$, with $\operatorname{corr}^2_{\max}(P_F) = c'\,S(P_F)^{-1}c/h$.*

Corollary 4.6 (Recovery ordering equals tracking ordering)**.** *By the Sherman–Morrison identity,*

$$\operatorname{corr}^2_{\max}(P_F) = \frac{1}{h}\left[c'S_0^{-1}c \;-\; \frac{P_F\,(c'S_0^{-1}w)^2}{1+P_F\,w'S_0^{-1}w}\right], \tag{41}$$

decreasing in P_F, strictly so whenever $c'S_0^{-1}w \neq 0$—which holds at every point of the calibrated precision grid; and for a common fixed map m with $m'w \neq 0$ the correlation $m'c/\sqrt{m'S(P_F)m\cdot h}$ is likewise strictly decreasing in P_F. Hence the recovery ranking of any set of filters—under their optimal maps or under one shared map—coincides with their tracking ranking, and the ceilings of Appendix B are the evaluation of (41) *at the exact filter's P.*

Corollary 4.7 (Filter-invariance of the shock content)**.** *The vector c depends only on the data-generating process, not on the filter: the pass-through and the level channel sit in* every *filter's innovations. Filter choice enters the recovery only through the scalar P_F.*

Proofs: Appendix D. The law explains, in one line each, the recovery patterns the simulations document. Corollary 4.7 is why the integral-form EDT and the Euler filter recover the shock nearly identically on the flow arm despite entirely different innovation-covariance models—the shock content of the innovation is data-side, and the two filters differ only through their (similar) P_F; the "endpoint surrogate" of Section 3 remains

the correct account of Euler's *tracking* anomaly but plays no separate role in recovery. Corollary 4.6 replaces per-filter recovery narratives with a single monotone map from tracking error to recovery accuracy. And the ρ-profile of $\mathrm{corr}^2_{\max}$—high on the flow-precise arm through c^f, high on the stock-precise arm through c_s, troughed between—is the U-shape of the stacked recovery, here derived rather than simulated.

4.7 The Two Primitives in One Display

The factors of this section are consequences of one covariance object, worth collecting in a single display. The two primitives share every mean block—the same e^{Ah}, the same accumulator loadings $\Gamma_h = \int_0^h e^{As} ds$, the same observation matrix, the same measurement variances $h\eta_j^2$—and differ in exactly one object: the joint covariance of the per-window reduced-form innovations (s_k, u_k), with s_k the state innovation and u_k^j the flow-accumulator innovations in rate units. To leading order in h,

	Primitive II (rate form)	Primitive I (integral form)
$\mathrm{Var}(s_k)$	$\Sigma_{\eta^x} = \int_0^h e^{As}\Omega\, e^{A's}\, ds$	identical
$\mathrm{Var}(u_k^j)$	$C^j\Omega(C^j)'\, h^{-1}$	$C^j\Omega(C^j)'\, h/3$
$\mathrm{Cov}(u_k^y, u_k^\pi)$	$C^y\Omega(C^\pi)'\, h^{-1}$	$C^y\Omega(C^\pi)'\, h/3$
$\mathrm{Cov}(s_k, u_k^j)$	$\Omega(C^j)'$	$\Omega(C^j)'\, h/2$

(42)

Every rate-form entry is the corresponding integral-form entry plus the pass-through contribution, and the variance ratios of order h^{-2} are the $O(h_b^{1/2})$-against-$O(h_b^{3/2})$ scaling of the mixed-frequency analysis (Section 5.6) squared. The display makes the anatomy of the comparison self-evident: diagonal release-noise parameters η_j can shift only the $\mathrm{Var}(u_k^j)$ row, so the off-diagonal rows—the state–flow block $\Omega(C^j)'$, the paper's organising object, and the flow–flow cross-moment—carry the discriminating power untouched by any free parameter. Everything downstream of (42) is induced: a different innovation covariance gives a different Riccati fixed point, hence different gains and different innovation sequences. (Euler differs through a second channel as well—it approximates the mean blocks, $e^{Ah} \approx I + Ah$—and is the only filter carrying genuine discretisation error on top of the primitive choice.)

4.8 Summary: Four Factors and a Law

The section's objects, each closed-form in the calibration and each attached to one mechanism: the *temporal-aggregation factor* $\Phi(h) = \frac{1}{2}h + O(h^2) \to 0$, the resolution at which the integral channel carries the cross-covariance—removable by finer sampling; the *pass-through penalty* $\Pi(\rho) \geq 1$, the fine-sampling limit of the integral-form filter's excess MSE—irreducible, because it prices the omitted representation term rather than its resolution, and equal to one only at the single ratio ρ^* where the misspecification errors cancel exactly in the gain; the *structural-penalty pair* $\Psi_{\mathrm{I}}(h) \to 2$, $\Psi_{\mathrm{II}}(h) \to -1$, the Euler shortcut's flow-variance error under the two primitives—sign-flipping, so the cost of the naive scheme is not even directionally primitive-free; and the *level-resolution factor* $\Lambda(h) = 1 - (Ah)^2/24 + O(h^4)$, the stock arm's near-immunity to sampling. The recovery-correlation law then maps any filter's tracking error into its recovery accuracy

through one strictly monotone scalar function, making the recovery ranking a corollary of the tracking ranking rather than a separate empirical finding. Figure 2 displays the four factors under the calibration; Tables 3 and 2 tabulate them on the sampling and precision grids. Together they are the algebraic content of the simulations that follow: Φ and Π organise the CT–EDT comparison, $\Psi_{\mathrm{I}}, \Psi_{\mathrm{II}}$ the Euler comparison, and Λ with the recovery law the shock-recovery results.

Remark 4.8 (Frequency effects across channels)**.** *The factors are channel-local, so mixed sampling assigns each channel its own: moving a flow from quarterly to monthly divides its resolution factor $\Phi_j(h_j)$ roughly by three, while the level channel improves* quadratically*, $1 - \Lambda$ falling ninefold. Under mixed sampling the exact representation's steady state is periodic within the quarter, so the recovery-correlation law delivers one attainable ceiling per within-quarter observation pattern—an* identification wave*. With a monthly flow present, the quarter-end increment of the wave is second-order: the quarterly flow's pass-through weights the three monthly increments equally, so its marginal information about the target month beyond the monthly channels is $O(h)$. The wave becomes first-order when the intra-quarter channels are stocks alone. The mixed-frequency simulations document the wave; its counterpart on observed release calendars is developed in (Parra-Alvarez et al., 2026).*

5 Simulation Design and Results

This section evaluates the three filters of Section 3 by Monte Carlo, on the two counts the online criterion of Section 3.1 sets: tracking of the latent state and recovery of the contemporaneous structural increment. The evidence covers two economies. The *feedback-rule* (FR) model of Section 2—one latent state, one structural shock—carries the factor verification; the *partial-adjustment* (PA) model of Section 5.4—two states, two shocks—tests the taxonomy where shock separation becomes substantive. Within each economy, observation designs share one data-generating process. In the *uniform* design all three observables arrive quarterly and synchronously, isolating precision heterogeneity from frequency heterogeneity; in the *mixed* design inflation and the interest rate move to monthly observation while the output gap remains quarterly, the empirically standard configuration in which fast channels surround a slow flow (the partial-adjustment study extends the mixed design to daily sampling of the stock, Section 5.6). The closed-form factors of Section 4 organise the results throughout: Φ and Π the continuous-time-versus-EDT comparison, Ψ_{II} the Euler comparison, Λ and the recovery-correlation law the shock-recovery results, and the channel-local factor arithmetic of Remark 4.8 the passage from uniform to mixed sampling.

5.1 Simulation Design

5.1.1 Data Generating Process

The latent demand shock follows the Ornstein–Uhlenbeck process (1) under the calibration of Table 4, and the flow observables are genuine flows of the same shock: at fine resolution the state and each flow increment share the common Brownian increment, as in (4), so each flow channel carries the cross-covariance $\Sigma_{xy,j} = C_j \sigma_d^2$ of (12) with the state innovation. The demeaned state $x_t = d_t - 1$ is the filtering target.

Simulation is *exact*: for each observation window the triple consisting of the state innovation η_k^x, the flow-integral news $u_k = \int \Gamma(h-s)\,\sigma_d\,dW_s$, and the structural increment ΔW_k is drawn jointly from its closed-form Gaussian covariance, whose blocks reproduce the Van Loan moments of Section 2.2 to machine precision. The quarter-end state then follows the exact recursion $x_k = e^{Ah}x_{k-1} + \eta_k^x$, the structural flow integral is $\Gamma_h x_{k-1} + u_k$, and the published flow rate adds the pass-through $\sigma_d \Delta W_k$ inside $C_j(\cdot)/h$. No simulation discretisation error therefore enters the ground truth, the observables, or the recovery target $\int_k dW$: their joint law is that of the continuous-time model at the observation grid. The uniform design draws at the quarterly grid; the mixed design draws at the monthly grid and aggregates, which yields the identical law for all quarterly functionals.

Measurement noise follows the per-release convention of Section 2.2: each published rate—flow and stock alike—carries one discrete draw of standard deviation $\sqrt{h}\,\eta$, attached at the release, with h the window of that release. Observations are synchronous at fixed frequencies within each design, so that the factors isolate a mechanism independent of timing frictions; what varies is the precision ratio between flow and stock channels, which modulates the pass-through penalty and determines where filter choice is material.

5.1.2 Precision Heterogeneity Grid

The precision ratio $\rho = \eta_{\text{flow}}/\eta_{\text{stock}}$ varies across seven values spanning four orders of magnitude,

$$\rho \in \{0.01,\ 0.1,\ 0.5,\ 1,\ 2,\ 10,\ 100\}, \tag{43}$$

with intensities parameterised symmetrically around $\eta_{\text{base}} = 0.03$: $\eta_{\text{flow}} = \eta_{\text{base}}\sqrt{\rho}$ for both flow channels and $\eta_{\text{stock}} = \eta_{\text{base}}/\sqrt{\rho}$ for the stock, so the geometric mean $\sqrt{\eta_{\text{flow}}\,\eta_{\text{stock}}} = \eta_{\text{base}}$ is constant across the grid. At $\rho = 0.01$ the flows (output gap, inflation) are the precise channels; at $\rho = 100$ the point-in-time interest rate is. The grid brackets both the near-singular flow-precise corner of Section 3.4 and the stock-dominated Anderson–Moore limit.

5.1.3 Monte Carlo Configuration

Each precision ratio uses $N = 1000$ replications over $T = 150$ quarters, with a burn-in of 50 quarters excluded from all reported statistics; replications share random draws across filters, so every comparison is paired. Tracking is scored by steady-state mean squared error against the true simulated state at quarter ends; paired t-statistics accompany every filter comparison in the table notes. Recovery is scored by the correlation between the recovered and true structural increments, pooled over the steady-state window—correlation rather than MSE because the increment is identified only up to proportionality through Σ_{xy}.

5.1.4 Recovery Protocol

Recovery follows the operators of Section 3.3, applied identically to every filter's own innovations. Innovations are formed by the exact prediction maps of (21)—$C_j\Gamma_h/h$ against the window-start estimate for a flow, $C_s e^{Ah}$ for a stock—so that differences across filters reflect only the quality of the state estimate entering the prediction, as the recovery-correlation law (Proposition 4.5) assumes. Two maps are evaluated per filter: the *flow-only* projection (24), which reads the increment from the pass-through in the two flow innovations, and the *stacked* map, which adds the stock innovation with weight

proportional to $C_s\sigma_d\Gamma_h$—the level channel of Section 2.3, whose near-perfect alignment with the increment is the factor $\Lambda(h) = 0.9995$ at the quarterly interval (Remark 4.4). Both maps are fixed, computed once per design from the steady innovation moments of the exact representation, and applied online; the corresponding closed-form attainable ceilings of Appendix B are reported next to every empirical correlation. Under mixed sampling the exact representation's steady state is periodic within the quarter, so the stacked map and its ceiling are computed per within-quarter observation pattern—one map and one ceiling per position—which is the operational form of the identification wave anticipated in Remark 4.8.

5.2 The Feedback-Rule Model: Calibration

The feedback-rule economy is the scalar system of Section 2. Table 4 reports the baseline parameter values drawn from (Posch, 2020; Fernandez et al, 2012) and standard New Keynesian calibrations.

Table 4: Model Calibration (Feedback-Rule Model)

Parameter	**Description**	**Symbol**	**Value**
Latent State Dynamics			
Mean reversion	OU persistence	ρ_d	0.4214
Diffusion volatility	Shock standard deviation	σ_d	0.03
Observation Loadings			
Output gap loading	GDP response to shock	$\bar{y}_d$	0.0669
Inflation loading	Price response to shock	π_d	0.1311
Taylor coefficient	Interest rate response to inflation	ϕ_π	4.0
Steady States			
Inflation target		π_{ss}	0.02
Interest rate		i_{ss}	0.03
Measurement Error			
Base noise scale	Release-noise scale (per-release std $\sqrt{h}\,\eta$)	η_{base}	0.03
Noise convention	One draw per published release	—	§2.2
Precision ratio	Flow/stock variance ratio	ρ	$[0.01, 100]$

The mean-reversion parameter $\rho_d = 0.4214$ implies a half-life of approximately 1.6 years, consistent with business cycle persistence. The observation loadings generate realistic magnitudes: a one-standard-deviation demand shock ($\sigma_d = 0.03$) moves the output gap by $0.03 \times 0.0669 \approx 0.2$ percentage points and inflation by $0.03 \times 0.1311 \approx 0.4$ percentage points. The Taylor coefficient $\phi_\pi = 4$ satisfies the Taylor principle, ensuring determinacy of the underlying New Keynesian equilibrium.

5.3 The Feedback-Rule Model as the Factor Laboratory

With one state and one shock, the feedback-rule economy is where every filter gap is priced by a closed form, and its role in the paper is exactly that: to verify the factors of Section 4 empirically before the partial-adjustment model puts the taxonomy to substantive work.

This subsection reports the verification in compact form; the complete tables and figures are collected in Appendix E.

5.3.1 Uniform quarterly frequency

All three observables arrive quarterly and synchronously ($h = 0.25$), so any separation among filters is attributable to precision heterogeneity and to the representation of the flow, not to observation timing. The tracking results (Table 12, Appendix E) realise the factor predictions point by point. The continuous-time filter's advantage lives where flows are precise and follows the pass-through penalty: its MSE sits a factor of over 115 below EDT's at $\rho = 0.01$, a factor of 2.3 at $\rho = 0.1$, within 5 percent by $\rho = 2$, and the two filters coincide to numerical precision beyond $\rho = 10$—the Anderson–Moore limit, reached because the two-step filter treats the precise stock as an exact discrete update rather than a smeared rate. At the quarterly interval EDT still retains the fraction $\Phi(0.25) = 0.113$ of the cross-covariance through its Van Loan cross block, so its measured deficit sits below the h-free limit $\Pi(\rho)$, exactly as Proposition 4.2 prescribes. The Euler filter, modelling 4.96 percent of the rate-form flow-disturbance variance ($\Psi_{\mathrm{II}}(0.25) = -0.95$, Proposition 4.3), trails the correctly specified filters by up to +11,362 percent in the flow-precise corner; its one local irregularity—sitting below EDT on parts of the flow side—is the endpoint surrogate of Section 3, an accidental, mis-scaled contemporaneous channel, not a substitute for the pass-through.

Recovery reproduces the tracking taxonomy through the recovery-correlation law (Table 13, Appendix E). Three properties organise it. *Attainment*: the CT recovery sits on its closed-form ceiling at every grid point, under both maps, to within 0.002. *Complementarity*: the flow-only recovery declines monotonically in ρ, from 0.986 to 0.066, while the stacked recovery is U-shaped—its right arm restored by the level channel, which lifts the recovery to 0.966 at $\rho = 100$ because the stock innovation aligns with the increment at $\Lambda(0.25) = 0.9995$ (Remark 4.4). *Attribution*: filter choice matters only on the flow arm, where the CT recovery leads EDT and Euler by up to 0.11; from the balanced region rightward the three filters' recoveries coincide, as Corollaries 4.6 and 4.7 require.

5.3.2 Mixed observation frequency and the identification wave

The mixed design S1 moves inflation and the interest rate to monthly observation while the output gap remains quarterly—the canonical configuration in which fast price and financial data surround a slow national-accounts release—rerun on the identical simulated paths as the uniform anchor S0, so every comparison is paired at the path level. The frequency effect splits by representation exactly where the factors say it should (Tables 14 and 15, Appendix E). On the stock-precise side the upgrade is filter-invariant: every filter's S1/S0 MSE ratio falls to 0.34 at $\rho = 100$, the level channel improving quadratically ($1 - \Lambda$ falling ninefold). On the flow-precise side the ratios split by measurement model: the correctly specified filter converts the upgraded monthly inflation channel into its largest gain, an S1/S0 ratio of 0.13 at $\rho = 0.01$, while the integral-form and Euler filters deteriorate outright, with ratios of 3.8 and 3.4—the escalation of the primitive penalty under base-grid refinement. Centrally for the paper's argument, the CT–EDT separation does not close under finer sampling of the flow channel: +333,586 percent under S1 against +11,604 percent under S0 at $\rho = 0.01$, since $\Pi(\rho)$ is h-free by Proposition 4.2 while the omitted pass-through gains relative weight as the window shrinks; Euler's monthly penalty

is governed by $\Psi_{\mathrm{II}}(1/12) = -0.994$—the naive filter models 0.6 percent of the monthly flow-disturbance variance, against 5.0 percent of the quarterly one.

On the recovery side, intra-quarter information raises the recovery of the quarterly increment—by up to 0.27 at $\rho = 1$, largest in the balanced region where the uniform design left the most identification unrealised—while moving the target to the monthly increment lowers the attainable correlation, because a finer target divides the same identifying information across three windows (Table 16, Appendix E). Within the quarter the exact representation's steady state is periodic, and the attainable recovery has one closed-form ceiling per within-quarter position: the identification wave (Figure 12 and per-position values in Table 19, both in Appendix E). Its quarter-end increment—the marginal identification content of the quarterly flow beyond the monthly channels—never exceeds 0.0007 anywhere on the grid, confirming the second-order prediction of Remark 4.8, and the CT empirical correlations attain the per-position ceilings to within 0.008: the wave is an information-set object—the sampling pattern itself, not the method, makes identification oscillate within the quarter.

5.3.3 The anatomy of the filter gaps

Table 5 closes the chain at the filter practitioners actually run: the total CT–Euler gap decomposes multiplicatively into a *resolution* component (CT against EDT$^+$, the exact discrete filter of the data-generating process), a *pass-through* component (EDT$^+$ against EDT), and a *naive-implementation* component (EDT against Euler). The three components behave in three different ways, and the decomposition is the empirical form of the Section 4 asymmetry. The resolution component is economically nil in both designs—statistically indistinguishable from zero under quarterly sampling and bounded by five hundredths of one percent under the mixed design, an $O(h)$ difference in the gain costing only $O(h^2)$ in MSE. The pass-through component *is* the gap: +11,604 percent at $\rho = 0.01$ under quarterly sampling, inflated to +333,586 percent under the mixed design. The naive-implementation component is sign-indefinite—Euler sits significantly above EDT only in a band around $\rho = 0.1$ and significantly below it elsewhere on the flow side, the endpoint surrogate at work—while the totals a practitioner faces against the correctly specified filter remain +11,443 percent (S0) and +291,697 percent (S1) in the flow-precise corner, and +10.3 against +47.6 percent at the balanced ratio. Frequency upgrades are therefore not a substitute for representation: they remove the removable factor, amplify the primitive penalty, and leave the filter ranking intact. The implementation itself is validated separately: Appendix C shows the ODE-realised filter attaining the EDT$^+$ ceiling to within 1.3 percent everywhere on the grid.

Table 5: Anatomy of the Filter Gaps: Resolution, Pass-Through, and Naive Implementation

ρ	CT vs. EDT$^+$ (%)	EDT$^+$ vs. EDT (%)	EDT vs. Euler (%)	Total: CT vs. Euler (%)
Panel A: S0 uniform quarterly				
0.01	+0.01 ($t = -0.81$)	+11604.25 ($t = -220.78$)	−1.39 ($t = +25.35$)	+11443.28
0.10	+0.01 ($t = -1.47$)	+127.02 ($t = -98.33$)	+29.43 ($t = -112.24$)	+193.86
0.50	+0.00 ($t = -1.10$)	+21.19 ($t = -46.63$)	−4.52 ($t = +42.33$)	+15.71
1.00	+0.00 ($t = -0.68$)	+13.29 ($t = -40.01$)	−2.60 ($t = +40.29$)	+10.34
2.00	+0.00 ($t = -0.69$)	+6.32 ($t = -30.78$)	−1.08 ($t = +25.19$)	+5.18
10.00	−0.00 ($t = +1.18$)	+0.71 ($t = -12.61$)	+0.03 ($t = -1.04$)	+0.74
100.00	−0.00 ($t = +0.40$)	+0.01 ($t = -2.08$)	+0.03 ($t = -3.04$)	+0.05
Panel B: S1 mixed				
0.01	−0.05 ($t = +2.38$)	+333585.68 ($t = -222.43$)	−12.51 ($t = +171.04$)	+291697.20
0.10	−0.02 ($t = +2.10$)	+7396.92 ($t = -208.97$)	+12.30 ($t = -186.30$)	+8317.94
0.50	−0.01 ($t = +2.11$)	+89.51 ($t = -114.90$)	+3.03 ($t = -23.47$)	+95.25
1.00	−0.00 ($t = +2.60$)	+56.30 ($t = -98.26$)	−5.56 ($t = +81.83$)	+47.60
2.00	−0.00 ($t = +3.32$)	+32.74 ($t = -80.61$)	−3.08 ($t = +85.62$)	+28.65
10.00	−0.00 ($t = +0.94$)	+4.36 ($t = -31.52$)	−0.33 ($t = +27.42$)	+4.02
100.00	−0.00 ($t = +0.40$)	+0.06 ($t = -3.08$)	+0.00 ($t = -0.47$)	+0.06

Notes: Decomposition of the total CT–Euler steady-state MSE gap into the three components of Section 4, obtained by inserting the exact discrete filter of the data-generating process (Appendix C) and the EDT filter as intermediate benchmarks; the sign convention throughout is the excess MSE of the second-listed filter in percent of the first-listed filter's MSE, and the three components compose multiplicatively to the total up to rounding. The *resolution component* is the CT filter's gap against the exact filter, sized by $\Phi(h)$ (Proposition 4.1): insignificant at every ρ in Panel A ($|t| \leq 1.5$), and in Panel B statistically detectable ($|t| \leq 3.3$) with the positive sign the minimum-MSE property of the exact filter requires, but bounded by 0.05 percent—significant yet economically nil. The *pass-through component* is the exact filter's gap against EDT, the h-persistent content priced by $\Pi(\rho)$ (Proposition 4.2), largest in the flow-precise corner and vanishing as the stock dominates; the mixed design inflates it on the flow side. It is significant at $p < 0.01$ for all $\rho \leq 10$ in both panels; at $\rho = 100$ it remains detectable in Panel B while economically nil (+0.01–0.06 percent). Monte Carlo standard errors of this component (delta method, paired replications) are ±102 (Panel A) and ±2,945 (Panel B) percentage points at $\rho = 0.01$, declining to below one point for $\rho \geq 0.5$. The *naive-implementation component* is Euler's gap against EDT, combining the $O(h^2)$ transition approximation, the innovation-covariance approximation Ωh, and the treatment of flows as point-in-time stocks (Proposition 4.3); unlike the other components it is not sign-definite: in the flow-precise corner Euler's endpoint surrogate places it significantly below EDT—a misspecified stand-in rather than a substitute, as the total column shows both filters remaining orders of magnitude above the correctly specified pair there. Only the gap percentages are reported; the verification filter's own levels are virtually identical to the CT filter's and are documented once, in Appendix C.

5.4 The Partial-Adjustment Model: Two States, Two Shocks

The second economy stays within the same continuous-time New Keynesian framework (Posch, 2020) but replaces the feedback rule with partial adjustment of the policy rate, so that the rate carries its own dynamics and its own structural shock. Under partial adjustment two latent variables evolve autonomously, and the minimal state is two-

dimensional,

$$x_t = \begin{bmatrix} i_t - i_{ss} \\ d_t - 1 \end{bmatrix}, \qquad dx_t = A\,x_t\,dt + B\,dW_t, \qquad W_t = \begin{bmatrix} W_t^i \\ W_t^d \end{bmatrix}, \tag{44}$$

with

$$A = \begin{bmatrix} \theta(\phi_\pi \pi_i + \phi_y \bar{y}_i - 1) & \theta(\phi_\pi \pi_d + \phi_y \bar{y}_d) \\ 0 & -\rho_d \end{bmatrix} = \begin{bmatrix} -1.5752 & 0.5018 \\ 0 & -0.4214 \end{bmatrix}, \qquad B = \begin{bmatrix} \sigma_i & 0 \\ 0 & \sigma_d \end{bmatrix}. \tag{45}$$

At the calibration $\rho = 0.03$, $\rho_d = 0.4214$, $\phi_\pi = 4$, $\phi_y = 0$, $\kappa = 0.8842$, $\theta = 0.5$, the rational-expectations policy functions have coefficients $\bar{y}_i = -0.9761$, $\bar{y}_d = 0.4332$, $\pi_i = -0.5376$, $\pi_d = 0.2509$ (the unique real root of the associated cubic; verified numerically at startup of both codes). The eigenvalues $\{-1.5752,\ -0.4214\}$ are both Hurwitz, so the system is mean-square stable and carries no marginal mode. The block-triangular structure means d_t evolves autonomously while i_t mean-reverts under the influence of d_t through A_{12}—a genuine second state rather than a static image of the first. Volatilities are parameterized by $\sigma_d = 0.03$ and the ratio $\lambda = \sigma_i/\sigma_d$.

The equilibrium map is $y_t = \tilde{C} x_t$ with

$$\tilde{C} = \begin{bmatrix} \bar{y}_i & \bar{y}_d \\ \pi_i & \pi_d \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} C^f \\ C_s \end{bmatrix}, \qquad C^f = \begin{bmatrix} \bar{y}_i & \bar{y}_d \\ \pi_i & \pi_d \end{bmatrix}, \quad C_s = [\,1,\ 0\,]. \tag{46}$$

The stock/flow assignment carries over unchanged from the feedback-rule economy: output gap and inflation are flows, the policy rate is the stock. But the stock now reads the first state directly—C_s is a coordinate selector, not a scaled image of d_t. This changes the level channel of the recovery: the stock innovation embeds the state innovation of the i-equation, whose diffusion loading is σ_i, so the level channel becomes primarily a monetary-shock channel.

The flows are specified under the rate form, their increments carrying both Brownian increments at the policy-function loadings,

$$dy_t^F = C^f x_t\,dt + C^f B\,dW_t, \qquad C^f B = \begin{bmatrix} \bar{y}_i \sigma_i & \bar{y}_d \sigma_d \\ \pi_i \sigma_i & \pi_d \sigma_d \end{bmatrix}. \tag{47}$$

This is the central structural upgrade relative to the feedback-rule model: the pass-through matrix $C^f B$ is 2×2 and generically full rank (it is rank one in the scalar economy). The two flows no longer reinforce a single shock; they carry a two-dimensional imprint that, in principle, separates the monetary from the preference shock contemporaneously. Correspondingly,

$$\Sigma_{xy} \equiv \tfrac{1}{dt}\operatorname{cov}\big(dx_t,\, dy_t^F\big) = B(C^f B)' = BB'(C^f)' = \begin{bmatrix} \sigma_i^2\,\bar{y}_i & \sigma_i^2\,\pi_i \\ \sigma_d^2\,\bar{y}_d & \sigma_d^2\,\pi_d \end{bmatrix} \tag{48}$$

is now a matrix rather than a vector, and the flow-disturbance variance rate $R^f = (C^f B)(C^f B)'$+ release noise is generically rank two before noise (rank one plus noise in the scalar economy, so the Sherman–Morrison collinearity structure behind the scalar $\Pi(\rho)$ argument does not carry over as such). In collected ABCD form (with D loading on

the same Brownian pair as B),

$$C = \begin{bmatrix} C^f \\ C_s \end{bmatrix}, \qquad D = \begin{bmatrix} C^f B \\ 0_{1\times 2} \end{bmatrix}, \qquad BD' = BB'\big[(C^f)',\ 0_{2\times 1}\big], \tag{49}$$

and the integral-form primitive remains exactly the restriction $D^f = 0$; the stock row of D is zero for the same reason as in the scalar economy.

Every published figure carries one discrete draw, exactly as in Section 2.2: the release block stacks the two flow rates and the sampled stock with $\Sigma_{v,h} = h\,\mathrm{diag}(\eta_y^2,\, \eta_\pi^2,\, \eta_i^2)$, h the window (or sampling interval) of the respective release, and precision heterogeneity parameterized by $\rho = \eta_{\text{flow}}/\eta_{\text{stock}}$ around the geometric mean 0.03. The accumulator augmentation for the exact-discretization machinery is four-dimensional ($n + m_f = 2 + 2$),

$$d\begin{bmatrix} x_t \\ y_t^F \end{bmatrix} = \underbrace{\begin{bmatrix} A & 0_{2\times 2} \\ C^f & 0_{2\times 2} \end{bmatrix}}_{\mathcal{A}\ (4\times 4)} \begin{bmatrix} x_t \\ y_t^F \end{bmatrix} dt + \underbrace{\begin{bmatrix} B \\ C^f B \end{bmatrix}}_{\mathcal{B}\ (4\times 2)} dW_t, \tag{50}$$

with $\Omega = BB'$ fed to the Van Loan integrals exactly as in the scalar case. The rate-form exact filter discretizes $(\mathcal{A}, \mathcal{B})$ as written; the integral-form filter (EDT, the exact-discretization construction of (Christensen et al., 2024)) zeroes the bottom two rows of $\mathcal{B}$; the Euler filter point-samples the flows. All filters adapt without modification of form—only dimensions change.

Three conditions that are automatic in the scalar economy are substantive here, and are exactly the assessment criteria of (Rátfai, 2026). First, observability and fundamentalness: $\mathrm{rank}(\mathcal{O}) = 2$ holds at the calibration, and $\tilde{C}B = \tilde{C}\,\mathrm{diag}(\sigma_i, \sigma_d)$ has full column rank iff both volatilities are positive and $\tilde{C}$ has rank two; $\det C^f = \bar{y}_i\pi_d - \bar{y}_d\pi_i = -0.0120$ is nonzero but small, so the flow loadings are nearly collinear and shock separation leans on the stock channel. Second, well-posedness of the correlated-noise reduction: $\bar{Q} = \Omega - SR^{-1}S'$ with $S = [\Sigma_{xy},\, 0]$ is positive semidefinite by the shared-loading (Schur-complement) structure $S = \Omega\,[(C^f)', 0]$, and the stability of the closed loop is verified numerically per grid point. Third, the degenerate limit: as $\sigma_i \to 0$ the system collapses continuously toward the feedback-rule economy driven by d_t alone, with i_t an endogenous filter of d rather than a shock carrier, so the partial-adjustment runs nest the feedback-rule results.

5.5 Uniform Quarterly Design: the Partial-Adjustment Model

The uniform study (`PrecUniform_PA.m`) runs the three filters on the $\lambda \times \rho$ grid, $\lambda \in \{0.1, 0.5, 1, 2, 10\}$[5], $\rho \in \{0.01, \ldots, 100\}$[6], at quarterly frequency with synchronous releases; simulation is exact at the quarterly grid, and the closed-form per-shock stacked recovery maps and ceilings are computed from the exact steady innovation moments. The mean-square orderings are clean across all thirty-five cells: the rate-form filter weakly dominates EDT and Euler everywhere (zero violations), and its empirical recoveries attain the closed-form ceilings up to Monte Carlo error (the S0 columns of Table 6; Figures 13–14, Appendix E). The size of the rate-form advantage in tracking d_t is not monotone in either precision or volatility: it peaks in the *interior* of the grid, at roughly +1,200 to +1,600 percent of the rate-form MSE in the central cells, and reaches its overall maximum of

[5] $\lambda = \sigma_i/\sigma_d$, the shock mix of Section 5.4.

[6] $\rho = \eta_f/\eta_s$, the ratio of flow to stock observation noise.

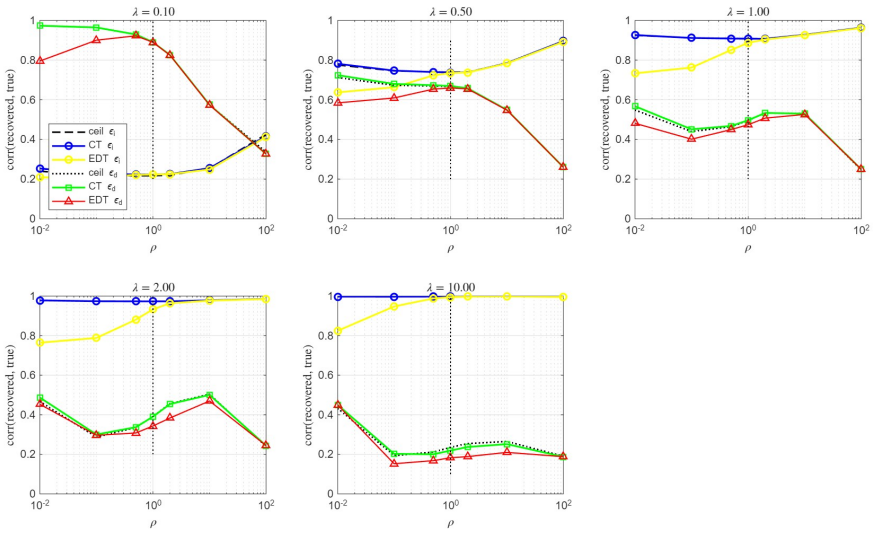


Figure 3: Contemporaneous structural-shock recovery: correlation of the projection estimates with the true quarterly Brownian increments, against the closed-form ceilings (dashed). The monetary shock ε_i is recovered near its ceiling essentially everywhere (0.91–0.999 for $\lambda \geq 1$) because the stock innovation embeds the i-equation state innovation—the level channel is a monetary-shock channel in this model. The preference shock ε_d rides on the flow arm and inherits its sensitivity to ρ and to the near-collinearity of C^f ($\det C^f = -0.0120$). The rate-form recoveries attain the ceilings up to Monte Carlo error; the integral-form recoveries approach them only in the stock-dominated (high-ρ) region.

+11,311 percent at $(\lambda, \rho) = (0.1, 0.01)$, where precise flows carry a strong shared diffusion loading that the integral-form blocks anchor to the window boundary. At the extremes the filters reconverge: with very noisy flows the pass-through content is drowned by release noise, and with a dominant stock all filters lean on the same level channel.

Two findings deserve emphasis. First, misspecified filtering can *destroy* a quasi-observable: at $(\lambda, \rho) = (2, 1)$ the EDT and Euler estimates of the interest-rate state are worse than the raw stock reading itself (MSE ratios to the reading of about 2.7 and 8.0), while the rate-form filter improves on the reading by a factor of three (0.34; Figure 15, Appendix E). Misattributing the flow pass-through forces spurious corrections into a state the data essentially report directly. Second, the two shocks are recovered through different channels: the monetary shock ε_i is recovered near its ceiling essentially everywhere (0.91–0.999 for $\lambda \geq 1$) because the stock innovation embeds the i-equation state innovation—the level channel is a monetary-shock channel in this model—whereas the preference shock ε_d rides on the flow arm and inherits its sensitivity to ρ and to the near-collinearity of C^f ($\det C^f = -0.0120$); see the S0 columns of Table 7 and Figure 3.

5.6 Mixed Observation Frequency: the Partial-Adjustment Model

The mixed-frequency study (`PrecMixed_PA.m`) compares three designs on one set of exact daily paths (30/360 calendar, $h_d = 1/360$): S0, all channels quarterly (the uniform anchor); S1^m, inflation and interest rate monthly with the output gap quarterly; and S1, the paper's design—output gap quarterly, inflation monthly, the interest rate sampled *daily*. (The design labels are redefined for this study: S1 here samples the stock daily, whereas in

the feedback-rule study of Section 5.3.2 it denoted monthly sampling of inflation and the rate.) Filters are real-time consistent by construction: both flows enter through accumulator states observed at their release dates and reset afterwards, the stock updates every base step, and no observation is assimilated before its release. With all information discretely timed, the rate-form filter coincides with the exact discrete filter of the rate-form representation, so the comparison isolates the measurement primitive: rate-form innovation blocks, integral-form blocks, and Euler run on identical machinery with periodic precomputed gains. Grids: $\lambda \in \{0.1, 1, 10\}$, full ρ grid. The data configuration of S1—daily financial readings alongside monthly and quarterly macroeconomic releases—is the one (Posch et al, 2016) use on the estimation side, where the daily rate strengthens identification of the structural parameters. The present design studies the same configuration on the filtering side.

Daily sampling transforms the tracking problem. At $\lambda = 1$, $\rho = 0.1$, the quarter-end steady MSE of $\hat{d}$ falls from 6.3×10^{-4} in S0 to 2.3×10^{-5} in S1—a factor of about twenty-eight—even though d_t is never observed at daily frequency (Table 6; Figure 16, Appendix E). The mechanism is specific to the partial-adjustment structure: d_t enters the drift of the i-equation through A_{12}, so ninety readings of i per quarter identify d through the dynamics of i itself, not only through the flows. The filtered interest-rate state, in turn, beats even the daily raw reading at every grid point (steady MSE between 0.8 and 99.6 percent of the reading variance $h_d \eta_i^2$), and the within-quarter profile of the $\hat{d}$ MSE shows the information geography of the design directly: a slow decay between releases punctuated by drops at month ends and a fall by a factor of about 3.7 on the quarter-end release day.

The integral-form penalty grows under refinement of the base grid. The rate-form advantage over EDT in tracking d escalates from about +490 percent in S0 to +3,900 percent in S1^m and +28,200 percent in S1 at $(\lambda, \rho) = (1, 1)$, approaching +95,000 percent at $\rho = 0.1$—and, when the monetary shock dominates, reaching +3.6 million percent ($\lambda = 10$; Table 6, Panel C), where the integral-form filter's error exceeds the unconditional variance of the state, so the misspecified filter is worse than reporting the unconditional mean while the rate-form error continues to fall. Euler behaves like EDT at the daily step, where its discretization error is negligible and only the temporal-aggregation misspecification remains; its near-instability under the monthly design at $\lambda = 10$ (errors reaching 16.8 at $\rho = 0.01$) is the same misspecification pushed past the stable range. One reversal is instructive rather than anomalous: at $\lambda = 0.1$ and $\rho \leq 1$ the S1 excess falls below the S1^m excess (+1,042 to +483 percent at $\rho = 1$), because the daily stock identifies d through the drift channel so effectively that even the misspecified filter improves faster than the rate-form filter's already-small error. The mechanism is a scaling property of the primitives: per accumulation window the omitted pass-through contributes $O(h_b^{1/2})$ against integral news of $O(h_b^{3/2})$, so refining the base step multiplies the relative weight of the omission.[7]

[7]Both orders follow from a two-line computation on a window of length h_b. The pass-through term $C^f B\, \Delta W_k$ has variance $C^f \Omega (C^f)'\, h_b$, magnitude $O(h_b^{1/2})$—the universal diffusion scaling. The unpredictable content of the accumulated flow beyond x_t is, for small h_b, approximately $C^f B \int_t^{t+h_b} (W_s - W_t)\, ds$, and $\mathrm{Var}\big(\int_t^{t+h_b} (W_s - W_t)\, ds \big) = h_b^3/3$, magnitude $O(h_b^{3/2})$: over a short window the state barely departs from its predictable path, so averaging it adds little news, while the raw increment is pure news at full diffusion strength. The ratio is $O(h_b^{-1})$ in standard deviations, $O(h_b^{-2})$ in variances—halving the base step doubles the relative weight of the omission. The $h^3/3$ law is the engine of the temporal-aggregation literature: the exact double-integral covariance blocks appear in (Bergstrom, 1984) and (Harvey and Stock, 1993), with (Working, 1960) as the classical antecedent on averaging destroying high-frequency

Table 6: Partial-Adjustment Model: Steady-State Tracking of the Demand State $\hat{d}_t$ across Observation Designs

	CT steady MSE ($\times 10^{-4}$)			**EDT excess over CT** (%)		
ρ	S0	S1^m	S1	S0	S1^m	S1
Panel A: $\lambda = 0.10$						
0.01	0.51	0.41	0.11	+11,306	+27,575	+99,570
0.10	0.61	0.43	0.06	+4,450	+10,389	+54,208
0.50	0.94	0.43	0.10	+615	+3,222	+4,026
1.00	1.26	0.48	0.16	+92.5	+1,042	+483
2.00	1.69	0.59	0.27	+15.8	+140	+71.2
10.00	2.60	1.18	0.79	+22.1	+47.4	+59.0
100.00	2.24	1.74	1.49	+2.9	+8.9	+8.6
Panel B: $\lambda = 1.00$						
0.01	5.20	2.88	0.47	+257	+1,325	+76,793
0.10	6.30	3.30	0.23	+277	+1,697	+94,829
0.50	4.18	1.71	0.17	+472	+3,629	+54,976
1.00	3.28	1.25	0.21	+491	+3,933	+28,231
2.00	2.65	0.99	0.32	+365	+3,231	+10,651
10.00	3.02	1.37	1.20	+38.4	+435	+503
100.00	6.59	4.82	4.81	+9.2	+17.0	+17.2
Panel C: $\lambda = 10.00$						
0.01	5.84	3.12	1.01	+86.5	+3,154	$+2.65 \times 10^6$
0.10	7.96	4.27	0.51	+1,486	+32,270	$+3.60 \times 10^6$
0.50	7.26	3.13	0.24	+3,335	+64,258	$+3.49 \times 10^6$
1.00	7.05	2.82	0.24	+3,662	+69,486	$+2.34 \times 10^6$
2.00	6.94	2.66	0.34	+3,555	+64,838	$+1.07 \times 10^6$
10.00	7.03	2.88	1.21	+1,850	+26,940	+84,090
100.00	8.41	5.64	5.16	+95.1	+1,005	+1,161

Notes: Quarter-end steady-state MSE of the filtered demand state against the true simulated state, over $N = 1000$ replications with the burn-in excluded, on identical exact daily paths across designs: S0, all channels quarterly; S1^m, inflation and interest rate monthly with the output gap quarterly; S1, the interest rate sampled daily. $\lambda = \sigma_i/\sigma_d$ is the shock mix and $\rho = \eta_f/\eta_s$ the flow-to-stock precision ratio. Left block: the level for the correctly specified rate-form filter—densifying the designs cuts the CT error monotonically in every cell, by up to a factor of twenty-eight ($\lambda = 1$, $\rho = 0.1$), although d_t is never observed at daily frequency. Right block: the excess MSE of the integral-form filter in percent of the CT MSE—the primitive penalty, which base-grid refinement amplifies by two to four orders of magnitude on the flow side, in line with the $O(h_b^{1/2})$-against-$O(h_b^{3/2})$ scaling of the omitted pass-through; at $\lambda = 10$ under daily sampling the misspecified filter's error exceeds the unconditional variance of the state, while the CT error continues to fall. Paired tests reject equality at $p < 0.001$ in every cell of every design and panel. The Euler filter is omitted for space: at fine base grids its discretisation errors are negligible and it tracks EDT closely, so the reported excess is the primitive penalty for both; full per-filter results appear in the appendix figures. The S0 column is an independent replication of the uniform study on this script's own draws; the two agree within Monte Carlo error cell by cell (e.g., +11,306 against +11,311 percent at $(\lambda, \rho) = (0.1, 0.01)$).

The result is a genuine DGP-mismatch statement, not an implementation artifact: on the counterfactual integral-form DGP (flows accruing without the pass-through term) the ordering flips and EDT beats the rate-form filter, as it must—each filter is exact for its own measurement primitive, and the built-in validation confirms it numerically.

Daily sampling of the stock strengthens identification—of whichever shock is weak. Quarterly per-shock recovery correlations rise monotonically across S0 $\rightarrow$ S1^m $\rightarrow$ S1 in every cell up to Monte Carlo error—two reversals of at most 0.002 among eighty-four comparisons—(Table 7, Figure 4). The striking pattern is that the daily stock rescues the *weakly identified* shock of the configuration: at $\lambda = 0.1$, where the monetary shock is small and hard to see, ε_i recovery jumps from 0.22 to 0.85; at $\lambda = 10$, where the flows are dominated by the monetary shock, ε_d recovery jumps from 0.24 to 0.93; at $\lambda = 1$ the preference shock improves from 0.50 to 0.94. Separation, not just precision, is what the dense stock stream buys: with the level of i pinned down daily, the monetary content of the flow innovations is absorbed and the residual flow information is freed to identify ε_d. The same channel read through the misspecified filters completes the picture: under daily sampling the integral-form and Euler recoveries of the preference shock collapse to at most 0.06 at $\lambda = 10$, against the rate-form filter's 0.77–0.93—densification under the wrong primitive loses the recovery as decisively as the tracking—while in the demand-dominated corner ($\lambda = 0.1$, $\rho = 0.01$) they remain respectable (0.58–0.69), because there both flows are carried by the same shock and separation, the thing the primitive delivers, is not yet at stake.

The identification wave has ninety positions and two shocks. The attainable per-position recovery ceilings, computed in closed form from the periodic Riccati of the exact representation, display the within-quarter information pattern: the ε_i profile is flat near its ceiling every day (the level channel operates daily), while interior-day ε_d ceilings are essentially zero—within a single day the preference shock reaches the interest rate only through $A_{12}h_d$—so the ε_d wave consists of spikes at the flow releases, at month ends and above all at quarter ends (Figure 17, Appendix E). The empirical rate-form recoveries attain the ceilings position by position up to Monte Carlo error (mean absolute deviation below 0.002). One convention should be recorded: the per-position maps are optimal for the *daily* increments, and the monthly and quarterly recoveries reported alongside are summation aggregates of the daily estimates against the aggregated targets, not re-optimized coarse-grid estimators; their high correlations (e.g., ε_d monthly around 0.79 at $\lambda = 1$, $\rho = 0.1$, against near-zero interior daily ceilings) reflect the concentration of the monthly signal at the release positions.

information; the display (42) records the same scaling as the leading-order innovation covariances of the two primitives.

Table 7: Partial-Adjustment Model: Quarterly Per-Shock Recovery across Observation Designs (Stacked Map)

	ε_i **(monetary), CT**			ε_d **(preference), CT**			ε_d**, S1 misspecified**	
ρ	S0	S1^m	S1	S0	S1^m	S1	EDT	Euler
Panel A: $\lambda = 0.10$								
0.01	0.245	0.251	0.413	0.974	0.977	0.978	0.576	0.689
0.10	0.223	0.222	0.683	0.965	0.974	0.977	0.835	0.846
0.50	0.220	0.237	0.828	0.929	0.959	0.964	0.938	0.935
1.00	0.221	0.256	0.852	0.890	0.943	0.947	0.930	0.931
2.00	0.223	0.289	0.850	0.826	0.911	0.915	0.900	0.902
10.00	0.257	0.413	0.757	0.573	0.744	0.748	0.739	0.741
100.00	0.425	0.556	0.631	0.333	0.418	0.458	0.454	0.455
Panel B: $\lambda = 1.00$								
0.01	0.925	0.926	0.935	0.553	0.601	0.771	0.195	−0.010
0.10	0.912	0.913	0.974	0.444	0.485	0.886	0.319	0.353
0.50	0.910	0.920	0.989	0.467	0.558	0.937	0.395	0.415
1.00	0.909	0.928	0.993	0.497	0.622	0.938	0.418	0.436
2.00	0.909	0.939	0.995	0.535	0.690	0.920	0.434	0.450
10.00	0.927	0.971	0.995	0.530	0.716	0.767	0.431	0.445
100.00	0.964	0.981	0.984	0.243	0.371	0.372	0.298	0.303
Panel C: $\lambda = 10.00$								
0.01	0.997	0.999	0.997	0.438	0.636	0.665	−0.003	−0.034
0.10	0.997	0.999	0.999	0.197	0.323	0.770	0.039	0.042
0.50	0.998	0.999	1.000	0.213	0.413	0.908	0.048	0.051
1.00	0.998	0.999	1.000	0.235	0.480	0.926	0.051	0.054
2.00	0.999	0.999	1.000	0.254	0.535	0.916	0.053	0.055
10.00	0.999	0.999	1.000	0.264	0.565	0.767	0.056	0.058
100.00	0.996	0.999	1.000	0.187	0.345	0.370	0.056	0.059

Notes: Correlation between the recovered and true quarterly structural increments under the stacked map, over $N = 1000$ replications with the burn-in excluded; recoveries in the mixed designs are computed at the native daily granularity and aggregated to the quarterly target, per the convention of Section 5.6. The CT columns strengthen monotonically across S0 $\to$ S1^m $\to$ S1 in every cell up to Monte Carlo error (two reversals of at most 0.002 among eighty-four comparisons), and the daily stock rescues the *weakly identified* shock of each configuration: at $\lambda = 0.1$ the monetary shock rises from 0.22 to 0.85, at $\lambda = 10$ the preference shock from 0.24 to 0.93, and at $\lambda = 1$ the preference shock from 0.50 to 0.94 (all at $\rho = 1$). The CT recoveries attain the closed-form per-position ceilings with mean absolute deviation at most 0.002 across all ρ and daily positions (identification-wave figures in the appendix). The last two columns report the S1 recoveries of the misspecified filters for the preference shock—the object the flow arm must deliver: under daily sampling the integral-form and Euler filters lose it almost entirely at $\lambda = 10$ (≤ 0.06 against the CT filter's 0.77–0.93), and in the flow-precise corner their point estimates turn (insignificantly) negative—the primitive penalty operating on recovery exactly as on tracking.

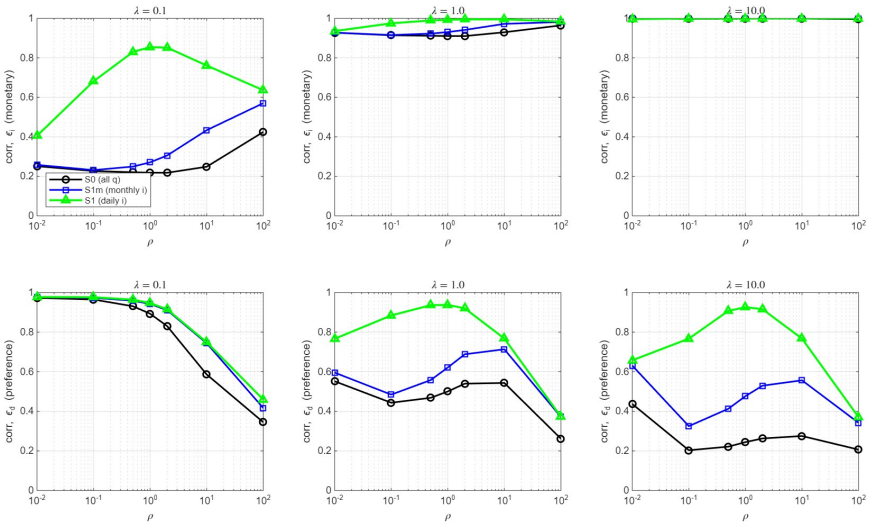


Figure 4: Per-shock quarterly recovery correlations across designs S0 $\rightarrow$ S1^m $\rightarrow$ S1: monotone strengthening in every cell up to Monte Carlo error (two reversals of at most 0.002 among the eighty-four comparisons of Table 7). The daily stock rescues the *weakly identified* shock of each configuration: at $\lambda = 0.1$ the monetary shock jumps from 0.22 to 0.85; at $\lambda = 10$ the preference shock jumps from 0.23 to 0.93. Separation, not just precision, is what the dense stock stream buys—with the level of i pinned daily, the monetary content of the flow innovations is absorbed and the residual flow information is freed to identify ε_d.

5.7 Take-aways: What the Two Model Simulations Establish

The feedback-rule study is the factor laboratory. With one shock and one state, every gap between the filters is priced by the closed forms of Section 4: the CT–EDT separation decomposes into the removable resolution component $\Phi(h)$ and the irreducible pass-through penalty $\Pi(\rho)$, material on the flow-precise arm and equal to one only at the interior ratio ρ^*; the Euler comparison realises the structural-penalty pair $(\Psi_{\mathrm{I}}, \Psi_{\mathrm{II}})$, whose strictly positive limit no sampling frequency removes; the stock channel's near-immunity is the level-resolution factor $\Lambda(h)$; and the recovery results reproduce the tracking taxonomy through the recovery-correlation law, with the stacked operator spanning the precision range through its two complementary arms. At the stock-dominated extreme the filters converge—the Anderson–Moore limit—and at the flow-precise extreme they separate maximally, exactly as the factors prescribe.

The partial-adjustment study then shows the same taxonomy operating where separation becomes substantive, and two channels organize everything it found. The first is the *primitive channel*, moving from the integral-form filters (Euler, EDT) to the rate form: on the rate-form economy the rate-form filter weakly dominates in every one of the thirty-five uniform cells and attains the closed-form recovery ceilings up to Monte Carlo error; the advantage is non-monotone, peaking in the interior of the (λ, ρ) grid where precise flows carry a strong shared diffusion loading; and misspecification is not a graceful degradation—it can destroy a quasi-observable outright, producing interest-rate estimates worse than the raw reading the filter was handed. The second is the *frequency channel*, moving from uniform quarterly to mixed daily observation: daily sampling of the stock

alone cuts the steady tracking error of the never-daily-observed demand state by a factor of about twenty-eight, beats the raw daily reading at every grid point, and—the finding with the most empirical consequence—rescues the *weakly identified* shock of whichever configuration prevails, lifting recovery correlations from the 0.2–0.5 range to 0.85–0.94 by absorbing the monetary content of the flow innovations and freeing the residual flow information for the preference shock: separation, not merely precision. The channels are complements with an ordering: refining the observation grid *amplifies* the primitive penalty (from roughly +490 percent quarterly to +28,200 percent daily at the central cell—and to millions of percent when the dominant shock is monetary—by the $O(h_b^{1/2})$-against-$O(h_b^{3/2})$ scaling of the omitted pass-through), so the design lesson is to correct the measurement primitive before refining the observation grid—densification under the wrong primitive makes the filter worse faster. Two guarantees make these statements more than simulation folklore: every ceiling is computed in closed form from the exact innovation moments (uniform case) or the periodic Riccati equation (mixed case), with empirical recoveries attaining them position by position; and the validation runs in both directions—on the counterfactual integral-form economy the ordering flips exactly as it must, so the ranking follows the data-generating process, not a preference for either primitive.

What the exercise teaches about filtering, in one sentence per model: the single-shock economy establishes that the measurement primitive, not continuous versus discrete time per se, organises filter performance, and prices each gap by a closed-form factor; the two-shock economy establishes that the same organisation survives when the observables must separate increments rather than reinforce one, with the stock channel carrying the separation and dense sampling amplifying—never repairing—a misassigned primitive. What the simulations cannot decide is which primitive actual published releases follow; that is an empirical property of the publication chain; the tests that decide it are stated in Section 6 and taken to observed data in (Parra-Alvarez et al., 2026).

6 Misspecification and Its Tests

The identification of this paper is carried by a full structural commitment, and the natural critique is the classic one: a fully structural device is only as good as its specification. The framework's answer is that the commitment is testable at every point where it could fail. The measurement primitive—the one specification choice on which the filters of this paper genuinely divide—is a testable restriction on observable moments before any filtering; a misassigned primitive is detectable from estimation output after; and the validation of Section 5.7 is two-sided, so the framework detects rather than assumes its own specification. The section runs from localising tests to general ones: the moment restrictions that assign the primitive (Section 6.1); the innovation diagnostics implied by the complete law of the observables (Section 6.2); the testable properties of the recovered shocks themselves (Section 6.3); parameter stability across sampling frequencies (Section 6.4); and the two-sided validation (Section 6.5).

6.1 The primitive as a testable moment restriction

The two primitives share every conditional-mean block and differ in exactly one object, the joint covariance of the per-window innovations (42). Because the release-noise parameters η_j are diagonal and free, the discriminating power sits in the moments they cannot reach.

To leading order in h, restating the two discriminating rows of (42),

$$\begin{aligned} \operatorname{Var}\left(u_k^j\right) &= C^j \Omega \left(C^j\right)' \times \begin{cases} h^{-1} & \text{(rate form)}, \\ h/3 & \text{(integral form)}, \end{cases} \\ \operatorname{Cov}\left(u_k^y,\ u_k^\pi\right) &= C^y \Omega \left(C^\pi\right)' \times \begin{cases} h^{-1}, \\ h/3, \end{cases} \end{aligned} \tag{51}$$

with u_k^j the flow innovations in rate units, net of the (common, computable) carry-over and release-noise terms. Two tests follow. The *frequency-scaling test* uses the variance row (the same moment identifies the pass-through scale τ of Remark 2.1): refining the release window from h_1 to h_2 multiplies the excess innovation variance by h_1/h_2 under the rate form and by h_2/h_1 under the integral form—opposite scalings that no free parameter can reconcile, so observing the same channel at two frequencies identifies the primitive outright.[8] The *cross-channel test* uses the covariance row: two flows over a common window share the pass-through, so their innovation cross-moment is of order h^{-1} under the rate form and $h/3$ under the integral form—and since diagonal release noise cannot generate an off-diagonal moment at all, the comparison discriminates within a single frequency.[9]

6.2 Innovation diagnostics

Under the null—the maintained model, parameters, and primitive—the filter's one-step innovations ν_k form a martingale-difference sequence with the closed-form covariance S_ν of (23). The maintained structure thereby implies a complete, known law for the observables, and everything about that law is testable; with calibrated parameters the tests are sharp, since no estimation step is available to absorb a violation. Three diagnostics follow. *Whiteness*: the standardised innovations $S_\nu^{-1/2}\nu_k$ are serially uncorrelated under the null, so a portmanteau statistic per channel and jointly tests the entire dynamic specification at once—wrong persistence, a misassigned primitive, and omitted dynamics all leave memory in the innovations.[10] *Covariance match*: the sample innovation covariance is compared with the model-implied S_ν—the general form of the targeted moments of Section 6.1.[11] With more observables than shocks, S_ν carries overidentifying structure—the low-rank pass-through block of (16) plus diagonal noise—so the near-singular configuration that

[8] The statistic is a variance ratio. With T_1 and T_2 windows at the two frequencies, the sample excess variances $\hat{v}_1, \hat{v}_2$ are scaled χ^2 under Gaussianity, and the standardised log ratio $\left[\log(\hat{v}_1/\hat{v}_2) - \log r_p\right]/\sqrt{2/T_1 + 2/T_2}$ is asymptotically standard normal under primitive p, with null ratios $r_{\mathrm{II}} = h_2/h_1$ and $r_{\mathrm{I}} = (h_2/h_1)^{-1}$ (up to the $O(h)$ corrections of the exact blocks, which the implementation uses in full); the test selects the primitive whose null is not rejected. Innovation serial dependence under a misspecified filter is accommodated by a HAC variance in place of $2/T$.

[9] The statistic is the delta-method t ratio of the sample cross-moment of one-step-ahead flow innovations against its value under each null, with a HAC standard error; equivalently a one-restriction GMM distance. Both tests are stated for known parameters, the setting of this paper; under estimated parameters the usual first-stage correction applies.

[10] A Ljung–Box statistic on the standardised innovations, per channel and on the stacked vector. Under mixed sampling the exact representation's steady state is periodic, so standardisation is per within-quarter position, using the position-specific innovation covariances already computed for the ceilings; the joint statistic is χ^2 with the usual degrees of freedom.

[11] A quadratic-form χ^2 distance on the distinct elements of the sample covariance around S_ν, HAC-corrected for the serial dependence a misspecified filter induces.

weakens likelihood estimation strengthens the audit: proximity to singularity multiplies the testable restrictions. *Density calibration*: the predictive distributions are fully specified, so the probability integral transforms of the standardised innovations are i.i.d. uniform under the null (Diebold et al., 1998).[12] The comparative form of this diagnostic is already in the empirical record: on US releases the integral form fails exactly this calibration—predictive standard deviations roughly eight-fold too small on the inflation channel—where the rate form does not (Parra-Alvarez et al., 2026). One caveat governs the layer: these diagnostics test the joint of model, parameters, and primitive, so a rejection does not localise by itself—which is why the section runs the localising tests first—and under calibrated parameters a rejection may indict the calibration rather than the structure, a distinction the estimation extension converts into standard post-estimation diagnostics. The diagnostic has power precisely where the tracking comparison loses it: at ρ^* the two errors of the misspecified filter cancel in the gain but not in the predictive distribution, whose variance remains understated, so density calibration separates the primitives at the one precision ratio where MSE cannot.

6.3 Recovered-shock diagnostics

The recovered increments themselves carry two model-implied properties that require no knowledge of the true shocks. *Unpredictability*: a structural increment postdates the window's opening, so the recovered series must be orthogonal to every observable dated t_{k-1} or earlier—a zero-coefficient regression of $\hat{\varepsilon}_k$ on lagged data, where the information set may include variables outside the filter's own observables.[13] The test connects the framework to diagnostics both neighbouring literatures already run: structural innovations recovered from temporally aggregated data are substantially predictable from lagged daily information (Snudden, 2024), and informational-sufficiency tests of the same form are the SVAR literature's own instrument (Forni and Gambetti, 2014). Under the rate form the recovery reads the increment through the pass-through, with unit kernel and contemporaneously, so its residual is orthogonal to the past by construction (Section 4.4); the regression verifies that property on data. *Scale*: the variance of the recovered series is pinned by the ceiling—$\mathrm{Var}(\hat{\varepsilon}_k) = c\, S_\nu^{-1} c'/h$ per shock, the squared attainable correlation of Appendix B, not unity and not anything else—so an over- or under-scaled recovery indicts the specification before any external benchmark is consulted.[14] The standardisation evidence of (Parra-Alvarez et al., 2026)—a correctly scaled rate-form recovery against a badly overscaled integral-form one at the identical calibration—is this test in comparative form.

6.4 Parameter stability across frequencies

With parameters held fixed—the setting of this paper—misassignment costs filtering accuracy and nothing more, the graceful degradation priced by $\Pi(\rho)$. With parameters free, quasi-maximum likelihood moves them to whatever values best reflect the missing covariance structure, and no interior value can: the omitted pass-through contributes flow

[12] Berkowitz's likelihood-ratio version on the normal-transformed transforms tests uniformity and independence jointly.

[13] A standard zero-restriction Wald test with HAC standard errors; power comes from including lagged high-frequency information, the design of (Snudden, 2024).

[14] A variance test of the recovered series against its closed-form target, χ^2 under Gaussianity with HAC correction.

moments of order σ_d^2/h where the integral-form model generates moments of order $\sigma_d^2 h$ (42), a ratio no marginal bias closes. The pseudo-true point therefore migrates toward a boundary, and the exact population likelihood locates it.[15] Over the regular range of the precision grid the destination is a single corner: the fitted release-noise variances absorb the pass-through channel by channel, $\hat{\eta}_j^2 = \eta_j^2 + C_j^2\sigma_d^2/h^2$, while $\hat{\rho}_d$ and $\hat{\sigma}_d$ remain essentially at their true values—exactly so as the window shrinks, because the stock channel pins the dynamics. The fitted model declares the flows nearly pure noise, sets their gains toward zero, and drops from the information set the very channels that are most informative in truth. A second route, absorbing the cross-moments through an inflated shock scale and extinguished persistence, is not available, as Gerrit Meyerheim showed in a comment on an earlier draft: the small-window approximation that suggests it fails at exactly the rates it proposes, the path-news variance collapsing to order h where the target is of order h^{-1}. Only in the flow-precise corner of the grid, where the stock is too noisy to anchor the dynamics, does the pseudo-true point leave the noise corner and distort $\hat{\rho}_d$ and $\hat{\sigma}_d$ outright; no closed form applies there, and the point is characterised numerically. What no parameter can absorb at any point of the grid is the flow–flow cross-moment, which diagonal release noise cannot generate, and its weight grows as the window shrinks: the population Kullback–Leibler divergence per observation rises from 0.04 at quarterly to 0.29 at monthly sampling under the calibration—the estimation-side form of the result that densification under the wrong primitive amplifies the misspecification. The noise corner is a degenerate edge at which the state-space family collapses into a qualitatively different one—flows as exogenous noise—so a misassigned primitive changes the model class rather than merely biasing parameters within it; this degeneration is documented on observed data in (Parra-Alvarez et al., 2026).

This phenomenon converts frequency invariance into a specification test. Under the correct primitive, exact discretisation makes the QML estimand frequency-free: (Christensen et al., 2024) prove it for the integral form, and the exact discrete filter of the rate form (Appendix C) defines an equally frequency-invariant likelihood, so parameter estimates are stable across sampling frequencies. (Filter *output* is invariant under neither primitive—finer sampling delivers more information—and the correct primitive guarantees only that it improves monotonically, whereas under misassignment the integral-form filter's flow-precise error worsens toward the floor $\Pi(\rho)$ as the interval refines, Proposition 4.2.) Under a misassigned primitive the pseudo-true point drifts at a computable rate, and the drift sits in the noise block: the fitted flow-noise variances at two sampling windows differ by the factor $\left(\eta_j^2 + C_j^2\sigma_d^2/h_2^2\right)/\left(\eta_j^2 + C_j^2\sigma_d^2/h_1^2\right)$, which approaches the squared frequency ratio—nine, between quarterly and monthly—under fine sampling and equals 2.7 for inflation and 1.5 for the output gap at the calibration, while $\hat{\rho}_d$ and $\hat{\sigma}_d$ barely move. Re-estimating the model at two frequencies and comparing is therefore a specification test of the primitive, its discriminating statistic concentrated on the release-noise estimates—the estimation-side form of the moment restrictions of Section 6.1.[16] Viewed this way, the cross-frequency consistency that (Christensen et al., 2024) report for US national-accounts

[15] Exactly, not by simulation: truth and candidate model are both linear Gaussian state spaces, so the expected log-likelihood of any candidate parameter point under the true law follows from the candidate filter's stationary innovation moments, one Lyapunov equation per point, and the pseudo-true point maximises it. The correctly specified model recovers the true parameters with zero divergence; the computation is implemented in `PseudoTrueKL.m`.

[16] A Hausman-type contrast: under the correct primitive both frequency-specific estimators are consistent for the same point and their standardised contrast is asymptotically normal around zero. Under misassignment the flow-noise block diverges at the ratio above.

data corroborates the assignment rule rather than testing the rate form: their construction is nested here as the restriction $D^f = 0$ in the ABCD form (10); their data—consumption and hours built by accumulation and averaging—sit squarely inside that restriction; and the only market observable in their panel, the real interest rate, enters as a point-sampled stock, which carries no pass-through under either primitive and hence no discriminating power. A panel that discriminates requires a rate-form flow channel—a window increment of a cumulative market observable—which is the design of the companion empirical work (Parra-Alvarez et al., 2026) and of the estimation extension deferred there.

6.5 Two-sided validation

The tests above concern the data; the validation concerns the method. On a synthetic economy of either primitive, the matched filter wins and the mismatched one loses: on the rate-form economy the rate-form filter weakly dominates in every cell and attains the closed-form ceilings (Sections 5.5–5.6), and on the counterfactual integral-form economy the ordering flips exactly as optimality requires, verified numerically in the simulation codes. The framework is therefore a taxonomy equipped with its own instrumentation: it supplies the ceilings that state what any specification can deliver, the factors that price what a misassignment costs, and the moment restrictions that decide the assignment on data. What it does not do is repair a wrong structural model: misspecification beyond the primitive—wrong dynamics, an incomplete shock menu—is not priced by the factors, and its diagnosis runs through external validation against benchmarks that share no input with the filter, part of the deployment of (Parra-Alvarez et al., 2026).

A final observation places the section in the identification tradition. The two routes to structural shocks differ in where their testable content lies: identification under structure imposed *ex post* concentrates its restrictions at that very step, leaving little for the data to falsify—by design, and defensibly so—while the full commitment maintained here spreads its restrictions across the entire law of the observables, so that whiteness, covariance, density calibration, predictability, scale, and frequency stability all become stated tests. What the framework offers the practitioner is that its tests run online, release by release, on the same information the filter processes.

7 Discussion: Economic Real-Time Navigation

The framework of this paper is designed for real-time navigation: at each release the filter returns, from data observed through the current instant alone, an estimate of the latent state and an estimate of the contemporaneous structural increment—the two objects the online criterion of Section 3.1 requires.

The practical case for the pair is the following. Between two GDP releases a policymaker must decide on rates, facilities, and communication over a quarter in which the official picture does not move; the tracker replaces that information-free interval with an estimate that updates at every release on any calendar—each monthly price print, each day's market close. And knowing where the economy is settles only half of the decision, because the right response depends on why it moved: the same fall in activity calls for easing if it reflects a demand shock, for patience if it is measurement noise, and for a different instrument altogether if it is monetary. The recovered contemporaneous shock is precisely that attribution, delivered when the decision is made rather than years later from revised data and full-sample estimation. The ceilings supply the discipline that belongs next to

any such estimate: a computable statement of how much of today's shock the available releases can reveal at all, so that the confidence a reading deserves—and the value of adding faster data—can be assessed before acting, not after.

The second object is not available by standard means. Because the structural shock is a Brownian increment, it is a martingale difference with respect to the observation history: its conditional expectation given past data is zero, so it cannot be forecast, and can only be recovered contemporaneously or reconstructed after the fact. Structural practice has taken the retrospective route. In the SVAR tradition, structural shocks are obtained by rotating reduced-form residuals under identifying restrictions—timing, long-run, or sign restrictions—that are choices of research design (Ramey, 2016), are imposed after estimation, and recover the shocks only when these are invertible from the history of the observables (Fernández-Villaverde et al., 2007); in nowcasting practice, structural interpretation, when offered, is likewise an ex-post step (Bànbura et al., 2013).

The present framework takes the contemporaneous route and relocates the identifying assumptions rather than removing them. The practitioner specifies the economic environment ex ante—a continuous-time structural model, the measurement primitive of each observable as its statistical construction dictates, and the associated precision structure—and the filter then operates within that environment as data arrive. Identification is carried by the specification itself: the pass-through places the structural increment in the innovation of every filter (Corollary 4.7), a fixed linear map reads it out online, and the recovery-correlation law (Proposition 4.5) together with the ceilings of Appendix B states, in closed form, how much of the increment the current information set can deliver. The invertibility requirement of the retrospective approach, which either holds or fails, is thereby replaced by a graded and computable measure of attainable recovery. In the scalar model the identification questions are settled by construction; the conditions a multi-shock extension must satisfy, where the observables must separate distinct increments rather than reinforce a single one, are recorded in Appendix A.

Since the structural shock is recoverable under both primitives, the framework might appear to advocate for the rate form. Recovery under the integral form runs through the drift cross-moment alone—the projection exists, its flow arm is simply closed, consistent with the ex-post route of (Christensen et al., 2024)—while the rate form opens the flow arm through the pass-through, raising the attainable ceiling and constituting the only recovery channel alive at high observation frequency, where interior-position integral-form ceilings are structurally near zero (Section 5.6). The recommendation is therefore not an advocacy but an assignment rule: matching the correct primitive to the release—and the primitive is a property of the *measurement route*, not of the economic variable. The flows at the centre of this paper, output and inflation, are precisely those whose underlying object is a cumulative accrual: income is paid and prices are charged transaction by transaction, so the window's structural shocks are present in the cumulative observable itself—the statement the rate-form primitive (4) takes at face value. Whether a given *release* preserves that content is decided by its publication chain. A figure built as the window increment of a cumulative or point-sampled observable—accrual totals, price changes computed from point-in-time index readings—is rate-form by construction; a figure built by surveying, averaging, interpolating, or smoothing, the typical national-accounts construction, is integral-form, and the exact discrete filter of (Christensen et al., 2024) is its optimal implementation (also implemented in this paper with the online criterion). The same economic flow can therefore be observed through either primitive depending on the measurement route, and the assignment is testable release by release

(Section 6) rather than assumed.

Whether the recovery is informative in a given application is determined by quantities computable before any data are processed. The precision ratio ρ and the persistence-scaled interval Ah locate the application within the closed forms of Section 4: precise flow observables identify the increment through the pass-through, a precise stock identifies it through the level channel (Remark 4.4), and the attainable correlation is lowest in the balanced region where neither channel dominates. The representation of flow variables remains the one implementation requirement for which observation frequency does not offer any aid: the cost of treating a flow as point-sampled does not vanish as the sampling interval shrinks, so the measurement representation must be specified correctly, alongside the other ex-ante choices the framework requires.

Table 8: The primitive test on US releases, by vintage

Release	Vintage	T	F	r_1	$R(0)$	$R(1)$	θ	$\bar{\tau}$	τ_{hi}	τ_{I}	$\bar{c}(0)$	$\bar{c}(\tau_{\text{hi}})$	$F_q \mid F^3$
CPI-U (SA)	first print	498	0.210	0.427	0.405	0.021	0.27	0.00 (0.00)	0.00	0.026	0.596	–	0.84 | 0.01†
	12-month print	498	0.280	0.481	0.469	0.026	0.27	0.00 (0.00)	0.00	0.028	0.565	–	–
	latest vintage	498	0.222	0.472	0.416	0.022	0.33	0.00 (0.00)	–	0.027	0.591	–	–
PCE prices	first print	313	0.384	0.434	0.556	0.034	0.06	0.04 (0.02)	0.02	0.031	0.343	0.639	0.78 | 0.06†
	12-month print	313	0.306	0.416	0.491	0.028	0.13	0.03 (0.03)	0.00	0.029	0.426	0.469	–
	latest vintage	313	0.466	0.511	0.621	0.040	0.06	0.04 (0.02)	–	0.033	0.335	–	–
Core PCE prices	first print	313	0.970	0.269	0.980	0.160	-0.86	0.67 (0.19)	0.48	0.041	0.005	0.595	0.92 | 0.91
	12-month print	313	0.968	0.317	0.979	0.155	-0.83	0.56 (0.14)	0.35	0.041	0.007	0.471	–
	latest vintage	313	0.963	0.436	0.975	0.144	-0.78	0.39 (0.08)	–	0.041	0.012	–	–
Industrial production (SA)	first print	500	0.989	0.993	0.993	0.288	0.27	0.00 (0.05)	0.00	0.042	0.403	–	0.97 | 0.97
	12-month print	500	0.988	0.991	0.992	0.265	0.18	0.02 (0.00)	0.00	0.041	0.349	–	–
	latest vintage	500	0.992	0.994	0.995	0.340	0.17	0.02 (0.00)	–	0.042	0.344	–	–
Real PCE	first print	500	0.992	0.992	0.994	0.332	0.01	0.03 (0.01)	0.00	0.042	0.256	–	0.98 | 0.98
	12-month print	500	0.992	0.993	0.995	0.349	0.03	0.03 (0.01)	0.00	0.042	0.264	–	–
	latest vintage	500	0.993	0.993	0.995	0.360	0.02	0.03 (0.01)	–	0.042	0.260	–	–
Real disposable income	first print	500	0.990	0.948	0.994	0.303	-0.52	0.13 (0.02)	0.00	0.042	0.058	–	0.98 | 0.97
	12-month print	500	0.980	0.889	0.987	0.199	-0.54	0.14 (0.02)	0.00	0.041	0.053	–	–
	latest vintage	500	0.980	0.889	0.986	0.198	-0.54	0.14 (0.02)	–	0.041	0.054	–	–
Retail sales ex-auto (SA)	first print	302	0.968	0.956	0.979	0.155	-0.15	0.05 (0.01)	0.00	0.041	0.184	–	0.94 | 0.91
	12-month print	302	0.962	0.948	0.974	0.142	-0.14	0.05 (0.01)	0.00	0.041	0.187	–	–
	latest vintage	302	0.968	0.959	0.978	0.154	-0.11	0.05 (0.01)	–	0.041	0.200	–	–
Real GDP (quarterly)	first print	138	0.949	0.939	0.966	0.439	-0.08	0.13 (0.03)	0.00	0.123	0.214	–	–
	12-month print	138	0.941	0.932	0.960	0.410	-0.07	0.13 (0.03)	0.00	0.122	0.222	–	–
	latest vintage	138	0.948	0.935	0.965	0.436	-0.11	0.14 (0.03)	–	0.123	0.203	–	–

Monthly releases enter as annualised log changes (prices) or as the monthly amount (quantities: cumulated first-print growth with a linear trend removed), regressed on month dummies; real GDP enters quarterly. First print, twelve-month print and latest vintage are evaluated on the same real-time sample (ALFRED). Standard errors in parentheses are delta-method from the ARMA fit. †: frequency check failed, $|F_q - F^3| > 0.1$; the root of such a release is not a persistence and its ceilings are not interpretable. "–": not defined for that vintage (no later vintage to revise against) or not reported.

- F: the autoregressive root of the release, $F = e^{-\rho_d h}$, from an ARMA$(1,1)$ fit; the implied persistence is $\rho_d = -\log F/h$ per year. Under (69) the release is ARMA$(1,1)$ with this root whatever τ is (Proposition F.1), so F carries no information on the primitive and is the one input the other columns take from the data.
- r_1: the first autocorrelation of the published series, γ_1/γ_0.
- $R(0)$, $R(1)$: the value r_1 would take at the release's own (ρ_d, h) under the integral form and under the rate form if the release carried no noise, $R(\tau; \rho_d, h)$ of (73). Both are free of the loading and of σ_d. Release noise can only lower r_1, so $R(1)$ is an upper bound on r_1 under the rate form and $R(0)$ an upper bound under the integral form (Corollary F.2).
- θ: the moving-average coefficient of the fit. For a growth-rate release a positive θ is the signature of a differenced within-window average ((Working, 1960)); a negative θ close to $-F$ means the release is near white noise around a small persistent component.
- $\bar{\tau}$: the no-noise bound, the solution of $R(\bar{\tau}; \rho_d, h) = r_1$. It is an upper bound on the pass-through scale, with equality only if $\eta = 0$; $\bar{\tau} = 0$ means $r_1 \geq R(0)$, the release is at or beyond the integral-form boundary.
- τ_{hi}: the bound with the revision variance as the noise floor, $R(\tau_{\text{hi}}) = \gamma_1/(\gamma_0 - h\underline{\eta}^2)$ with $\underline{\eta}^2 = \text{Var}(y_k^{\text{first}} - y_k^{\text{latest}})/h$. Revisions are the noise the agency later removed, so $\underline{\eta}^2$ is a lower bound on the first print's noise variance and $\tau_{\text{hi}} \leq \bar{\tau}$. Defined for the first and the twelve-month print.
- τ_{I}: the integral form's own window-shock covariance, expressed in units of τ: $\tau_{\text{I}} = (h - \Gamma_h)/(\rho_d h) \approx h/2$. A release with pass-through τ has window-shock covariance $(\tau_{\text{I}} + \tau)\, C\sigma_d$, so τ is to be read against τ_{I}: at a monthly window, $\tau = 0.04$ doubles the window content of the release.
- $\bar{c}(0)$, $\bar{c}(\tau_{\text{hi}})$: the flow-only attainable recovery ceiling $\text{corr}^2_{\max}$ of Appendix B, computed from the release's own (γ_0, γ_1, F) under two readings of its primitive. Given τ, (70) pins $(C\sigma_d)^2$, (71) pins $h\eta^2$, and the steady-state one-release filter gives the innovation variance; no calibration enters. $\bar{c}(0)$ is the integral reading with the noise it implies; $\bar{c}(\tau_{\text{hi}})$ the pass-through reading consistent with the revision floor, reported where $\tau_{\text{hi}} > 0$.
- $F_q \mid F^3$: the frequency check on first prints. A single state implies that the release aggregated to quarters has root F^3. F_q is the root estimated on the quarterly aggregate. † marks $|F_q - F^3| > 0.1$.

Procedure. For each release and vintage: (i) build the series (growth rate or amount, month dummies); (ii) fit ARMA$(1,1)$ by conditional-sum-of-squares maximum likelihood from several starting values, with the parameters in tanh space so that the root and the moving-average coefficient stay inside the unit circle, and take $(F, \theta, \gamma_0, \gamma_1, r_1)$ and the numerical-Hessian covariance from the fit; (iii) evaluate $R(0)$ and $R(1)$ at (ρ_d, h); (iv) solve $R(\tau) = r_1$ and $R(\tau) = \gamma_1/(\gamma_0 - h\underline{\eta}^2)$ for $\bar{\tau}$ and τ_{hi} by bisection, with delta-method standard errors; (v) compute the ceilings at $\tau = 0$ and $\tau = \tau_{\text{hi}}$; (vi) on first prints, aggregate to quarters and refit for F_q. Steps (ii)–(vi) use the release alone. The procedure's only assumption is persistence made transparent by (69).

8 Empirical Study

As discussed above, the SVAR literature that identifies structural shocks starts out from reduced-form surprises, then spends its efforts to rotate those surprises and translate them into economic structural shocks ((Ramey, 2016; Fernández-Villaverde et al., 2007)). Similarily, high-frequency identification works by restricting time-window around specific events, such as FOMC annouuncments ((Bauer and Swanson, 2023)). Hence, previous work asks the question, which structural shocks can be recovered through the restriction (identification), in contrast, this paper's empirical section deploys whether the surprises contain the structural shocks, at all, and if so, to which degree. Identification then follows freely from the structural model via the estiblashed results of the paper (4.6). This is exactly the testing machinery of primitives, whereas no primitive is being advocated, instead I show in this section, how to test for them, while stating transpaerntly what assumptions and commitments need to be made at certain stages. In other words, previous literature takes it for granted, that reduced-form surprises contain shocks, here, I propose a framework, which tests for it – specification-free with a commitment to the persistence of the latent state. In addition, casting model assumptions, this study shows concretely how to tell up to what scale (τ) a shock is inside a specific data release and what is the attainable ceiling of sturctural shock recovery.

Intuitively, the release's surprise is the part its own past did not predict. Section 2.1 suggests, under integral-form (averaging), the surprise is small and the release inherits the persistence of the state, but under pass-through (rate-form), the surprise is large and white, as it happens to be the structural shock itself at its full weight. Therefore, the test proposed here utilizes the auto-covariance of the releases in the first layer. Lag-one autocovariance over noise-corrected variance is a loading and σ-free statistic, which has null near the root under averaging, and is far below it under pass-through. Noise lowers the observed autocorrelation, so lower values can mean either large shock content or large noise. Accordingly, this test cannot tell those apart, which is formalized by the observational equivalence of the two primitives (F.1): for any τ there is an integral-form twin with a larger loading and a larger noise variance that has the same law. The pass-through variance reappears as noise, its covariance with the carry-over as loading. The §6 noise corner falls into this equivalence class, as well.

How can we then break the equivalence in practice – in order of degree in commitment? One option is to impose information on the noise calculated from the variance revisions versus first releases. This variance, then bounds the noise (η-s) from below and tightens τ – no economics involved. Another option is to include a new stock observable on the same state at the window's end: under the integral form the release's surprise is uncorrelated with the stock's increment, whatever the stock's loading or noise, while under pass-through it co-moves in proportion to τ (Remark F.3). Note, that the new variable is not an instrument/proxy, but tests the null directly and the absence vs presence of correlation is the tell. Finally, the third layer uses the highest order of assumptions, where model structure is introduced via loadings - by fixing them, τ is identified from the release's likelihood. This layer requires commitment to model structure and calibration.

Without the aim of deploying the real-time tracker at full scale (which is deferred to a companion paper), I demonstrate procedures and statistical test to analyze the content of several releases. I show when it is possible to decide without model assumptions and illustrate when release-only test is ambiguous. Finally I introduce the model structure to address final questions and assess identification ceilings of an example series, which

concludes the scope of this empirical study.

8.1 What a Release Carries: Primitive Test and Ceilings

To demonstrate whether a structural shock (or a part of them) is carried within a measure, I rely on the following mechanism that can be read off the releases. First, looking at the autocovarince relative to its own persistence (implied by an OU-process), if it sits close to the root, the shock is averaged in, if far below it, the shock is in the data. Release noise can mix into the shock, and increasing noise would generally lower autocorellation. Nontheless, revisions may help identify release noise, the rest of the whiteness is then attributed to the shock – persistence and noise together pin the degree of shocks carried by the data release.

Three hypotheses which allow for testing and checking them is available merely from the series itself (self-contained testing). 1. Later vintages should move the release toward the averaged reading if revisions remove noise 2. aggregating the release to quarters should give the AR root and 3. a variable published as both primitives from the same quotes should land in the two classes. Finally one can convert the reading into the scale of how much of the window's shock the release can deliver on the day it is published.

Table 8 reports statistics computed from the release alone under the single-state release representation (69): one Ornstein–Uhlenbeck state behind the release, loading C and release-noise scale η free. Note, that no economic content is attached to the state. T is the number of observations.

Three columns settle the primitive of the official releases, and the rate-form primitive is everywhere rejected except for an ambiguous case. First off, r_1 against $R(1)$: CPI first prints have $r_1 = 0.43$ against a rate-form maximum of 0.02, PCE prices 0.43 against 0.03, the quantity releases 0.94–0.99 against 0.15–0.44, real GDP 0.94 against 0.44. The rate form is rejected for every official release, in every vintage, at any noise level. Second, r_1 against $R(0)$: industrial production and real PCE sit on the integral-form boundary to three decimals, and CPI lies above it in every vintage (0.43 against 0.41 on first prints), with $\theta > 0$ throughout—the CPI is smoother than a window integral of one state, which is the within-month averaging of the price level that the gasoline rows of the last block reproduce from point-sampled quotes ($r_1 = 0.36$ for the CPI component, 0.36 for the EIA monthly average, 0.19 for the end-of-month increment). Third, τ_{hi}: once the revision variance is taken as the noise floor, the bound closes at zero for every quantity release, for CPI, and nearly so for PCE prices (0.02). The releases that the no-noise bound left at $\bar{\tau} = 0.05$–0.13 (retail sales, disposable income, GDP) are explained by the noise their own revisions removed.

Core PCE prices are the exception and the open case. The first print has $F = 0.97$ with $\theta = -0.86$, a near-cancelling fit, and $r_1 = 0.27$ against $R(0) = 0.98$ and $R(1) = 0.16$; the bound is $\bar{\tau} = 0.67\,(0.19)$ and, with the revision floor, $\tau_{\text{hi}} = 0.48$. The two ceilings make the stake explicit: read as integral form the first print is 73 percent release noise and delivers an online recovery correlation of 0.07 ($\bar{c}(0) = 0.005$). Read with τ_{hi} and the revision floor, the correlation is 0.77 ($\bar{c} = 0.595$). Proposition F.1 says the release itself cannot choose between these readings – this section delivers how the decision is made. Across vintages the bound falls, $0.67 \to 0.56 \to 0.39$, and the implied noise share under the integral reading falls with it, $0.73 \to 0.68 \to 0.55$: revision moves the release toward the integral form, as it should if revisions remove noise.

The frequency check separates headline from core. CPI ($F_q = 0.84$ against $F^3 = 0.01$)

and PCE prices (0.78 against 0.06) fail by an order of magnitude: their monthly root, $\rho_d \approx$ 12–19 per year, is the fast component rather than a persistence, and the single-state representation does not describe them. Core PCE (0.92 against 0.91), industrial production, real PCE and disposable income pass. Retail sales are marginal (0.94 against 0.91). The ceilings of the two headline price releases are reported for completeness and carry the †.

Table 9: Unrevised market prices measured as both primitives

Variable	Series	T	r_1	θ	b_m	b_{m-1}	R^2	R^2_{m-1}
USD/EUR	end-of-month increment	331	-0.010	-1.00	–	–	–	–
	monthly average	331	0.311	0.22	0.55 (0.02)	0.45 (0.02)	0.80	0.34
WTI crude	end-of-month increment	427	0.121	0.52	–	–	–	–
	monthly average	427	0.249	0.33	0.62 (0.03)	0.44 (0.03)	0.79	0.33
Gasoline	end-of-month increment	427	0.188	0.34	–	–	–	–
	monthly average	427	0.363	0.54	0.68 (0.02)	0.41 (0.02)	0.93	0.42
	CPI component (NSA)	427	0.359	0.50	0.61 (0.02)	0.45 (0.02)	0.90	0.48

Unrevised market prices, 1991–2026, monthly log changes. The end-of-month increment uses the last quote of each month (daily quotes for the dollar–euro rate and WTI; weekly Monday quotes for retail gasoline, so the "end of month" is the last Monday); the average is the published monthly mean of the same quotes; the CPI gasoline component (NSA) is the official index built from gasoline prices. r_1 and θ are computed on series regressed on month dummies; the regression columns use the raw quotes on both sides. Newey–West standard errors (three lags) in parentheses.
b_m, b_{m-1}: coefficients from regressing the averaged series' monthly change on the end-of-month increments of month m and of month $m-1$. Under the integral form the weights are $(n+1)/2n$ and $(n-1)/2n$ for n quotes per month, about 0.52 and 0.48; under the rate form they would be 1 and 0.
R^2: of that regression. R^2_{m-1}: of the regression on last month's increment alone, i.e. the share of what the averaged series reports for month m that was already observable, point-sampled, at the end of month $m-1$.

Table 8.1 displays unrevised market prices, each published in two ways from the same series: the end-of-month increment and the change in the monthly average, with the CPI gasoline component as the official average. The state behind a market price's change is its expected drift, which is too weak relative to the increment for a root to be estimated, so no τ bound is reported. The benchmark is (Working, 1960): $r_1 = 0$ for the increment of a random walk, 0.25 for its differenced average.

This second exhibit shows the mechanism on variables for which both measurement prmitives exist. The end-of-month increments have $r_1 = -0.01$ (dollar–euro), 0.12 (WTI) and 0.19 (gasoline), the first with $\theta = -1$, white as the increment of a random walk should be. The energy increments carry some autocorrelation of their own because retail prices adjust to crude over several weeks. Averaging the same series moves every series into the other class: r_1 rises to 0.31, 0.25 and 0.36 and θ turns positive in all three, by the amount the random-walk benchmark predicts. The official CPI gasoline component has $r_1 = 0.36$ and $\theta = 0.50$, against 0.36 and 0.54 for the EIA average and 0.19 and 0.34 for the increment: the official index is exactly the average construction.

The regression columns say what that construction does to timing. The averaged

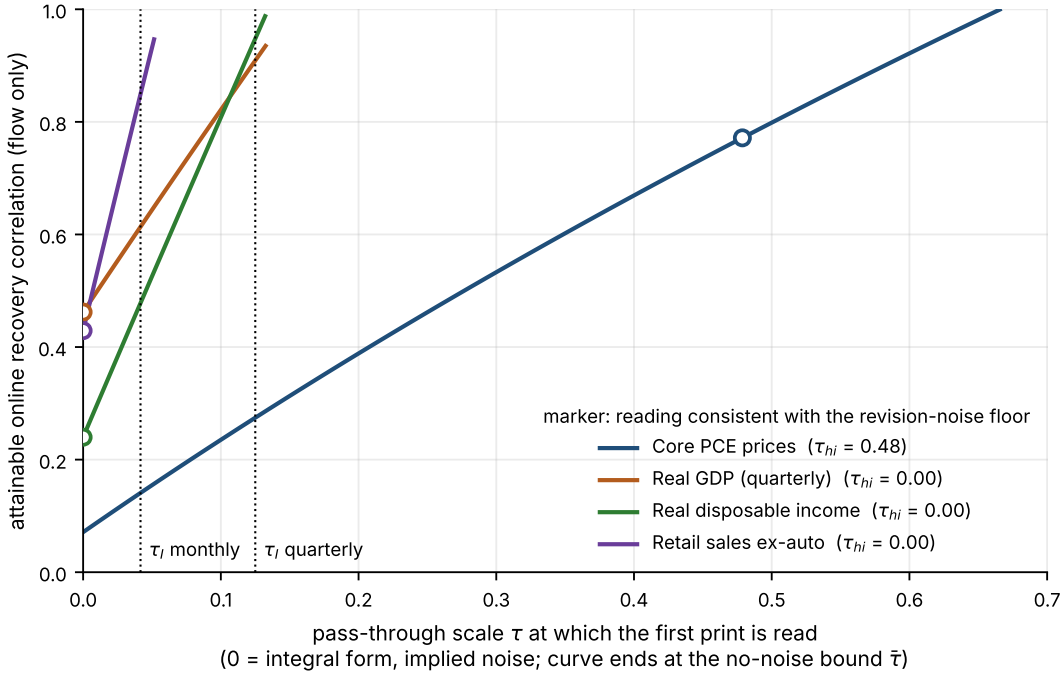


Figure 5: Attainable online recovery from one release, as a function of how its primitive is read. Each curve is the flow-only ceiling of Appendix B, $\mathrm{corr}_{\max} = \sqrt{c^2/(S_\nu h)}$, for one first print (core PCE prices, real GDP, real disposable income, retail sales ex-auto; samples as in Table 8), computed from the release's own (γ_0, γ_1, F) under the single-state release representation; no calibration enters. The horizontal axis is the pass-through scale τ at which the release is read: at $\tau = 0$ the integral reading with the noise it implies, at the right end the no-noise bound $\bar{\tau}$. Given τ, (70) fixes the signal scale, (71) the release noise, and the steady-state one-release filter the innovation variance. τ_{hi}, the reading consistent with the revision variance as noise floor. Dotted lines: $\tau_{\mathrm{I}} = (h - \Gamma_h)/(\rho_d h)$, the integral form's own window-shock content in units of τ, monthly and quarterly. Headline CPI and PCE prices are omitted because they fail the frequency check.

change loads on the current and the previous month's increment with weights 0.55/0.45 (dollar–euro) and 0.62/0.44 (WTI), at the integral-form kernel's 0.52/0.48; gasoline tilts to 0.68/0.41 because its last weekly quote falls several days before the month closes. The CPI gasoline component gives 0.61/0.45 with $R^2 = 0.90$, and last month's increment alone explains 0.48 of the variance of what the index reports this month. Between a third and a half of each averaged release was therefore observable, point-sampled, before the month it describes had begun. The increment carries the window's shock when it occurs; the average carries it at nearly half weight a window later. That is the pass-through content of Section 2, seen on data where the two routes can be compared directly, and it is why the categorization of Section 8 treats a point-sampled level as a stock to be fed on its own calendar rather than as a flow to be averaged.

Next, on Figure 5, each curve answers the following *real-time* quesiton: given first prints and that they are read along τ axis, how much of the window's structural shock could be recovered on the exact day it is published, a. k. a. first fed into the tracker. The tables in this section diagnosed three of the four displayed series as integral form, hence their flow arm is essentially closed ($\tau = 0$), so their marker sits on the far left,

which translates to integral form. The rate form lies on the far right, where the entire pass-through is included, the scale then stretches as a positive function between the two, capturing the fact the higher the pass-through content, the better the sturctural recovery gets. For core PCE the two readings are 0.07 and 0.77 at the marker, which means it has a great potential for recovery due to the large amount of whitness it contains, but the release alone cannot decide whether it is release noise or shock – only the identification step coming up in the next subsection.

The tests above reject rate form for every official aggregates – with one left as an open question (Core PCE) as to whether it retains a part of the shock content. Therefore, the attainable online recovery from the flow arm is what the integral form can deliver and whatever within-window inforamtion about structural shock has to come from the stock arm (or the core PCE with an open question). It reconciles with the paper mechanism as follows: the stock innovation and the flow pass-through are two readings of the same window increments. The simulations held the stock fixed and varied the flow, so their gaps are the value of the flow arm given the stock arm, hence the value of the shock content could be demonstrated which is still valid. However, the data close the flow arm for the official aggregates and leave the stock arm open by the result of unrevised market prices. The mixed-frequency results describe the US configuration the best, where the daily stock contributes the most to identification.

Everything by far was analyzed by the release only. The loadings of the stocks, the cross-moments between arms, and the separation of two shocks are model objects, therefore, require the introduction of the structure, which follows next. The PA calibration enters and the stacked ceilings are reported as functions of τ so that its weight becomes visible.

8.2 Identification and Recovery

To complement the release-only assessment, I address the question the core PCE series left unanswered by introducing the Partial Adjustment model structure[17], and demonstrate identification role as a flow variable on this specific series which concludes this empirical study.

The reason this section proceeds with the core PCE is that it left its categorization as a rate vs integral form unanswered. Previous tests suggested 73% white variance, $\bar{\tau} = 0.67, \tau_h i = 0.48$ and ceilings 0.07 vs 0.77 under the two primitives. Here, pair-statistics with a second variable (including the same shock) and structural likelihood test are run for the Partial Adjustment economy: A and loadings kept fixed, while $\sigma_i, \sigma_d, \eta_\pi, \eta_s, \tau$ is left free. Then the stacked ceilings (as opposed to the flow-only in the previous section) are shown through the data-assigned τ-s. Based on likelihood, rate or integral form can be rejected, and final identification ceilings reported for the specification of the paper.

To close the assessment of the core PCE, Figure 6 decides through paired likelihood, that rate-form is, in fact, rejected. The best fit is at $\tau = 0$, which is precisely the integral form. The curves show how quickly the fit deteriorates as τ rises, and the band shows which τ-s are statistically indistinguishable from the best. The slope is slighter with the stock loadings free and becomes steeper via imposing restrictions and introducing model content (dashed vs solid curves). On its own the pair statistics with the 2-year Tresury yield says little about τ, in contrast, the breakeven data pins the pass-through scale better, since both PCE and breakeven load on inflation state well. Note, that withour

[17]Calibration is borrowed from simulation

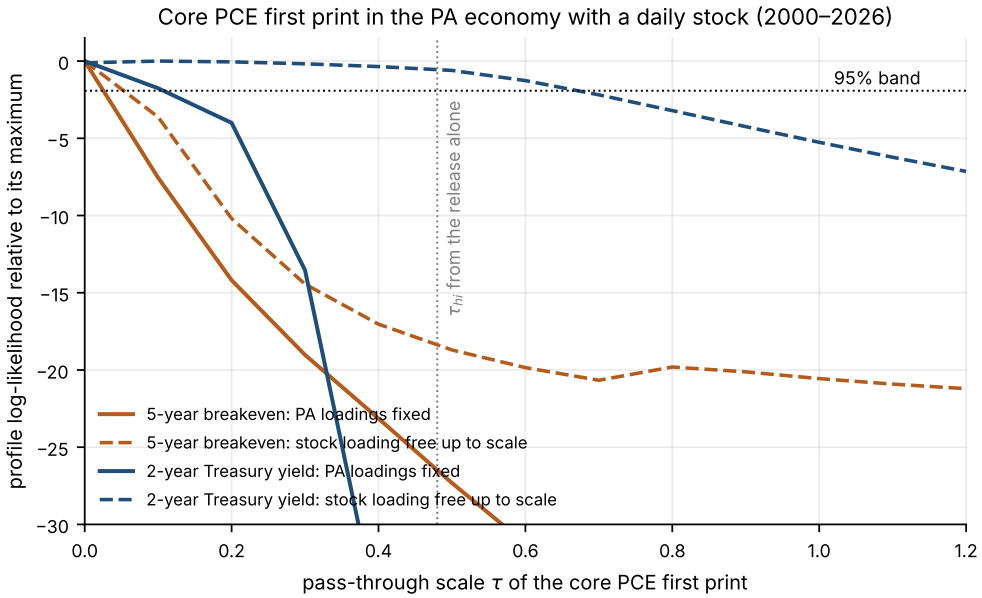


Figure 6: Closing the core PCE interval with a second series on the same state. Profile log-likelihood over the pass-through scale τ of the core PCE first print (2000–2026, monthly), relative to its maximum, in the two-state partial-adjustment economy with A at the calibration of Section 5.4 and $(\sigma_i, \sigma_d, \eta_\pi, \eta_s, \tau)$ free. The second observable is a daily stock sampled at the end of each month: the 2-year Treasury yield, loaded as the expected policy-rate path $(1/2)\, e_1'\Gamma(2)$, or the 5-year breakeven, loaded as expected inflation $(1/5)\, C_\pi\Gamma(5)$, with constant term premia. Solid lines fix these loadings; dashed lines leave the stock loading free up to a scale, so that identification comes from the pair alone. The dotted horizontal line is the 95% profile band; the vertical line is $\tau_{\text{hi}} = 0.48$, the bound the release alone delivers with the revision variance as noise floor (Table 8). All four profiles are maximised at or next to $\tau = 0$ and exclude τ_{hi} by more than ten log-likelihood points with the loadings fixed: the white component that left the release's own reading open is release noise, not pass-through.

adding further stock series to the display, Figure 7 would exhibit a flat line under the twin proposition. Even though the single-state test was unable to assign primitive category to core PCE, this figure depicts the procedure of adding more series and eventually model information rejects the rate-form representation.

Figure 7 illustrates the identification of structural shocks in different mixed frequency regimes. It reports the attainable recovery ceilings which can be computed without having to apply the framework (4.6). Monetary policy shock recovery remains table and flat across regimes and precision ratios (ρ) and the reason is that the daily stock data carries information about it, which is unaffected by the primitive.

As for the preference shock, the daily stock alone cannot rescue the identification, where recovery falls from 0.95 to 0.35 at balanced precision, and even further to 0.15 when flows are noisy. The counterfactural of (0,0.48)-curve tells the story of a partial pass-through in inflation, and still no pass-through in GDP, while the (1,1)-curve depicts the rate-form counterfactual – both rejected by the proposed test procedure and both obtain better shock identification than their (0,0) counterpart, the integral form. However, introducing a second daily series (daily breakeven), that also loads on the shock through

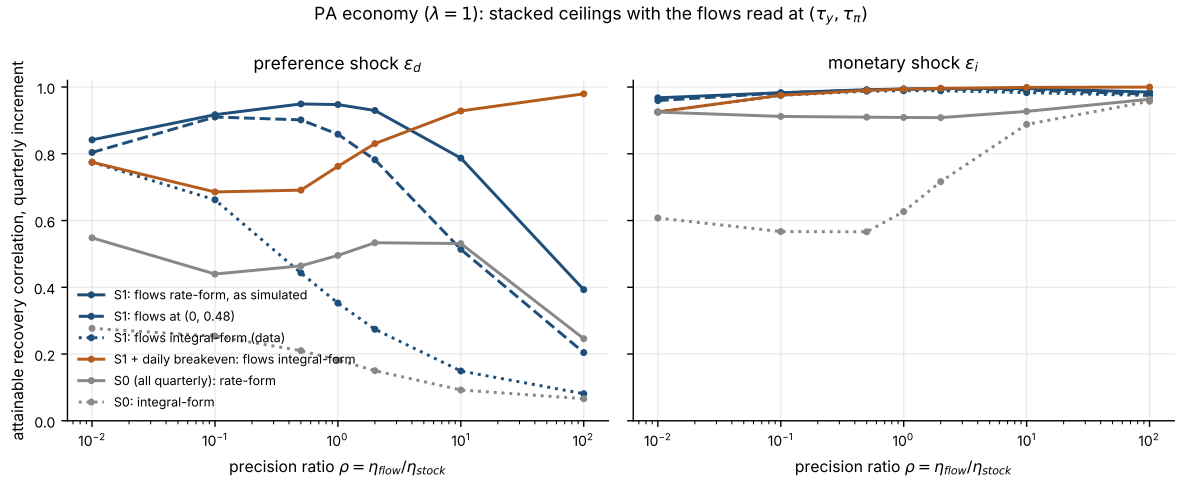


Figure 7: What the integral form costs the tracker (counterfactuals): attainable recovery of the quarterly structural increment in the partial-adjustment economy, $\lambda = 1$, as a function of the precision ratio $\rho = \eta_{\text{flow}}/\eta_{\text{stock}}$, with the two flows read at the pass-through scales (τ_y, τ_π). Left panel the preference shock, right panel the monetary shock. Ceilings are the stacked closed forms of Appendix B extended to the τ-kernel and to the quarterly target through the within-quarter innovation sequence; at $(1, 1)$ the S0 column reproduces Table 7. Blue: design S1 (output gap quarterly, inflation monthly, policy rate daily) with the flows rate-form as simulated (solid), at $(0, 0.48)$, the core PCE bound of the release alone (dashed), and integral-form as the data assign them (dotted). Orange: S1 with a daily 5-year breakeven added, flows integral-form. Grey: the uniform quarterly design S0 under the rate-form and integral-form readings. The monetary shock is recovered through the daily policy rate whichever way the flows are read; the preference shock loses most of its within-quarter recovery when the flow arm closes, unless a point-sampled level that loads on it is observed.

expected inflation, restores the identification in the mixed-frequency regime under integral form in the right corner, where the stocks are at least as precise as flows (which is, in practice, true).

8.3 Summary and Recommendations

In this section, I developed and showed the procedure to test for the categories of primitives this paper proposed. It is essential, since incorrect representation of flow variables leads to spurious results in structural filtering as shown via simulation in Section 5. This study testet GDP and inflation measures, and revealed that a chosen set of official US flow releases all reject the rate form primitive. The one open case left unanswered by the single-release test became settled via introducing other frequent series, or eventually model information. I showed variables that are published in both measurement ways, which ended up in the two different primitive category. What the simulation could not reveal, the empirical study did: the US flow-type releases do not include the pass-through content τ, in turn, the information can be recovered from adding point-in-time level data as the breakeven series. All this is confirmed without strictly introducing any economic or model structure, hence, the results remain robust and publishable.

The simulation study varied the flow primitive and left the stock fixed. The added

value of the rate form primitive lies in the pass-through content, which this empirical study recovered from adding an extra daily series, which loads on inflation when the diagnostics suggested no direct content within the flow variables. Thereby, the increment of the added stock recovered the pass-through content an eventual rate-form primitive would offer. The results and take-aways from the simulation remain valid and quantify exactly the gap between pass-through presence vs absence, merely by a different approach. The empirical study further confirmed the need for *ex ante* primitive testing and assignment, and developed a procedure for transparently doing it, entirely model-free if needed.

9 Conclusion

How should a macroeconomic model whose grid is quarterly read data that arrive daily? The answer is: on the data's own calendar, and under the correct measurement primitive. Establishing what separates those two conditions is the content of this paper. Exact discretisation makes parameter estimation invariant to the sampling frequency, while real-time filtering exhibits no such invariance, and the divergence is localised by continuous-time calculus entirely in the representation of flow variables. It suggests a measured flow is the window increment of a cumulative observable accruing at the equilibrium intensity, it inherits the state's Brownian increment by linearity—the pass-through—and the resulting flow–state cross-covariance is a structural feature of rate-form data, not a filtering option. Every difference among the filters studied here is a difference in how that single object is treated as data arrive, derived in closed-form.

The contributions reflect what calculus revealed, while empirical application confirms its validity. First, the paper casts three filters—the correlated-noise Kalman–Bucy filter, the exact discrete filter of the integral form, and the naive Euler scheme—as one taxonomy, ordered by how much of the pass-through each retains and at what temporal resolution. The general machinery is stated once, in the companion note (Rátfai, 2026). Second, four closed-form factors and a recovery-correlation law is derived for every comparison: the temporal-aggregation factor vanishes under finer sampling and is removable; the pass-through penalty is free of the interval and is not; the structural-penalty pair prices the Euler shortcut under each primitive—not even directionally alike; the level-resolution factor establishes the stock channel's near-immunity and the law makes the recovery ranking a corollary of the tracking ranking so that improving tracking accuracy would also improve recovery. Third, the mixed-frequency evidence carries the paper's practical message: daily observation of a single financial stock cuts the tracking error of a never-directly-observed state nearly thirtyfold and rescues the weakly identified shock of each configuration. The identifying information arrives within the quarter as a wave whose closed-form per-position ceilings the filter attains and densification under the wrong primitive amplifies the misspecification which cannot be repaired—the representation must be correctly assigned *before* the grid is refined. Fourth, the structural commitment the framework runs on is kept testable by the tests of Section 6: moment restrictions that assign the primitive release by release, innovation and recovered-shock diagnostics, parameter stability across sampling frequencies, and a validation agenda. Throughout, proximity to stochastic singularity—detrimental for likelihood-based estimation—is favourable for the online recovery, an asymmetry the closed forms make exact.

Fifth, the assignment rule is carried out on US real-time data (Section 8), and the results are addressed for the macro policy and macro-finance literature. The primitive

of a release can be read off the release itself, free of loadings and volatilities, and the reading rejects the rate form for every official flow release in the panel tested within this paper. The one open case, core PCE, is settled at the integral form once a daily yield or breakeven is read alongside it (Proposition F.1 says why a single release cannot decide). The within-window structural information of the US calendar therefore enters through its point-sampled levels, namely frequent asset price data, and not its flows: aggregating those levels to the macro grid moves them straight into the integral-form class and postpones half of the window's content a month late. While reading them daily, under the correlated-noise stacked representation recovers the shock on the day and rescues identification without imposing zero window-start error as in standard high-frequency identification ((Bauer and Swanson, 2023)). In the partial-adjustment economy the closed flow arm does not affect the monetary shock recovery substantially, since it is carried by the daily stock, but its effect is large for the preference shock – reduces from 0.95 to 0.35 at balanced precision, unless a daily level that loads on it is in the information set (addition of a second level-type series in the empirical study). This is a simulation-based, strong evidence that financial series can inform macroeconomic policy in real time, given the correct representation under which they are filtered. The conditions are testable robustly and model-free.

To address some limitations of the paper, the eocnomic systems used here are deliberately small-scale. The factors are exact for a scalar latent state driven by a single structural shock, the two-shock partial-adjustment economy shows the taxonomy survives where the observables must separate increments (two structural shocks). The closed-form factor expressions for the medium-scale economy left to future work. The conditions such a system must satisfy recorded are in Appendix A and are limited to linear state-space, but are applicable to any scale. The release-level test rests on a single persistent state behind each release, which headline inflation fails but e. g. an energy state would pass. Furthermore, the structural parameters are calibrated, not estimated: the empirical section estimates the measurement primitive with the loadings fixed where the structure enters, and the joint estimation of structural parameters within the same framework together with the full real-time deployment of the companion paper (Parra-Alvarez et al., 2026), are the subsequent steps.

Appendix

A The Identification Conditions in the Two Simulation Economies

The general conditions under which a rate-form system delivers the online objects (14)—observability of (A, C), the order condition $m \geq q$, fundamentalness of the instantaneous shock-to-observable map CB, shock separation through $\mathrm{cond}(CB)$, and well-posedness of the correlated-noise reduction—are stated as verifiable criteria, with proofs, in the companion note (Rátfai, 2026). This appendix records their verification in the two simulation economies.

The feedback-rule model. The scalar economy has $n = 1$ latent state driven by $q = 1$ demand shock, observed through output gap and inflation (flows) and the interest rate (stock). Each criterion holds unambiguously. Observability holds through any channel with nonzero loading. The order condition holds with margin, $m = 3 \geq q = 1$. Fundamentalness holds: with $q = 1$ the map CB is a nonzero vector of full column rank, and the single shock is recovered from the flow innovation through the scalar cross-covariance $(\Sigma_{xy})_j = \sigma_d^2 C_j$. Shock separation is vacuous, there being a single shock. Crucially, well-posedness holds for *any* positive flow noise: the effective process-noise variance rate $\bar{Q}^f$ of (25) is strictly positive throughout, vanishing only in the noiseless limit, so the singular point is excluded by the filter's own well-posedness condition and the joint tracking of state and shock holds across the entire precision grid.

The partial-adjustment model. With $n = q = 2$, the two requirements that are automatic in the scalar economy become substantive. First, shock separation: $\mathrm{rank}(\mathcal{O}) = 2$ and $\tilde{C}B$ has full column rank at the calibration, but $\det C^f = -0.0120$ is small, so the flow loadings are nearly collinear and the separation of the two increments rests on the stock channel—most powerfully under dense sampling of the rate, as Section 5.6 documents. Second, well-posedness: positive semidefiniteness of the correlated-noise reduction $\bar{Q} = \Omega - SR^{-1}S'$ holds automatically in any dimension by the shared-loading Schur-complement structure $S = \Omega\,[(C^f)', 0]$, proved in the companion note, so no positivity check is ever needed; what must be verified is the existence of the stabilizing solution of the filter Riccati equation, which is confirmed numerically at every point of the (λ, ρ) grid in Section 5.5.

B Attainable Recovery Ceilings: Application to the Simulation Economies

The ceilings quoted throughout Section 5 are the application, to the simulation economies, of the attainable-recovery result stated and proved in general form in the companion note (Rátfai, 2026): the correlation between any online linear estimate $m'\nu_k$ and the true increment is maximised at $m^\star \propto S_\nu^{-1} c$—the operators (24) are exactly $m^\star$ on their information sets (Proposition 4.5)—and, evaluated at the exact representation's steady

posterior variance P through the moments (22)–(23), the maxima take the closed forms:

$$\mathrm{corr}^2_{\max}(\text{flow-only}) = \frac{(c^f)'\, S_{ff}^{-1}\, c^f}{h}, \qquad \mathrm{corr}^2_{\max}(\text{stock-only}) = \frac{c_s^2}{S_{ss}\, h},$$
$$\mathrm{corr}^2_{\max}(\text{stacked}) = \frac{c'\, S_{\nu}^{-1}\, c}{h}, \tag{52}$$

all computable in closed form from the window moments and the steady posterior variance P of the exact representation (Appendix C). Three properties, documented in Section 5.3.1, follow. *Attainment*: the operators (24) sit on their ceilings at every point of the precision grid—the map is exact for its information set, and the correlation level measures the information content of that set, not method error. *Complementarity*: the marginal contributions of the two channels—stacked minus stock-only, and stacked minus flow-only—are simultaneously first-order only in the balanced-precision region $\rho \in [0.5, 2]$, exactly where the overall level troughs; at either precision extreme one channel spans the other. *Attribution*: filter choice matters only on the flow arm, where the innovations differ by their pass-through content; the stock arm is filter-invariant, an information-set result available to any filter once ν^s enters the operator.

C Simulation Validation against EDT$^+$, the Exact Discrete Filter of the Data-Generating Process

This appendix validates, on the simulated economies, that the continuous-time filter of Section 3.2 attains the minimum-MSE ceiling of its information set—the validation layer the general framework calls for but cannot itself supply, since the companion note contains no simulation. The instrument is EDT$^+$, the exact discrete filter of the rate-form process itself, constructed in general matrix form in (Rátfai, 2026): it shares with the EDT of Section 3.2 the augmented state, the measurement equation (19), the accumulator resets, the steady-state gain construction, the measurement-noise convention, and the Van Loan machinery, and differs in exactly one respect—it discretises the pair $(\mathcal{A}, \mathcal{B})$ of (11) rather than its integral-form restriction, so the Van Loan computation returns the EDT blocks *plus* the second moments of the pass-through, which in the scalar economy evaluate to:

$$\Sigma^+_{\eta^x\eta^f} \;=\; \Sigma_{\eta^x\eta^f} + \sigma_d^2\, \Gamma_h, \qquad \Sigma^+_{\eta^f} \;=\; \Sigma_{\eta^f} + \sigma_d^2\, h + 2\,\sigma_d^2\, A^{-1}(\Gamma_h - h)\,, \tag{53}$$

with $\mathbf{Q}^+$ assembled from these blocks exactly as in (18), in the same rate units. No new machinery is required: the pass-through enters as bookkeeping in the innovation covariance, not as an algorithmic change.

Role in the validation. By optimality of the conditional expectation, EDT$^+$ is the minimum-MSE filter on the release information set of its calendar, and its retained cross-covariance is the instantaneous Σ_{xy} at full strength but boundary-anchored—an $O(h)$ resolution deficit sized by $\Phi(h)$ (Section 4.1), so this, not the integral-form EDT, is the representation for which the collapse onto the continuous-time filter holds; both properties are established in the companion note. What the simulations add is the validated decomposition. The CT–EDT$^+$ MSE gap is numerically indistinguishable from zero already at the quarterly interval, so the pairwise comparisons decompose cleanly: the CT filter versus EDT$^+$ isolates the (negligible) resolution effect; EDT$^+$ versus the EDT isolates the pass-through content; the EDT versus Euler isolates the cost of discarding

the flow integral altogether. The validation comparisons are made at steady state: both the CT filter and EDT$^+$ run constant gains from different initialisations, so the transient is not ordered by optimality; the MMSE property holds in the steady state, which is what the simulations report.

Online structural-shock recovery. EDT$^+$ applies the operators (24) to its own innovations. Since its flow innovation carries the pass-through at full (boundary-anchored) strength, the recovery essentially coincides with the continuous-time recovery, differing only at the $O(h)$ resolution order.

Table 10: Implementation Validation: the ODE-Realised CT Filter and EDT against EDT$^+$

ρ	**EDT$^+$ steady MSE**	**CT excess (%)**	p	**EDT excess (%)**	p
0.01	1.203×10^{-5}	+0.02	0.33	+11,542.90	< 0.001
0.10	8.119×10^{-5}	−0.01	0.40	+126.33	< 0.001
0.50	2.113×10^{-4}	+0.20	< 0.001	+21.59	< 0.001
1.00	2.136×10^{-4}	+0.81	< 0.001	+13.48	< 0.001
2.00	1.705×10^{-4}	+1.30	< 0.001	+6.33	< 0.001
10.00	6.103×10^{-5}	+0.57	< 0.001	+0.65	< 0.001
100.00	7.849×10^{-6}	+0.02	< 0.001	+0.00	0.64

Notes: Steady-state comparison over $N = 1000$ replications on identical exactly-simulated draws; excess columns report the paired excess MSE of the named filter in percent of the EDT$^+$ MSE, with two-sided paired p-values. All filters run constant steady-state gains with their own initialisations, so the transient is not ordered by optimality and is excluded. EDT$^+$, the exact discrete filter of the data-generating process, cannot be beaten in steady state; the near-zero CT column therefore validates that the ODE realisation—a Riccati solve plus a linear-ODE integration plus an endpoint update—attains the ceiling, while the EDT column is the cost of the integral-form flow model, persistent under fine sampling (Section 5.3.3). Both columns have exact theoretical counterparts: the EDT$^+$ steady MSE is the fixed point of the augmented DARE and the EDT excess the corresponding mismatched-gain Lyapunov value—1.227×10^{-5} and +11,384 percent at $\rho = 0.01$—and the estimates attain them within Monte Carlo error at every ρ. Empirical steady MSE levels sit up to two percent below the infinite-horizon fixed points, the residual of upward convergence within the finite steady window; the residual cancels in every reported ratio. The CT deficit is the constant-rate assimilation cost of the scheme; it peaks mid-grid and vanishes at both precision corners.

Implementation validation: the ODE realisation against EDT$^+$. For the practitioner, the recommended implementation of the continuous-time filter requires none of the machinery of this appendix. The recipe is three steps: solve the correlated-noise CARE once for the steady flow gain (a scalar quadratic here; any standard Riccati routine, or integration of the Riccati ODE to convergence, in general); over each window, integrate the *linear* filter ODE $d\hat{x}/dt = A\hat{x} + K^f_\infty(r_k - C^f\hat{x})$ with any standard solver; apply the discrete stock update at the endpoint. No state augmentation, no accumulators, and no within-window moment blocks appear in the running filter—those belong to the verification layer, which this appendix supplies once so that the practitioner need not build it. Table 10 records the validation: the ODE realisation, run paired against EDT$^+$ on identical exactly-simulated draws, attains the minimum-MSE ceiling of the

quarterly information set to within 1.3 percent everywhere on the precision grid and to within 0.1 percent at both corners. The mid-grid deficits are statistically resolved at $N = 1000$ yet economically bounded; their sign is uniformly the one the minimum-MSE property requires, and their source is the constant-rate assimilation of the release across the window—a representation cost of the scheme, not solver error. The solver itself is safely inside its tolerance: the window ODE is linear, so its closed-form solution is the limit every solver converges to, and a deliberately coarse fixed-step integration quantifies the approach—two Euler steps per quarter inflate the MSE by 16 percent, five steps by 2.0 percent, twenty steps by under 0.1 percent, and beyond one hundred steps the result is indistinguishable from the closed form, where adaptive solvers at default tolerances already sit. This yields a self-check requiring no benchmark at all: refine the step until the MSE stabilises.

D Derivations: the Four Factors and the Recovery Law

This appendix derives the closed forms and expansions stated in Section 4. Throughout, $q \equiv Ah$ denotes the persistence-scaled interval, and the window's three primitive stochastic integrals are

$$\eta_k^x = \int_0^h e^{A(h-s)}\sigma_d\, dW_s^x, \qquad u_k = \int_0^h \Gamma(h-s)\,\sigma_d\, dW_s^x, \qquad \Delta W_k = \int_0^h dW_s^x, \tag{54}$$

the state innovation, the path news of (6), and the structural increment. The level-form flow innovation of the exact discrete representation coincides with the path news: writing it as the iterated integral $\eta_k^f = \int_0^h \big[\int_0^s e^{A(s-u)}\sigma_d\, dW_u^x\big] ds$ and exchanging the order of integration (stochastic Fubini), $\eta_k^f = \int_0^h \big[\int_u^h e^{A(s-u)} ds\big]\sigma_d\, dW_u^x = \int_0^h \Gamma(h-u)\,\sigma_d\, dW_u^x = u_k$, since $\int_u^h e^{A(s-u)} ds = A^{-1}(e^{A(h-u)} - 1) = \Gamma(h-u)$. All second moments below therefore follow from the Itô isometry applied to (54).

D.1 Window Moments and Expansions

The isometry $\mathbb{E}[dW_u^x\, dW_r^x] = \delta(u-r)\, du$ collapses every double integral to a single one over the window. Substituting $\tau = h - s$,

$$\begin{aligned} \Sigma_{\eta^x} &= \sigma_d^2 \int_0^h e^{2A\tau} d\tau = \sigma_d^2\, \frac{e^{2Ah}-1}{2A}, \\ \Sigma_{\eta^f} &= \sigma_d^2 \int_0^h \Gamma(\tau)^2\, d\tau = \frac{\sigma_d^2}{A^2}\left[\frac{e^{2Ah}-1}{2A} - 2\Gamma_h + h\right], \end{aligned} \tag{55}$$

$$\begin{aligned} \Sigma_{\eta^x\eta^f} &= \sigma_d^2 \int_0^h e^{A\tau}\Gamma(\tau)\, d\tau = \frac{\sigma_d^2}{A}\left[\frac{e^{2Ah}-1}{2A} - \Gamma_h\right], \\ \operatorname{cov}(\eta_k^x, \sigma_d \Delta W_k) &= \sigma_d^2 \int_0^h e^{A\tau} d\tau = \sigma_d^2\, \Gamma_h, \end{aligned} \tag{56}$$

$$\operatorname{cov}(u_k, \sigma_d \Delta W_k) = \sigma_d^2 \int_0^h \Gamma(\tau)\, d\tau = \sigma_d^2\, \frac{\Gamma_h - h}{A}. \tag{57}$$

The rate-form flow disturbance is $u_k + \sigma_d \Delta W_k$, whose variance and whose covariance with the state innovation assemble from these pieces:

$$v^+ = \Sigma_{\eta^f} + \sigma_d^2 h + 2\sigma_d^2\, \frac{\Gamma_h - h}{A}, \qquad x^+ = \Sigma_{\eta^x \eta^f} + \sigma_d^2\, \Gamma_h, \tag{58}$$

which are exactly the moments v^+, x^+ of the innovation covariance (23) and the EDT$^+$ blocks (53): the three constructions—recovery operator, verification filter, and the factors below—share one set of window moments. The expansions used repeatedly are, from $e^{Ah} = \sum_n (Ah)^n/n!$,

$$\Gamma_h = h + \tfrac{Ah^2}{2} + \tfrac{A^2h^3}{6} + \tfrac{A^3h^4}{24} + O(h^5), \qquad \frac{e^{2Ah} - 1}{2A} = h + Ah^2 + \tfrac{2}{3}A^2h^3 + \tfrac{1}{3}A^3h^4 + O(h^5), \tag{59}$$

whence

$$\Sigma_{\eta^f} = \sigma_d^2 \left[\tfrac{h^3}{3} + \tfrac{Ah^4}{4} + \tfrac{7A^2h^5}{60} + O(h^6) \right], \qquad \Sigma_{\eta^x \eta^f} = \sigma_d^2 \left[\tfrac{h^2}{2} + \tfrac{Ah^3}{2} + \tfrac{7A^2h^4}{24} + O(h^5) \right], \tag{60}$$

$$v^+ = \sigma_d^2 \left[h + h^2 + \left(\tfrac{1}{3} + \tfrac{A}{3} \right) h^3 + O(h^4) \right]. \tag{61}$$

The mixed powers in (61) are not an accident of bookkeeping: the pass-through variance $\sigma_d^2 h$ and the path-news variance $\sigma_d^2 h^3/3$ scale differently with the window, because the rate-form primitive (4) attaches the diffusion loading $C_j \sigma_d$ in the units in which the model is specified. The quantities $\Phi(h)/h$, $\Sigma_{\eta^f}/(\sigma_d^2 h^3)$, and $1 - \Lambda(h)$ are pure functions of $q = Ah$ and hence unit-free; Ψ_{II} is not, and all its numerical statements are in the annual units of the calibration — the scale τ of Remark 2.1 in explicit form.

D.2 The Temporal-Aggregation Factor

The identity (27) is the first member of (56): $\operatorname{cov}(\eta_k^x, \eta_k^f) = \sigma_d^2 \int_0^h e^{A(h-u)} \Gamma(h-u)\, du$, which reproduces the corresponding Van Loan block to machine precision and, divided by $\sigma_d^2 h$, gives the closed form of Proposition 4.1,

$$\Phi(h) = \frac{1}{Ah} \left[\frac{e^{2Ah} - 1}{2A} - \Gamma_h \right]. \tag{62}$$

Subtracting the two series in (59) and dividing by Ah yields $\Phi(h) = \frac{1}{2}h + \frac{1}{2}Ah^2 + \frac{7}{24}A^2h^3 + O(h^4)$, the expansion of (28) and (38). The factor is not globally monotone: under the calibration it attains a maximum of 0.483 at $h \approx 2.98$ years and thereafter decays as $1/(2A^2h)$, the endpoint discounting $e^{A(h-u)}$ eventually dominating the accumulation; on the calibrated range $h \le 2$ it is increasing, which is the statement of Proposition 4.1.

D.3 The Structural-Penalty Pair

Integral form. Both Σ_{I}^f and $\Sigma_{\mathrm{Euler}}^f$ carry the loading C_y^2 and the rate normalisation h^{-2}, so the penalty reduces to a ratio of level-unit variances, $\Psi_{\mathrm{I}}(h) = \sigma_d^2 h^3/\Sigma_{\eta^f} - 1$. From (55), $\Sigma_{\eta^f}/(\sigma_d^2 h^3) = \left[(e^{2q} - 1)/(2q) - 2(e^q - 1)/q + 1\right]/q^2$ is a pure function of $q = Ah$, with

series $\frac{1}{3}+\frac{q}{4}+\frac{7q^2}{60}+O(q^3)$ by (60). Inverting the series,

$$\frac{\sigma_d^2 h^3}{\Sigma_{\eta^f}} = \frac{3}{1+\frac{3}{4}q+\frac{7}{20}q^2+O(q^3)} = 3\left[1-\tfrac{3}{4}q+\left(\tfrac{9}{16}-\tfrac{7}{20}\right)q^2\right]+O(q^3) = 3-\tfrac{9}{4}\,q+\tfrac{51}{80}\,q^2+O(q^3), \tag{63}$$

using $\frac{9}{16}-\frac{7}{20}=\frac{45-28}{80}=\frac{17}{80}$ and $3\cdot\frac{17}{80}=\frac{51}{80}$; subtracting one gives (36).

Rate form. The rate-form flow disturbance is $u_k+\sigma_d\Delta W_k$ with variance v^+ of (58), so $\Psi_{\mathrm{II}}(h)=\sigma_d^2h^3/v^+-1$. By (61), the fraction of the true rate-form flow-disturbance variance that the Euler filter models is

$$f(h) \;\equiv\; \frac{\sigma_d^2 h^3}{v^+} = \frac{h^2}{1+h+\left(\frac{1}{3}+\frac{A}{3}\right)h^2+O(h^3)} = h^2-h^3+\left(\tfrac{2}{3}-\tfrac{A}{3}\right)h^4+O(h^5), \tag{64}$$

whence $\Psi_{\mathrm{II}}=-1+h^2-h^3+O(h^4)$, with the persistence A absent through third order: the omission penalty is set by the scale mismatch between the pass-through, of order σ_d^2/h in rate units, and Euler's assignment σ_d^2h, not by the within-interval curvature of the path. At the quarterly interval $f(\frac{1}{4})=0.0496$. The exact expression changes sign where $\sigma_d^2h^3=v^+$: Euler's assigned variance grows as h^3 in level units while v^+ grows essentially linearly, so at sufficiently coarse intervals the understatement turns into an overstatement; under the calibration the crossing sits at $h^*=1.722$ years, far coarser than any interval of applied interest, and $\Psi_{\mathrm{II}}(2)=+0.219$ as Table 3 records.

D.4 The Level-Resolution Factor

With $\Lambda(h)=\sigma_d\Gamma_h/\sqrt{\Sigma_{\eta^x}h}$, write $N(q)=\Gamma_h/h=1+\frac{q}{2}+\frac{q^2}{6}+\frac{q^3}{24}+O(q^4)$ and $D(q)^2=\Sigma_{\eta^x}/(\sigma_d^2h)=1+q+\frac{2}{3}q^2+\frac{1}{3}q^3+O(q^4)$, so that $\Lambda=ND^{-1}$ is a pure function of q. Expanding $D^{-1}=1-\frac{q}{2}+\frac{q^2}{24}+\frac{q^3}{48}+O(q^4)$ and multiplying,

$$\Lambda(q)=1+\left(\tfrac{1}{2}-\tfrac{1}{2}\right)q+\left(\tfrac{1}{6}-\tfrac{1}{4}+\tfrac{1}{24}\right)q^2+\left(\tfrac{1}{24}-\tfrac{1}{12}+\tfrac{1}{48}+\tfrac{1}{48}\right)q^3+O(q^4)=1-\frac{q^2}{24}+O(q^4)\,: \tag{65}$$

the linear and, less obviously, the *cubic* coefficients vanish identically, so the remainder in (39) is $O(h^4)$, one order better than the flow-side expansions. Numerically, $\left[(1-\Lambda)-q^2/24\right]/q^4$ is constant at -0.0033 across $h\in[0.125,1]$, confirming the vanishing cubic. The quadratic loss—against the $O(h)$ resolution effects of the flow side—is the algebraic content of the level collinearity (13): at $h=\frac{1}{4}$, $1-\Lambda=4.6\times10^{-4}$, at $h=\frac{1}{12}$ it is 7.4×10^{-5} (the ninefold fall quoted in Remark 4.4), and at a daily interval 1.4×10^{-8}.

D.5 The Pass-Through Penalty

Setup. Fix the flow-channel measurement intensity at its quarterly-calibrated value, $R_c=rI_2$ with $r=(\bar{h}\,\eta_f)^2$ and $\bar{h}=\frac{1}{4}$; the per-release variance at assimilation interval h is then r/h, so that refining h leaves the observation process—the information flow per unit time—unchanged and varies only the assimilation interval. In the continuous-record limit the flow block (8)–(9) reads $dz_t=C^f x_t\,dt+dv_t$ with $dv_t=C^f\sigma_d\,dW_t^x+dv_t^c$, so the total observation-noise variance rate and its correlation with the process noise are

$$R_{\mathrm{tot}}=\sigma_d^2\,C^f(C^f)'+R_c, \qquad \mathbb{E}\left[\sigma_d\,dW_t^x\,dv_t'\right]=\Sigma_{xy}'\,dt, \qquad \Sigma_{xy}=\sigma_d^2C^f. \tag{66}$$

The exact filter. The correlated-noise Kalman–Bucy filter (Kalman-Bucy, 1962; Anderson and Moore, 1979) carries the gain $K = (PC^{f\prime} + \Sigma'_{xy})R_{\text{tot}}^{-1}$ and the Riccati equation $0 = 2AP + \sigma_d^2 - (PC^{f\prime} + \Sigma'_{xy})R_{\text{tot}}^{-1}(C^f P + \Sigma_{xy})$. Since $\Sigma_{xy} = \sigma_d^2 C^f$, the three quadratic forms collapse onto $a_* = (C^f)'R_{\text{tot}}^{-1}C^f$: the cross term is $2P\sigma_d^2 a_*$ and the pure noise term $\sigma_d^4 a_*$, so the Riccati equation is the scalar quadratic $a_* P^2 - 2\bar{A}_* P - \bar{Q}_* = 0$ with $\bar{A}_* = A - \sigma_d^2 a_*$ and $\bar{Q}_* = \sigma_d^2 - \sigma_d^4 a_*$, whose stabilising (positive) root is (30).

The integral-form limit. At interval h the integral-form filter models the flow disturbance as the path news alone: its modelled disturbance variance rate is $\Sigma_{\eta^f}/h^2 = \sigma_d^2 h/3 + O(h^2)$ and its modelled state-flow cross-covariance is $\Sigma_{\eta^x\eta^f}/h = \sigma_d^2\,\Phi(h)$, both $O(h)$, while the modelled release-noise variance rate is fixed at R_c. The modelled system therefore converges, as $h \to 0$, to the *uncorrelated* pair $(A, \sigma_d^2; C^f, R_c)$, and the filter's gain to the uncorrelated-noise Kalman–Bucy gain $K_0 = P_0(C^f)'R_c^{-1}$ with P_0 the positive root of $a_0 P^2 - 2AP - \sigma_d^2 = 0$, $a_0 = (C^f)'R_c^{-1}C^f = \|C^f\|^2/r$. The convergence of the actual discrete filters is verified below.

The mismatched steady state. Running the constant gain K_0 against the rate-form truth gives the error dynamics $de_t = (A - K_0 C^f)\,e_t\,dt + \sigma_d\,dW_t^x - K_0\,dv_t$, whose stationary variance solves the Lyapunov equation

$$0 = 2(A - K_0 C^f)\,P_{\text{mis}} + \sigma_d^2 + K_0 R_{\text{tot}} K_0' - 2K_0\Sigma_{xy}. \tag{67}$$

With $R_c = rI_2$ the two gain forms reduce to scalars: $K_0 C^f = a_0 P_0$, $K_0\Sigma_{xy} = \sigma_d^2 a_0 P_0$, and $K_0 R_{\text{tot}} K_0' = (P_0^2/r^2)\big(\sigma_d^2\|C^f\|^4 + r\|C^f\|^2\big) = \sigma_d^2 a_0^2 P_0^2 + a_0 P_0^2$. The numerator of (67) assembles as $\sigma_d^2 - 2\sigma_d^2 a_0 P_0 + \sigma_d^2 a_0^2 P_0^2 + a_0 P_0^2 = \sigma_d^2(1 - a_0 P_0)^2 + a_0 P_0^2$, giving (31). The closed form solves (67) to machine precision at every grid point.

Inequality and the equality point. P_* is the minimum steady-state MSE over all causal filters of the observation process, and the constant-gain filter with K_0 is causal, so $\Pi = P_{\text{mis}}/P_* \geq 1$; equality holds exactly when the constant gain attains the optimum, i.e. when K_0 equals the optimal steady gain $K_* = (P_* C^{f\prime} + \Sigma'_{xy})R_{\text{tot}}^{-1} = (P_* + \sigma_d^2)\,(C^f)'R_{\text{tot}}^{-1}$.By the Sherman–Morrison identity, $(C^f)'R_{\text{tot}}^{-1}$ is proportional to $(C^f)'R_c^{-1}$ for *every* positive-definite R_c, so the two gains are always collinear and the filter is sensitive to the scalar effective gain alone; the symmetric design $R_c = rI_2$ is a convenience of the calibration, not a requirement. (I owe this observation, and the condition that follows, to a comment by Gerrit Meyerheim.) Equality of the gains requires $P_* = (1 + \sigma_d^2 a_0)P_0 - \sigma_d^2$, and substituting into the exact filter's Riccati equation, using the mismatched filter's own, leaves $-\sigma_d^2(2A + \sigma_d^2 a_0) = 0$: equality holds if and only if $a_0 = (C^f)'R_c^{-1}C^f = 2\rho_d/\sigma_d^2$. Since a_0 is strictly decreasing in the noise scale, the crossing is unique, interior, and explicit; under the calibration $\rho^* = 0.411$, at which the computed $\Pi(\rho^*) - 1$ is below 10^{-15}. On either side the penalty reopens quadratically in the gain error, which is why the tabulated $\Pi(0.5) = 1.004$ sits just above one.

Numerical convergence of the actual filters. To verify that the fine-sampling limits are the ones the proposition names, construct at each h the discrete rate-form truth from the window moments—state transition e^{Ah} with innovation variance Σ_{η^x}, observation $(C^f/h)\big(\Gamma_h x_{k-1} + u_k + \sigma_d\Delta W_k\big)$ plus release noise of variance $(r/h)I_2$—and evaluate two filters: the integral-form filter (modelled disturbance u_k alone, gain from its own Riccati recursion, true error from the mismatched discrete Lyapunov equation with the moments v^+, x^+) and the exact rate-form filter. Table 11 reports the true steady-state MSEs relative to their claimed limits: both ratios reach one, uniformly in ρ. The convergence is one-sided in the flow-precise corner: at $\rho = 0.01$ the integral-form filter's quarterly MSE is 0.886 of

its own limit—the retained fraction $\Phi(h)$ of the cross-covariance still helps at finite h—so refining the interval strictly *worsens* that filter until it settles at P_{mis}, the one-sidedness noted in Section 4.2. Everywhere else on the grid the quarterly interval is already within a fraction of a percent of the limit.

Table 11: Fine-sampling convergence of the actual filters to the limits of Proposition 4.2: true steady-state MSE relative to the claimed limit (P_{mis} for the integral-form filter, P_* for the exact rate-form filter), flow block, variance rates fixed at quarterly-calibrated values.

	Integral-form filter, MSE/P_{mis}				Exact rate-form filter, MSE/P_*			
h	ρ=0.01	0.1	1	10	ρ=0.01	0.1	1	10
1/4	0.886	1.000	1.002	1.000	1.360	1.033	1.002	1.000
1/12	0.986	1.000	1.000	1.000	1.040	1.004	1.000	1.000
1/50	0.999	1.000	1.000	1.000	1.002	1.000	1.000	1.000
1/365	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

D.6 The Recovery-Correlation Law

The formula. The window-start error $\tilde{e}_{k-1}$ is a function of data and shocks through t_{k-1} and the release noises are independent of W^x, so both are independent of the window's shocks $(u_k, \Delta W_k, \eta_k^x)$. From the innovations (21), the covariance of the stacked innovation with the target is filter-free—$\text{cov}(\nu_k, \Delta W_k) = c$ with c exactly as in (22), using (57) and (56) for the flow and stock entries respectively—and the variance decomposes along the same independence: the $\tilde{e}_{k-1}$ terms load on the innovation through the vector $w = \left[C_y\Gamma_h/h,\ C_\pi\Gamma_h/h,\ C_s e^{Ah}\right]'$, contributing $P_F\, ww'$, while the window-shock and release-noise terms contribute the matrix S_0 of (23) evaluated at $P = 0$. Hence $\text{var}(m'\nu_k) = m'(S_0 + P_F ww')m$ and, with $\text{var}(\Delta W_k) = h$, the correlation formula (40) follows. The maximisation is the optimal-read lemma of the companion note applied at $S(P_F)$.

Monotonicity. The Sherman–Morrison identity applied to $S(P_F) = S_0 + P_F\, ww'$ gives

$$c'S(P_F)^{-1}c = c'S_0^{-1}c - \frac{P_F\,(c'S_0^{-1}w)^2}{1 + P_F\, w'S_0^{-1}w}, \tag{68}$$

which is (41); its derivative in P_F is $-(c'S_0^{-1}w)^2/(1 + P_F\, w'S_0^{-1}w)^2 \le 0$, strictly negative if and only if $c'S_0^{-1}w \neq 0$. Under the calibration $c'S_0^{-1}w$ is positive at every point of the precision grid, ranging from 5.0 (at $\rho = 0.5$) to 30.4 (at $\rho = 100$), so the decrease is strict throughout. For a common fixed map m, $\text{var}(m'\nu_k) = m'S_0m + P_F(m'w)^2$ is increasing in P_F while $m'c$ is constant, so the correlation is decreasing, strictly whenever $m'w \neq 0$. The ranking statement of Corollary 4.6 follows in both regimes, and evaluating (68) at the exact filter's P reproduces the stacked ceiling of (52); the flow-only and stock-only ceilings are the same evaluation on the corresponding sub-blocks.

Filter invariance. Every entry of c is a covariance between ΔW_k and a data-side component of the innovation—pass-through, path news, or the state innovation embedded in the sampled level—and the filter enters ν_k only through $\tilde{e}_{k-1}$, which is uncorrelated with ΔW_k. Hence c is a property of the data-generating process alone, which is Corollary 4.7.

Numerical verification. A direct Monte Carlo of the window (2×10^5 windows of 250 substeps at $\rho = 1$, $h = \frac{1}{4}$, with $\tilde{e}_{k-1}$ drawn at an arbitrary P_F) reproduces c and

$S(P_F)$ within simulation error and gives $\mathrm{corr}(m'\nu_k, \Delta W_k) = 0.6035$ under the optimal map against the closed-form 0.6042, with (68) agreeing with the direct inverse to machine precision.

E Additional Simulation Tables and Figures

This appendix collects the complete feedback-rule simulation record summarised in Section 5.3: the uniform-design tracking and recovery results (Tables 12 and 13, Figures 8 and 9), the mixed-design results and the frequency effect (Tables 14, 15, and 16, Figures 10 and 11), the benchmark comparison against the naive Euler filter (Table 17), the per-filter flow-only recoveries under mixed frequency (Table 18), the feedback-rule identification wave (Figure 12, Table 19), and the partial-adjustment design figures not shown in the text (Figures 13, 14, 15, 16, and 17).

Table 12: Uniform Quarterly Frequency: Steady-State Filtering MSE by Precision Ratio

	MSE			**MSE ratio**	
ρ	**CT**	**EDT**	**Euler**	**CT/EDT**	**CT/Euler**
0.01	0.000012	0.001400	0.001379	0.0086	0.0087
0.10	0.000081	0.000184	0.000238	0.4418	0.3415
0.50	0.000212	0.000257	0.000245	0.8241	0.8631
1.00	0.000215	0.000242	0.000236	0.8884	0.9122
2.00	0.000173	0.000181	0.000179	0.9528	0.9633
10.00	0.000061	0.000061	0.000061	0.9992	0.9990
100.00	0.000008	0.000008	0.000008	1.0002	0.9998

Notes: Steady-state mean squared error of the filtered state against the true simulated state at quarter ends, over $N = 1000$ replications of $T = 150$ quarters with a 50-quarter burn-in excluded. All reported MSEs in this and the following tables are steady-state: the initialization transient is excluded because it depends on the arbitrary starting covariance and is not ordered by optimality. The two right-hand columns report the MSE ratio of the continuous-time filter to each alternative; values below one favour the CT filter. CT is the correlated-noise Kalman–Bucy filter of Section 3; EDT is the Christensen–Neri–Parra-Alvarez exact discrete filter of the integral-form flow model (Christensen et al., 2024); Euler is the naive discretisation. Paired t-tests of CT against EDT reject equality at $p < 0.001$ for every $\rho \leq 2$ and are insignificant beyond; full test statistics accompany Table 17.

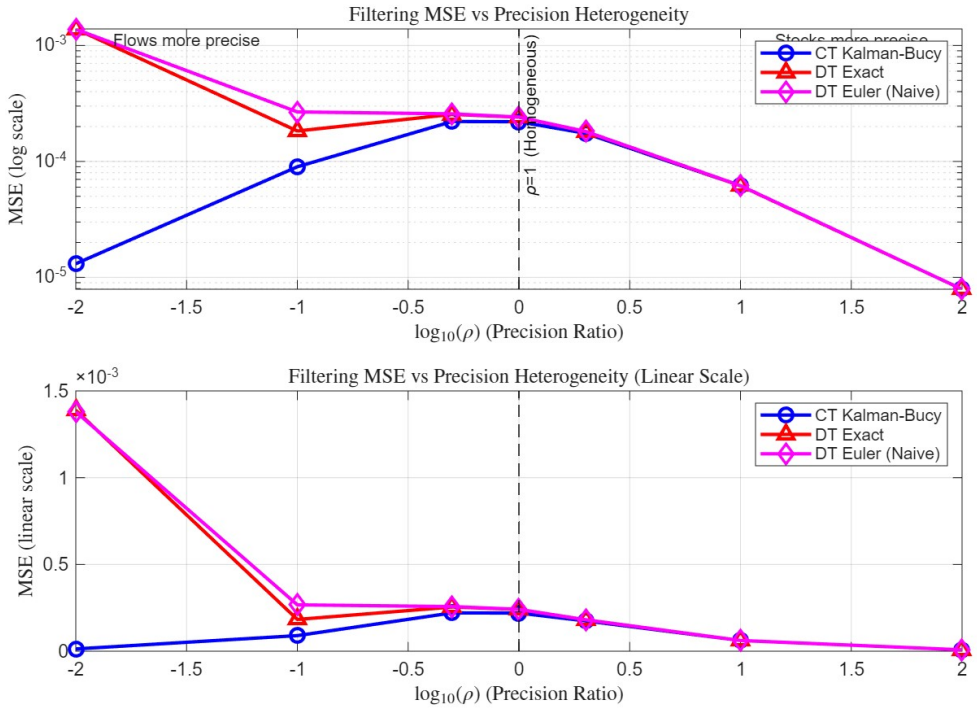


Figure 8: Uniform Quarterly Frequency: filtering MSE versus precision heterogeneity, log scale (upper panel) and linear scale (lower panel). The continuous-time filter separates from the exact discrete filter where flows are the precise channel ($\rho < 1$), the region where the pass-through carries the most information, and the two converge to numerical equivalence as the precise channel shifts to the point-in-time stock.

Table 13: Uniform Quarterly: Online Structural-Shock Recovery—Flow-Only and Stacked Maps against Their Ceilings

	Ceilings		**CT**		**EDT**		**Euler**	
ρ	flow	stacked	flow	stacked	flow	stacked	flow	stacked
0.01	0.9857	0.9857	0.9856	0.9856	0.8753	0.8753	0.8767	0.8767
0.10	0.8945	0.8947	0.8938	0.8940	0.8866	0.8866	0.8828	0.8827
0.50	0.6750	0.6916	0.6737	0.6902	0.6724	0.6876	0.6727	0.6883
1.00	0.5469	0.6123	0.5457	0.6110	0.5452	0.6084	0.5454	0.6090
2.00	0.4214	0.6001	0.4204	0.5991	0.4203	0.5973	0.4203	0.5977
10.00	0.2042	0.7823	0.2036	0.7824	0.2036	0.7823	0.2036	0.7823
100.00	0.0658	0.9655	0.0655	0.9656	0.0655	0.9656	0.0655	0.9656

Notes: Correlation between the online-recovered structural increment and the true simulated increment $\int_k dW$, pooled over the steady-state window across replications. The ceiling columns are the closed-form attainable correlations of Appendix B for the flow-only and stacked information sets—theoretical quantities computed from the exact representation's steady innovation moments, requiring no filter run. Each filter applies the identical fixed maps to its own innovations, formed by the exact prediction operators of (21); filters therefore differ only through the quality of the window-start estimate. The recovery is established within the simulation, where the true increment is known; recovery from observed data is developed in (Parra-Alvarez et al., 2026).

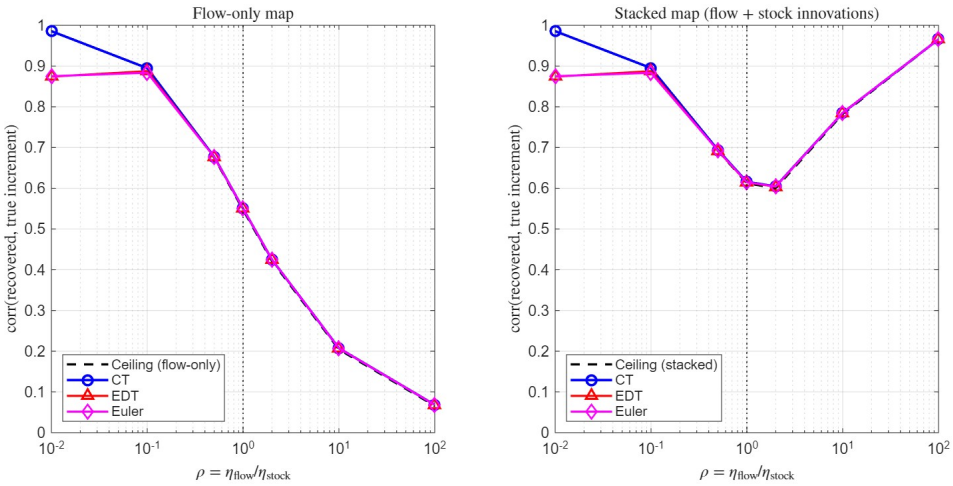


Figure 9: Uniform Quarterly: structural-shock recovery under the flow-only map (left) and the stacked map (right), each against its closed-form ceiling (dashed). The flow-only recovery declines monotonically in ρ as the pass-through channel loses precision; the stacked recovery is U-shaped, its right arm restored by the stock innovation through the level channel. Filters separate only on the left arm.

Table 14: Mixed Frequency (S1): Steady-State Filtering MSE by Precision Ratio

ρ	**CT**	**EDT**	**Euler**	**CT vs. EDT (%)**
0.01	1.6133×10^{-6}	5.2969×10^{-3}	4.6348×10^{-3}	+328227.97
0.10	1.2669×10^{-5}	9.1042×10^{-4}	1.0239×10^{-3}	+7086.35
0.50	4.2624×10^{-5}	7.7922×10^{-5}	7.9601×10^{-5}	+82.81
1.00	5.0627×10^{-5}	7.7621×10^{-5}	7.3144×10^{-5}	+53.32
2.00	4.7132×10^{-5}	6.1760×10^{-5}	5.9812×10^{-5}	+31.04
10.00	2.0083×10^{-5}	2.0881×10^{-5}	2.0804×10^{-5}	+3.97
100.00	2.6199×10^{-6}	2.6210×10^{-6}	2.6206×10^{-6}	+0.04

Notes: Steady-state MSE against the true state at quarter ends, over $N = 1000$ replications with a 50-quarter burn-in excluded. CT is the accumulator-augmented continuous-time filter; EDT and Euler are the monthly-base discrete filters with time-varying observation patterns. The final column is the fold-change improvement of CT over EDT, with the paired comparison rejecting equality at $p < 0.01$ for all $\rho \leq 10$ and the gap economically nil at $\rho = 100$.

Table 15: The Frequency Effect: MSE Ratio S1/S0 per Filter

ρ	**CT**	**EDT**	**Euler**
0.01	0.1349	3.8291	3.3966
0.10	0.1522	5.0203	4.3879
0.50	0.1961	0.2982	0.3190
1.00	0.2326	0.3171	0.3067
2.00	0.2739	0.3393	0.3322
10.00	0.3301	0.3412	0.3400
100.00	0.3362	0.3363	0.3361

Notes: Steady-state MSE under the mixed design relative to the uniform quarterly anchor, per filter, paired at the path level; values below one measure the tracking value of the intra-quarter observations. The gain is channel-local in the sense of Remark 4.8: moving the inflation flow to monthly divides its resolution factor roughly by three ($\Phi(1/12) = 0.042$ against $\Phi(1/4) = 0.113$), while the monthly stock improves quadratically ($1 - \Lambda$ falls ninefold).

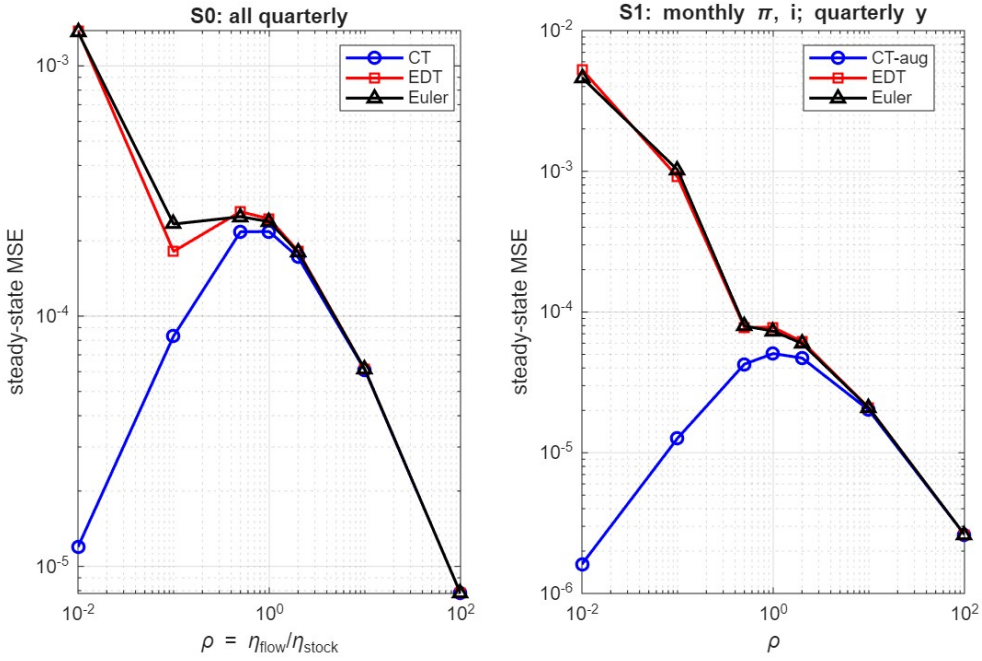


Figure 10: Steady-state filtering MSE across the precision grid under the uniform quarterly design S0 (left) and the mixed design S1 (right), log–log scale. Intra-quarter observation compresses all filters' errors; the CT–EDT separation on the flow-precise side persists under S1, because the pass-through penalty is a representation effect that finer sampling does not remove.

Table 16: Mixed Frequency: CT Stacked Recovery and Closed-Form Ceilings—the Two Jumps

	S0 ceilings		**S0 quarterly**		**S1 quarterly**		**S1 monthly**	
ρ	flow	stacked	flow	stacked	flow	stacked	flow	stacked
0.01	0.9857	0.9857	0.9857	0.9860	0.9979	0.9980	0.9977	0.9978
0.10	0.8945	0.8947	0.8951	0.8956	0.9825	0.9825	0.9808	0.9809
0.50	0.6750	0.6916	0.6736	0.6908	0.9222	0.9247	0.9155	0.9165
1.00	0.5469	0.6123	0.5430	0.6103	0.8598	0.8764	0.8502	0.8574
2.00	0.4214	0.6001	0.4149	0.5979	0.7624	0.8367	0.7531	0.7901
10.00	0.2042	0.7823	0.1942	0.7809	0.4508	0.9074	0.4551	0.8032
100.00	0.0658	0.9655	0.0546	0.9653	0.1451	0.9881	0.1571	0.9660

Notes: CT recovery under the flow-only and stacked maps at the three target–design pairs, with the S0 quarterly closed-form ceilings for reference; the S1 per-position ceilings appear in Table 19. Reading left to right displays the two jumps: adding intra-quarter information (S0 → S1 at the quarterly target) and moving the recovery to the monthly increment (S1 quarterly → S1 monthly), the resolution of the recovered object itself.

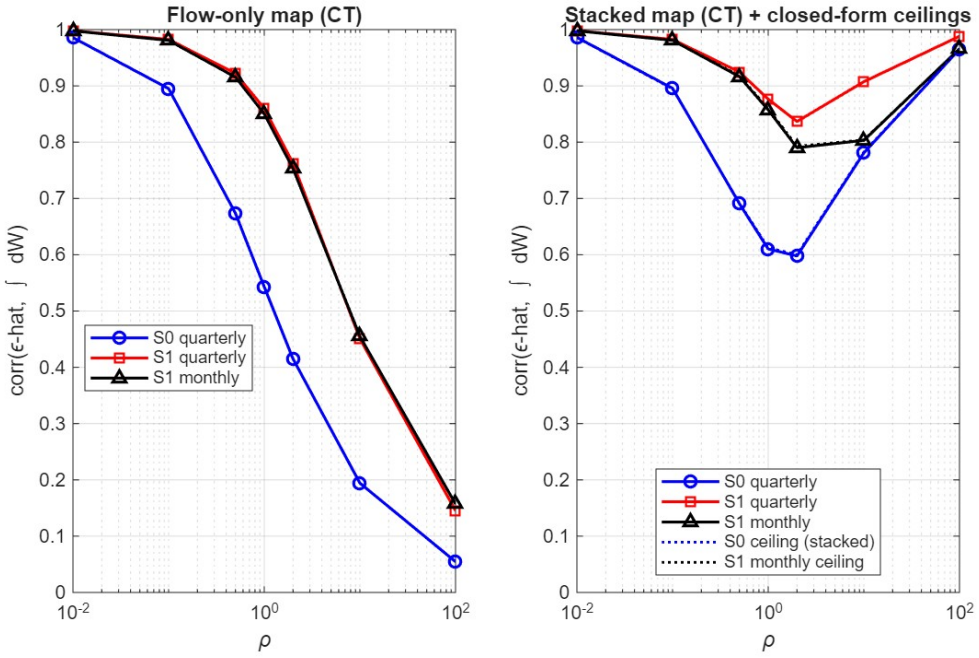


Figure 11: Mixed Frequency: CT shock recovery under the flow-only map (left) and the stacked map with closed-form ceilings (right), at the three target–design pairs.

Table 17: Uniform Quarterly: MSE Improvement over the Naive Euler Filter and over EDT

ρ	**CT vs. Euler (%)**	**EDT vs. Euler (%)**	**CT vs. EDT (%)**	p **(CT vs. EDT)**
0.01	+11362.12	-1.53	+11540.53	$< 10^{-300}$***
0.10	+192.84	+29.37	+126.35	$< 10^{-300}$***
0.50	+15.86	-4.53	+21.35	1.0×10^{-221}***
1.00	+9.62	-2.61	+12.56	1.8×10^{-147}***
2.00	+3.81	-1.09	+4.95	5.0×10^{-60}***
10.00	+0.10	+0.02	+0.08	0.423
100.00	+0.02	+0.04	-0.02	0.134

Notes: All MSEs are steady-state (burn-in excluded). The pass-through comparison has an exact theoretical value in closed form (Appendix C), which the CT vs. EDT column attains within Monte Carlo error at every ρ. Fold-change improvements, $(\text{MSE}_{\text{other}} - \text{MSE}_{\text{method}})/\text{MSE}_{\text{method}} \times 100$, positive when the method outperforms the comparison filter; this form does not saturate at 100%, so order-of-magnitude gaps read as such. The final column reports the paired t-test p-value for the CT–EDT comparison across replications. Both correctly specified filters dominate Euler wherever flows carry information, because Euler discards the flow-integral operator $\Gamma(h)$ entirely; the single reversal at $\rho = 0.01$ is the endpoint-surrogate effect discussed under Finding 4.

Table 18: Mixed Frequency: Flow-Only Shock Recovery by Filter—S0 Quarterly, S1 Quarterly, S1 Monthly

	S0 quarterly			**S1 quarterly**			**S1 monthly**		
ρ	CT	EDT	Euler	CT	EDT	Euler	CT	EDT	Euler
0.01	0.9857	0.8586	0.8595	0.9979	0.6248	0.6427	0.9977	0.8242	0.8369
0.10	0.8951	0.8850	0.8802	0.9825	0.8995	0.8880	0.9808	0.9455	0.9414
0.50	0.6736	0.6717	0.6720	0.9222	0.9174	0.9174	0.9155	0.9143	0.9142
1.00	0.5430	0.5422	0.5422	0.8598	0.8548	0.8550	0.8502	0.8495	0.8496
2.00	0.4149	0.4147	0.4147	0.7624	0.7580	0.7580	0.7531	0.7529	0.7529
10.00	0.1942	0.1942	0.1942	0.4508	0.4499	0.4498	0.4551	0.4551	0.4551
100.00	0.0546	0.0546	0.0546	0.1451	0.1451	0.1451	0.1571	0.1571	0.1571

Notes: Correlation between the flow-only recovered increment and the true increment, pooled steady-state. S1 recovers monthly (inflation alone intra-quarter; inflation and the quarterly output-gap rate jointly at quarter ends, a real-time convention that mixes windows and is documented as such); the S1 quarterly columns aggregate three monthly recoveries against the quarterly increment $\int_q dW$, the common target with S0.

Table 19: The Identification Wave: Per-Position Monthly Ceilings and CT Stacked Recovery

	Ceiling			**CT empirical**		
ρ	month 1	month 2	month 3	month 1	month 2	month 3
0.01	0.9978	0.9978	0.9978	0.9978	0.9978	0.9977
0.10	0.9812	0.9812	0.9813	0.9809	0.9812	0.9805
0.50	0.9176	0.9176	0.9181	0.9167	0.9179	0.9149
1.00	0.8593	0.8592	0.8599	0.8579	0.8596	0.8545
2.00	0.7925	0.7925	0.7931	0.7912	0.7932	0.7859
10.00	0.8040	0.8040	0.8040	0.8030	0.8079	0.7987
100.00	0.9661	0.9661	0.9661	0.9654	0.9675	0.9651

Notes: Closed-form attainable correlation of the stacked monthly recovery per within-quarter observation position (months one and two condition on the inflation and stock innovations; month three, the quarter end, additionally on the quarterly output-gap innovation), against the CT filter's empirical per-position correlations pooled over the steady state. The oscillation of the attainable recovery with the within-quarter observation pattern is the identification wave of the mixed design.

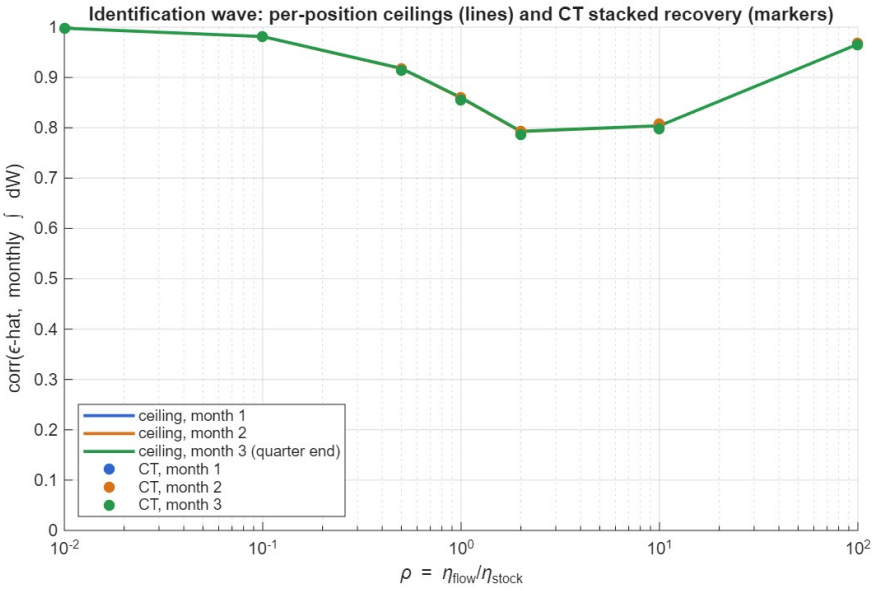


Figure 12: The identification wave: per-position closed-form ceilings (lines) and CT stacked per-position recovery (markers) across the precision grid. The quarter-end increment of the ceiling measures the marginal identification content of the quarterly flow beyond the monthly channels.

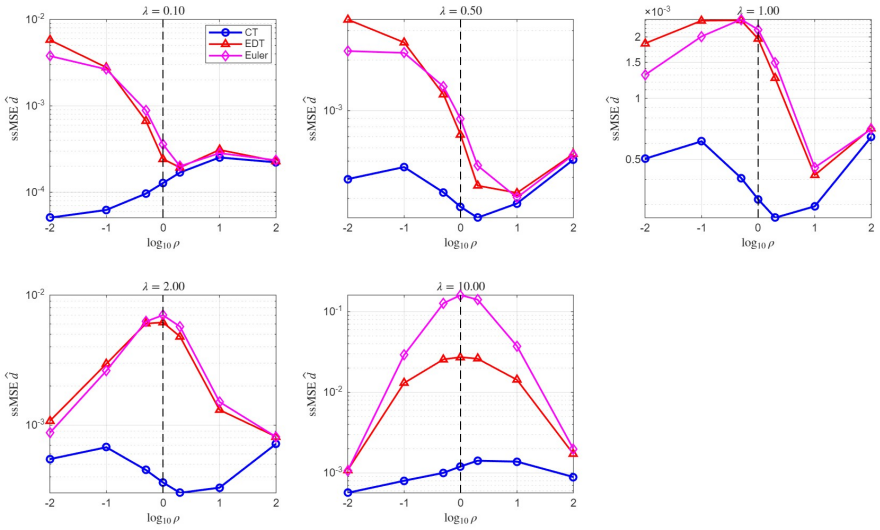


Figure 13: Steady-state MSE of the filtered demand state $\hat{d}_t$ under the three filters, across the precision grid $\rho = \eta_f/\eta_s$, one panel per shock mix λ. Low ρ means precise flows relative to the stock (and conversely); low λ means the monetary shock is weak relative to the demand shock. What to see: the rate-form filter (CT) attains the lowest MSE at every grid point—the ordering CT $\preceq$ EDT $\preceq$ Euler holds in all thirty-five cells with zero violations—and the filters reconverge at both extremes of ρ, where either release noise drowns the pass-through (high ρ) or all filters lean on the same near-noiseless level channel (low ρ).

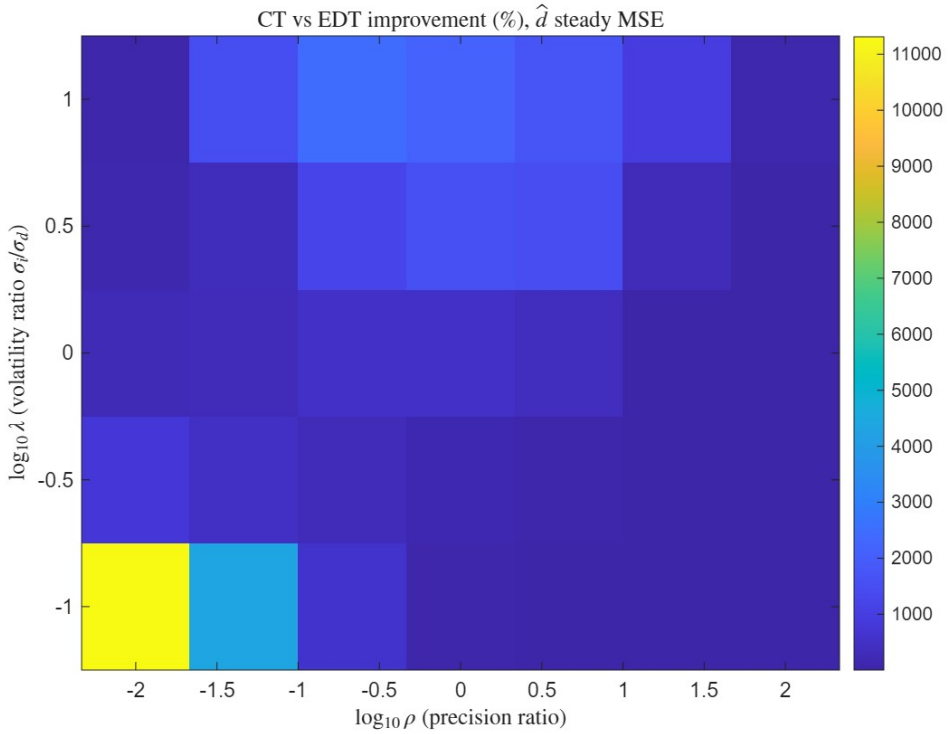


Figure 14: The rate-form advantage in tracking $\hat{d}_t$, expressed as the EDT excess MSE in percent of the rate-form MSE, on the same grid. The advantage is not monotone in either coordinate: it peaks in the *interior* (roughly +1,400 to +2,250 percent in the central cells, maximum $\approx$ +11,300 percent at $(\lambda, \rho) = (0.1, 0.01)$), where precise flows carry a strong shared diffusion loading that the integral-form innovation blocks anchor to the window boundary. Given the rate-form DGP this is a genuine misspecification cost, not an implementation artifact: on the counterfactual integral-form DGP the ordering flips, as validated in the code.

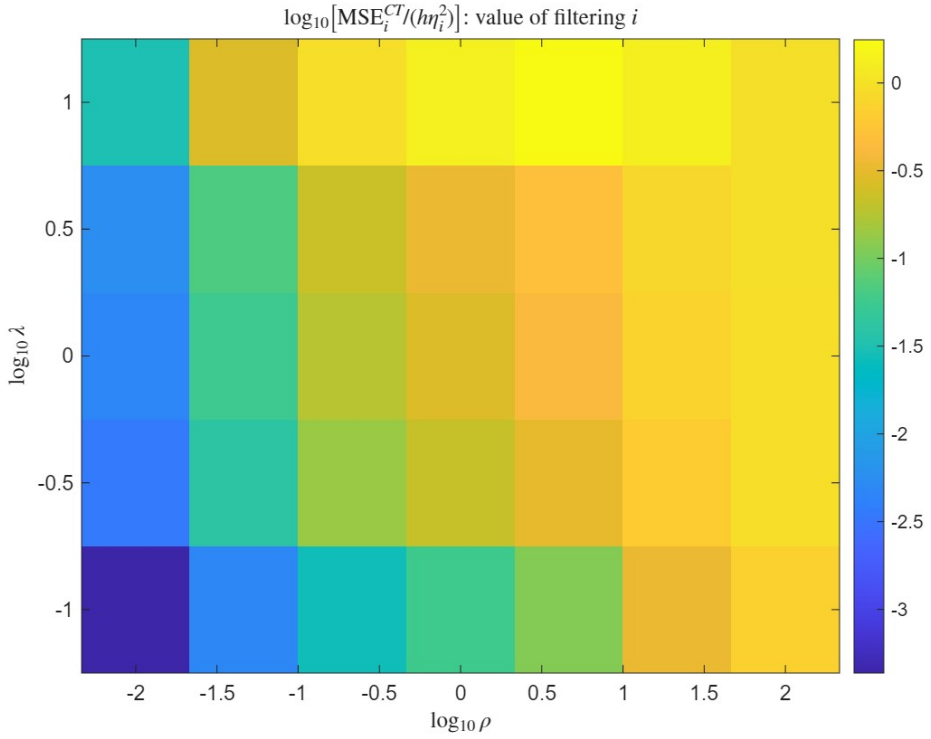


Figure 15: Filtering value for the quasi-observable interest-rate state: MSE of $\hat{\imath}_t$ relative to the variance of the raw stock reading ($h\eta_i^2$; unity means the filter is no better than reading the data). What to see: misspecified filtering can *destroy* a quasi-observable—at $(\lambda, \rho) = (2, 1)$ the EDT and Euler estimates are respectively about 2.7 and 8.0 times *worse* than the raw reading, because misattributing the flow pass-through forces spurious corrections into a state the data essentially report directly, while the rate-form filter improves on the reading by a factor of about three.

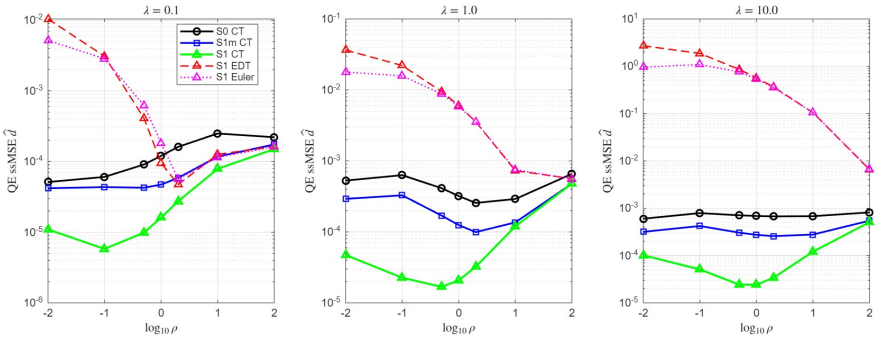


Figure 16: Quarter-end steady MSE of $\hat{d}_t$ across the three observation designs—S0 (all quarterly), S1^m (monthly π and i), S1 (daily i)—on identical exact daily paths. At $(\lambda, \rho) = (1, 0.1)$ the MSE falls from 6.3×10^{-4} to 2.3×10^{-5}, a factor of about twenty-eight, although d_t is never observed at daily frequency: ninety readings of i per quarter identify d through the drift channel A_{12} of the interest-rate equation. The figure also shows the escalation of the integral-form penalty under base-grid refinement (from +490 percent in S0 to +28,200 percent in S1 at $(1, 1)$): per window the omitted pass-through is $O(h_b^{1/2})$ against integral news of $O(h_b^{3/2})$.

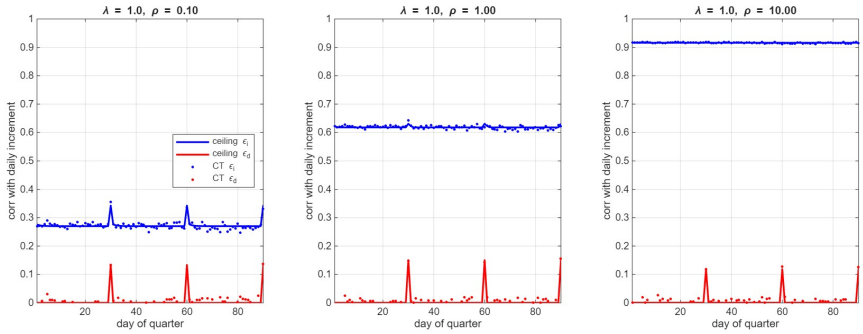


Figure 17: The identification wave: closed-form per-position recovery ceilings across the ninety daily positions of a quarter (lines), with the empirical rate-form recoveries attaining them position by position (markers; mean absolute deviation below 0.002). The ε_i profile is flat near its ceiling every day (the level channel operates daily), while interior-day ε_d ceilings are essentially zero—within a single day the preference shock reaches the interest rate only through $A_{12}h_d$—so the ε_d wave consists of spikes at the flow releases: month ends and, above all, the quarter-end GDP release.

F Observational Equivalence of the Primitives in a Single Release

This appendix states the identification content of one flow release on its own. Throughout, the release is driven by a single Ornstein–Uhlenbeck state, $dx_t = -\rho_d\, x_t\, dt + \sigma_d\, dW_t$, and we write $F \equiv e^{-\rho_d h}$ for the one-window autoregressive root and $\Gamma_h = (1-F)/\rho_d$ for the flow operator. To treat both primitives at once, we scale the pass-through term of (6) by $\tau \in [0,1]$:

$$y_k(\tau) \;=\; \frac{C}{h}\Big[\Gamma_h\, x_{k-1} + u_k + \tau\,\sigma_d\,\Delta W_k\Big] \;+\; \sqrt{h}\,\eta\,\varepsilon_k, \qquad u_k = \int_0^h \Gamma(h-s)\,\sigma_d\, dW_s, \quad (69)$$

so that $\tau = 0$ is Primitive I (integral form), $\tau = 1$ is Primitive II (rate form), and intermediate values describe a release whose construction passes through only part of the window's increment. The loading C, the release-noise scale η, and τ are the release-specific parameters; (ρ_d, σ_d) belong to the state.

Proposition F.1 (Autocovariances and the integral-form twin)**.** *Under* (69)*, $y_k(\tau)$ is a stationary Gaussian ARMA$(1,1)$ process with autoregressive root F, so that $\gamma_j = F\,\gamma_{j-1}$ for $j \geq 2$, and*

$$\gamma_1(\tau) = \Big(\frac{C\,\sigma_d\,\Gamma_h}{h}\Big)^2 \frac{1+2\rho_d\tau}{2\rho_d}, \quad (70)$$

$$\gamma_0(\tau) = \Big(\frac{C\,\sigma_d}{h}\Big)^2\Big[\frac{(h-\Gamma_h)(1+2\rho_d\tau)}{\rho_d^2} + \tau^2 h\Big] \;+\; h\,\eta^2. \quad (71)$$

Consequently, for every $\tau \in (0,1]$ the release $y_k(\tau)$ with parameters (C,η) has the same law as the integral-form release $y_k(0)$ with parameters

$$C_{\mathrm{I}}^2 = C^2\,(1+2\rho_d\tau), \qquad \eta_{\mathrm{I}}^2 = \eta^2 + \tau^2\,\frac{C^2\sigma_d^2}{h^2}. \quad (72)$$

In particular, with C and η free, τ is not identified from the second moments of a single release, at any observation frequency.

Proof. Write $y_k = (C/h)\Gamma_h x_{k-1} + \zeta_k$ with $\zeta_k = (C/h)(u_k + \tau\sigma_d \Delta W_k) + \sqrt{h}\eta\varepsilon_k$. The sequence ζ_k is white and independent of x_{k-j} for $j \geq 1$, while $x_{k-1} = Fx_{k-2} + \eta^x_{k-1}$ with η^x_{k-1} independent of y_{k-j} for $j \geq 2$. Hence $\gamma_j = F\gamma_{j-1}$ for $j \geq 2$, which is the autocovariance structure of an ARMA$(1,1)$ process with root F; the law of a stationary Gaussian ARMA$(1,1)$ is determined by (F, γ_0, γ_1).

The three window functionals are Itô integrals of the same Brownian path with kernels $e^{-\rho_d(h-s)}$ (state innovation η^x_k), $\Gamma(h-s)$ (path news u_k) and 1 (increment ΔW_k), so their moments follow from $\int_0^h$ of products of kernels, using $\rho_d\Gamma_h = 1 - F$:

$$\mathrm{Var}(x_k) = \frac{\sigma_d^2}{2\rho_d}, \quad \mathrm{Cov}(\eta^x_k, u_k) = \frac{\sigma_d^2\,\Gamma_h(1-F)}{2\rho_d}, \quad \mathrm{Cov}(\eta^x_k, \Delta W_k) = \sigma_d\,\Gamma_h,$$
$$\mathrm{Cov}(u_k, \Delta W_k) = \frac{\sigma_d\,(h - \Gamma_h)}{\rho_d}, \quad \mathrm{Var}(u_k) = \sigma_d^2\Big[\frac{h-\Gamma_h}{\rho_d^2} - \frac{\Gamma_h^2}{2\rho_d}\Big].$$

For the lag-one autocovariance only the carry-over of y_k is correlated with y_{k-1}:

$$\gamma_1 = \Big(\frac{C}{h}\Big)^2\Gamma_h\Big[\Gamma_h F\,\mathrm{Var}(x) + \mathrm{Cov}(\eta^x_{k-1}, u_{k-1}) + \tau\sigma_d\,\mathrm{Cov}(\eta^x_{k-1}, \Delta W_{k-1})\Big] = \Big(\frac{C\sigma_d\Gamma_h}{h}\Big)^2\Big[\frac{F+(1-F)}{2\rho_d} + \tau\Big],$$

which is (70). For the variance, $\gamma_0 = (C/h)^2\big[\Gamma_h^2\,\mathrm{Var}(x) + \mathrm{Var}(u_k) + 2\tau\sigma_d\,\mathrm{Cov}(u_k, \Delta W_k) + \tau^2\sigma_d^2 h\big] + h\eta^2$; the two $\Gamma_h^2/(2\rho_d)$ terms cancel and the remainder is (71).

For the twin, F is common to both parameterisations. Matching (70) at τ and at 0 requires $C_{\mathrm{I}}^2 = C^2(1 + 2\rho_d\tau)$. Substituting into (71) at $\tau = 0$ and equating with (71) at τ, the terms in $(h - \Gamma_h)(1 + 2\rho_d\tau)/\rho_d^2$ cancel exactly, leaving $h\eta_{\mathrm{I}}^2 = h\eta^2 + (C\sigma_d/h)^2\tau^2 h$, which is (72). Since (F, γ_0, γ_1) coincide and both processes are Gaussian, their laws coincide; the map (72) does not involve h except through the scale of η_{I}, so the conclusion holds at every observation frequency, and aggregating the release to a coarser window is a deterministic map of its law that adds no information. □

The map (72) has a direct reading. The pass-through term $(\tau C\sigma_d/h)\Delta W_k$ has per-release variance $\tau^2 C^2\sigma_d^2/h$; in the release-noise convention $\sqrt{h}\,\eta$ this is exactly the excess $\eta_{\mathrm{I}}^2 - \eta^2$. Its covariance with the state innovation, $\tau\sigma_d^2\Gamma_h$, is what raises γ_1 and is absorbed into the loading, $C_{\mathrm{I}}^2/C^2 = 1 + 2\rho_d\tau$. An estimator restricted to the integral form therefore records the pass-through variance as release noise and its covariance as loading, and fits the release exactly as well; the noise corner of Section 6 is the point (72) of this equivalence class, not an approximation to it.

Corollary F.2 (Loading-free statistic and the pass-through bound)**.** *The ratio of the lag-one autocovariance to the noise-corrected variance,*

$$R(\tau; \rho_d, h) \;\equiv\; \frac{\gamma_1(\tau)}{\gamma_0(\tau) - h\eta^2} \;=\; \frac{\rho_d\,\Gamma_h^2\,(1 + 2\rho_d\tau)}{2\,(h - \Gamma_h)(1 + 2\rho_d\tau) + 2\rho_d^2\tau^2 h}, \tag{73}$$

is free of C and σ_d, and is strictly decreasing in τ, with $R(0) = \rho_d\Gamma_h^2/\{2(h - \Gamma_h)\}$. Since the observed first autocorrelation satisfies $r_1 = \gamma_1/\gamma_0 \leq R(\tau)$, with equality if and only if $\eta = 0$: (i) $r_1 > R(1; \rho_d, h)$ rejects the rate form for any release-noise level; (ii) the solution $\bar\tau$ of $R(\bar\tau; \rho_d, h) = r_1$ is an upper bound on τ; (iii) a lower bound $\underline{\eta}^2$ on the release-noise

variance tightens it to $\tau_{\mathrm{hi}} \le \bar{\tau}$ *through* $R(\tau_{\mathrm{hi}}) = \gamma_1/(\gamma_0 - h\underline{\eta}^2)$*, and an upper bound on* η^2 *gives a lower bound on* τ *in the same way.*

Proof. Divide (70) by (71) less $h\eta^2$; the prefactor $(C\sigma_d/h)^2$ cancels. Writing the denominator as $2(h-\Gamma_h)+2\rho_d^2 h\,\tau^2/(1+2\rho_d\tau)$ shows monotonicity, since $\tau^2/(1+2\rho_d\tau)$ is increasing on $[0,\infty)$. The bounds follow from $\gamma_0 \ge \gamma_0 - h\eta^2$ and monotonicity. □

Remark F.3 (What breaks the equivalence)**.** *Proposition F.1 fails as soon as the loading is not free or a second observable shares the state. (a) If* C *is fixed by the structural model, the map* (72) *is unavailable and* τ *is identified from the release itself; this is why the pseudo-true analysis of Section 6, which holds the loadings at their calibrated values, finds a strictly positive Kullback–Leibler distance between the primitives. (b) If a stock* $s_k = C_s x_k + \sqrt{h}\eta_s \varepsilon_k^s$ *is observed at the window ends, its increment satisfies* $\mathrm{Cov}(s_k - s_{k-1},\, y_k) = \tau\, C_s C \sigma_d^2 \Gamma_h / h$ *and* $\mathrm{Cov}(s_{k-1},\, y_k) = C_s C \sigma_d^2 \Gamma_h/(2\rho_d h)$*, so that* $\tau = \mathrm{Cov}(\Delta s_k, y_k)/\{2\rho_d \,\mathrm{Cov}(s_{k-1}, y_k)\}$ *is identified without loadings,* σ_d*, or either noise scale. The zero at* $\tau = 0$ *is the time-reversibility of the stationary scalar Ornstein–Uhlenbeck process: the window integral of the state is equally correlated with its two endpoints, hence orthogonal to their difference. Both routes are used in Section 8.*

References

Aït-Sahalia, Y., Fan, J., and Xiu, D. (2010). High-frequency covariance estimates with noisy and asynchronous financial data. *Journal of the American Statistical Association*, 105(492):1504–1517.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *Review of Economic Studies*, 89(1):45–86.

Anderson, B. D. O. and Moore, J. B. (1979). *Optimal Filtering*. Prentice-Hall, Englewood Cliffs, NJ.

Andrle, M. (2013). Understanding DSGE filters in forecasting and policy analysis. IMF Working Paper No. 13/98, International Monetary Fund.

Bànbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2013). Now-casting and the real-time data flow. In *Handbook of Economic Forecasting*, volume 2A, pages 195–237. Elsevier.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2008). Designing realised kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica*, 76(6):1481–1536.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2011). Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics*, 162(2):149–169.

Bergstrom, A. R. (1966). Nonrecursive models as discrete approximations to systems of stochastic differential equations. *Econometrica*, 34(1):173–182.

Bergstrom, A. R. (1983). Gaussian estimation of structural parameters in higher order continuous time dynamic models. *Econometrica*, 51(1):117–152.

Bergstrom, A. R. (1984). Continuous time stochastic models and issues of aggregation over time. In Griliches, Z. and Intriligator, M. D., editors, *Handbook of Econometrics*, volume 2, chapter 20, pages 1145–1212. Elsevier.

Bergstrom, A. R. (1986). The estimation of open higher-order continuous time dynamic models with mixed stock and flow data. *Econometric Theory*, 2(3):350–373.

Bergstrom, A. R. (1990). *Continuous Time Econometric Modelling*. Oxford University Press.

Bok, B., Caratelli, D., Giannone, D., Sbordone, A. M., and Tambalotti, A. (2018). Macroeconomic nowcasting and forecasting with big data. *Annual Review of Economics*, 10:615–643.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Canova, F. and Sala, L. (2009). Back to square one: Identification issues in DSGE models. *Journal of Monetary Economics*, 56(4):431–449.

Chambers, M. J. (2009). Discrete time representations of cointegrated continuous time models with mixed sample data. *Econometric Theory*, 25(4):1030–1049.

Christensen, B. J., Posch, O., and van der Wel, M. (2016). Estimating dynamic equilibrium models using mixed frequency macro and financial data. *Journal of Econometrics*, 194(1):116–137.

Christensen, B. J., Neri, L., and Parra-Alvarez, J. C. (2024). Estimation of continuous-time linear DSGE models from discrete-time measurements. *Journal of Econometrics*, 244(2):105815.

Christiano, L. J. and Eichenbaum, M. (1987). Temporal aggregation and structural inference in macroeconomics. *Carnegie-Rochester Conference Series on Public Policy*, 26:63–130.

Corsi, F., Peluso, S., and Audrino, F. (2015). Missing in asynchronicity: A Kalman-EM approach for multivariate realized covariance estimation. *Journal of Applied Econometrics*, 30(3):377–397.

Diebold, F. X., Gunther, T. A., and Tay, A. S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39(4):863–883.

Fernández-Villaverde, J., Posch, O., and Rubio-Ramírez, J. F. (2012). Solving the New Keynesian model in continuous time. Unpublished manuscript, Federal Reserve Bank of Atlanta.

Forni, M. and Gambetti, L. (2014). Sufficient information in structural VARs. *Journal of Monetary Economics*, 66:124–136.

Foroni, C. and Marcellino, M. (2016). Mixed frequency structural vector auto-regressive models. *Journal of the Royal Statistical Society: Series A*, 179(2):403–425.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676.

Hansen, L. P. and Sargent, T. J. (2011). Robustness and ambiguity in continuous time. *Journal of Economic Theory*, 146(3):1195–1223.

Harvey, A. C. (1985). Trends and cycles in macroeconomic time series. *Journal of Business & Economic Statistics*, 3(3):216–227.

Harvey, A.C. and Stock, J.H. (1993). Estimation, smoothing, interpolation and distribution for structural time series models in continuous time. In Phillips, P.C.B., editor, *Models, Methods and Applications of Econometrics: Essays in Honour of A.R. Bergstrom*. Blackwell, Oxford.

Fernández-Villaverde, J., Rubio-Ramírez, J. F., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of understanding VARs. *American Economic Review*, 97(3):1021–1026.

Jacobson, M. M., Matthes, C., and Walker, T. B. (2023). Temporal aggregation bias and monetary policy transmission. Finance and Economics Discussion Series 2022-054r1, Board of Governors of the Federal Reserve System.

Kalman, R.E. and Bucy, R.S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, 83(1):95–108.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.

Komunjer, I. and Ng, S. (2011). Dynamic identification of dynamic stochastic general equilibrium models. *Econometrica*, 79(6):1995–2032.

Marcellino, M. (1999). Some consequences of temporal aggregation in empirical analysis. *Journal of Business & Economic Statistics*, 17(1):129–136.

Marcet, A. (1991). Temporal aggregation of economic time series. In Hansen, L. P. and Sargent, T. J., editors, *Rational Expectations Econometrics*, chapter 10. Westview Press.

Nüsken, N., Reich, S., and Rozdeba, P. J. (2019). State and parameter estimation from observed signal increments. *Entropy*, 21(5):505.

Øksendal, B. (2003). *Stochastic Differential Equations: An Introduction with Applications*. Springer, Berlin, 6th edition.

Phillips, P. C. B. (1978). The treatment of flow data in the estimation of continuous time systems. In Bergstrom, A. R., Catt, A. J. L., Peston, M. H., and Silverstone, B. D. J., editors, *Stability and Inflation*, pages 257–274. Wiley, New York.

Parra-Alvarez, J. C., Posch, O., and Rátfai, R. (2026). Kalman–Bucy filtering in the NK model. Working paper, Aarhus University and Universität Hamburg.

Posch, O. (2020). Resurrecting the New Keynesian model: (Un)conventional policy and the Taylor rule. CESifo Working Paper No. 6925.

Rátfai, R. (2026). What is in a flow? Real-time filtering of mixed stock–flow systems: A technical note. Working paper, Universität Hamburg.

Richardson, A., van Florenstein Mulder, T., and Vehbi, T. (2021). Nowcasting GDP using machine-learning algorithms: A real-time assessment. *International Journal of Forecasting*, 37(2):941–948.

Shannon, C. E. (1949). Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21.

Shephard, N. and Xiu, D. (2017). Econometric analysis of multivariate realised QML: Estimation of the covariation of equity prices under asynchronous trading. *Journal of Econometrics*, 201(1):19–42.

Snudden, S. (2024). Don't ruin the surprise: Temporal aggregation bias in structural innovations. Working paper, Wilfrid Laurier University.

Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier, Amsterdam.

Van Loan, C. F. (1978). Computing integrals involving the matrix exponential. *IEEE Transactions on Automatic Control*, 23(3):395–404.

Weiss, A. A. (1984). Systematic sampling and temporal aggregation in time series models. *Journal of Econometrics*, 26(3):271–281.

Whittle, P. (1983). *Prediction and Regulation by Linear Least-Square Methods*. University of Minnesota Press, Minneapolis, 2nd edition.

Wiener, N. (1949). *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. MIT Press, Cambridge, MA.

Working, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica*, 28(4):916–918.

Zadrozny, P. A. (1988). Gaussian likelihood of continuous-time ARMAX models when data are stocks and flows at different frequencies. *Econometric Theory*, 4(1):108–124.

Bansal, R. and Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509.

Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296.

Lucas, R. E., Jr. (1978). Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445.

Schorfheide, F., Song, D., and Yaron, A. (2018). Identifying long-run risks: A Bayesian mixed-frequency approach. *Econometrica*, 86(2):617–654.

Buccheri, G. (2025). Identification of continuous-time linear filters when only discrete-time data is available. *Econometric Reviews*, 44(9). doi:10.1080/07474938.2025.2516336.

Bauer, Michael D., and Eric T. Swanson. (2023). "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual*, 37(1), 1–74.