

SESSION 1

Exercise Hints & Worked Solutions

Nowcasting with High-Frequency Data

Contents of this companion

No.	TOPIC	PAGE
	Exercises and Solutions	2
	Exercise 1: Build the two-clock ledger	2
	Exercise 2: Open the release	3
	Exercise 3: Reproduce the unit test	5
	Exercise 4: Signal and noise	7
	Exercise 5: A ragged update	9
	Exercise 6: Why time alone does not narrow a fixed-target band	11
	Exercise 7: Derive the quarterly weights	12
	Exercise 8: Decompose the news	14
	Exercise 9: Evaluate 2008Q4	15
	Exercise 10: Further state-space practice	17
	References	20

This companion is meant to be used in stages. Try the exercise first. Open the hint if you need a direction, then compare your work with the detailed solution. The solutions emphasize the reasoning behind each calculation, not only the final number.

On the website, the hints and solutions are collapsed until selected. In the PDF, every hint and solution is printed in full so the document can be read offline or used as a conventional solutions manual.

Exercises and Solutions

Exercise 1: Build the two-clock ledger

Problem. [core] For one monthly macroeconomic series, record the reference period, first-release date, first value, next revision date, and revised value. Explain which cells would appear in vintages dated one day before each release and one day after it. Why would storing only a single date create either misalignment or look-ahead bias?

Hint

Draw a row for one fixed reference period, then make four vintage columns: $r_1 - 1$, $r_1 + 1$, $r_2 - 1$, and $r_2 + 1$, where r_1 is the first-release date and r_2 is the revision date. Ask what a forecaster could actually have read from the source by each cutoff.

Detailed solution

The numerical values depend on the series chosen, but every correct ledger has the same structure. Let τ denote the month described by the observation, r_1 its first-release date, and r_2 its next revision date. Suppose the first published value is $x_\tau^{(r_1)}$ and the revised value is $x_\tau^{(r_2)}$.

Vintage cutoff	Value stored for reference period τ
One day before r_1	missing
One day after r_1	$x_\tau^{(r_1)}$
One day before r_2	$x_\tau^{(r_1)}$
One day after r_2	$x_\tau^{(r_2)}$

For example, if a July observation is first released in August and revised in September, it still belongs to July in the model. August and September are the dates at which different values for that July cell became available.

The distinction explains why one date is not enough.

1. If the database keeps only the reference period, it knows where the number belongs economically but not when a forecaster could use it. A backtest can then place the July value in a July information set even though it was not published until August. That is look-ahead bias.
2. If the database keeps only the release date, it knows when the number became public but can mistakenly treat an August release as an August observation. The series is then shifted away from the month it measures.
3. If the database overwrites the first value with the revision, it cannot reconstruct the vintage seen between r_1 and r_2 .

A complete submission should therefore record at least the series identifier, reference period, release timestamp, vintage value, units, and transformation. For a real data source, it should also record the source URL and retrieval timestamp so another reader can reproduce the ledger.

Exercise 2: Open the release

Problem. [core, pencil] Before a release, $\alpha \sim \mathcal{N}(1.5, 0.36)$. The indicator satisfies $y = \alpha + \varepsilon$, with $H = 0.64$, and the released value is 0.5. Compute v, F, K, a^+ , and P^+ . Explain the result in words. Repeat for a release equal to 1.5 and distinguish the effect on the posterior mean from the effect on its variance.

Hint

In this scalar measurement equation, $Z = 1$. Start with $v = y - a^-$, $F = P^- + H$, and $K = P^-/F$. The observed value affects v , but it does not enter F or K .

Detailed solution

The prior mean and variance are

$$a^- = 1.5, \quad P^- = 0.36.$$

For the release $y = 0.5$, the innovation is the released value minus the model's prior expectation:

$$v = y - a^- = 0.5 - 1.5 = -1.0.$$

The innovation variance combines uncertainty about the state with measurement noise:

$$F = P^- + H = 0.36 + 0.64 = 1.00.$$

The Kalman gain is therefore

$$K = \frac{P^-}{F} = \frac{0.36}{1.00} = 0.36.$$

The posterior mean equals the prior mean plus the gain-weighted surprise:

$$a^+ = a^- + Kv = 1.5 + 0.36(-1.0) = 1.14.$$

Finally,

$$P^+ = (1 - K)P^- = (1 - 0.36)(0.36) = 0.2304.$$

The release is one percentage point below expectation, but the state estimate moves down by only 0.36 point. The model does not treat the indicator as a perfect observation of the state: 64 percent of the innovation variance is measurement noise. The posterior standard deviation falls from $\sqrt{0.36} = 0.60$ to $\sqrt{0.2304} = 0.48$.

Now set $y = 1.5$. Then

$$v = 1.5 - 1.5 = 0, \quad F = 1, \quad K = 0.36.$$

The posterior mean does not move:

$$a^+ = 1.5 + 0.36(0) = 1.5.$$

The posterior variance is still 0.2304. This is the central distinction. The realized surprise determines the direction and size of the mean revision. The precision of the measurement determines how much uncertainty is removed. A release that exactly matches the forecast can teach us that the prior estimate was accurate without changing its numerical center.

Exercise 3: Reproduce the unit test

Problem. [core, computational] Implement the local level filter and reproduce the two-period numerical unit test from the notes without hard-coding any intermediate values. Add assertions for both filtered means and variances. Then replace y_2 by a missing value and verify that the second update is skipped.

Hint

Carry two objects into each iteration: the predicted mean $a_{\tau|\tau-1}$ and predicted variance $P_{\tau|\tau-1}$. If an observation is missing, copy those predicted moments into the filtered moments and proceed directly to the next prediction.

Detailed solution

For the local level model, the measurement update is

$$v_{\tau} = y_{\tau} - a_{\tau|\tau-1}, \quad F_{\tau} = P_{\tau|\tau-1} + H, \quad K_{\tau} = \frac{P_{\tau|\tau-1}}{F_{\tau}},$$

followed by

$$a_{\tau|\tau} = a_{\tau|\tau-1} + K_{\tau} v_{\tau}, \quad P_{\tau|\tau} = (1 - K_{\tau}) P_{\tau|\tau-1}.$$

The prediction step is especially simple:

$$a_{\tau+1|\tau} = a_{\tau|\tau}, \quad P_{\tau+1|\tau} = P_{\tau|\tau} + Q.$$

One Python implementation is:

```
import numpy as np

def local_level_filter(y, Q, H, a0, P0):
    rows = []
    a_pred, P_pred = float(a0), float(P0)

    for period, observation in enumerate(y, start=1):
        if observation is None or np.isnan(observation):
            v = F = K = np.nan
            a_filt, P_filt = a_pred, P_pred
        else:
            v = observation - a_pred
            F = P_pred + H
            K = P_pred / F
            a_filt = a_pred + K * v
            P_filt = (1.0 - K) * P_pred

        P_next = P_filt + Q
        rows.append({
            "period": period,
            "v": v,
            "F": F,
            "K": K,
            "a_filt": a_filt,
            "P_filt": P_filt,
            "P_next": P_next,
        })
        a_pred, P_pred = a_filt, P_next

    return rows

rows = local_level_filter(
    y=np.array([2.4, 1.9]),
    Q=0.1024,
    H=1.0,
    a0=0.0,
```

The results, shown to six decimal places, are:

Quantity	$\tau = 1$	$\tau = 2$
v_τ	2.400000	-0.281818
F_τ	11.000000	2.011491
K_τ	0.909091	0.502856
$a_{\tau \tau}$	2.181818	2.040104
$P_{\tau \tau}$	0.909091	0.502856
$P_{\tau+1 \tau}$	1.011491	0.605256

For the missing-data check, replace the second observation by `np.nan`. The second innovation, variance, and gain are not computed. The filtered moments are copied from the prediction:

$$a_{2|2} = a_{2|1} = 2.181818, \quad P_{2|2} = P_{2|1} = 1.011491.$$

The next predicted variance is then $P_{3|2} = 1.011491 + 0.1024 = 1.113891$. The mean stays flat while uncertainty continues to accumulate. That is exactly what “skip the update” should mean; setting the missing observation to zero would produce a false negative surprise and is therefore not an acceptable shortcut.

Exercise 4: Signal and noise

Problem. [core, pencil] Derive the steady-state predicted variance and gain for the local level model. Compute $\bar{K}$ for $q \in \{0.01, 0.1, 1, 10\}$. For each value, state whether variation in the observed series is being attributed mainly to state movement or measurement noise.

Hint

Write the filtered variance as $P_{\tau|\tau} = P_{\tau|\tau-1}H/(P_{\tau|\tau-1} + H)$, add Q in the prediction step, and divide the entire recursion by H .

Detailed solution

Let $P_\tau = P_{\tau|\tau-1}$ denote the variance just before observing y_τ . The scalar gain is

$$K_\tau = \frac{P_\tau}{P_\tau + H}.$$

The filtered variance can be written as

$$P_{\tau|\tau} = (1 - K_\tau)P_\tau = \left(1 - \frac{P_\tau}{P_\tau + H}\right)P_\tau = \frac{P_\tau H}{P_\tau + H}.$$

Prediction adds the next state shock variance:

$$P_{\tau+1} = \frac{P_\tau H}{P_\tau + H} + Q.$$

Define the scaled predicted variance $\pi_\tau = P_\tau/H$ and the signal-to-noise ratio $q = Q/H$. Dividing by H gives

$$\pi_{\tau+1} = \frac{\pi_\tau}{\pi_\tau + 1} + q.$$

At a steady state, $\pi_{\tau+1} = \pi_\tau = \bar{\pi}$:

$$\bar{\pi} = \frac{\bar{\pi}}{\bar{\pi} + 1} + q.$$

Multiplying by $\bar{\pi} + 1$ and collecting terms yields

$$\bar{\pi}^2 - q\bar{\pi} - q = 0.$$

The quadratic formula gives

$$\bar{\pi} = \frac{q \pm \sqrt{q^2 + 4q}}{2}.$$

A variance cannot be negative, so the relevant root and gain are

$$\bar{\pi} = \frac{q + \sqrt{q^2 + 4q}}{2}, \quad \bar{K} = \frac{\bar{\pi}}{\bar{\pi} + 1}.$$

Substitution gives:

$q = Q/H$	$\bar{\pi}$	$\bar{K}$	Interpretation
0.01	0.1051	0.0951	Measurement noise dominates; the filter smooths heavily.
0.10	0.3702	0.2702	Noise still dominates; a release receives limited weight.
1.00	1.6180	0.6180	State and measurement variances are comparable; the filter is moderately responsive.
10.00	10.9161	0.9161	State movement dominates; the latest observation receives most of the weight.

The gain is not the fraction of total observed variance mechanically labeled “signal.” It is the optimal revision coefficient after accounting for the uncertainty carried into the period. Its monotonic rise with q nevertheless has the expected intuition: a rapidly moving state makes old information stale, while a noisy measurement makes any one release less trustworthy.

Exercise 5: A ragged update

Problem. [core, pencil] Let $\alpha \mid \Omega_{r-} \sim \mathcal{N}(0, 1)$, $Z = (1, 2)'$, and $H = \text{diag}(1, 4)$. At vintage r , only the first observation is released and its value is 1. At the next vintage, the second observation is released with value -1. Perform the two scalar updates in release order. Identify the reference-period rows and the

release dates separately.

Hint

For a scalar state observed through row z_i , use $v_i = y_i - z_i a$, $F_i = z_i^2 P + H_i$, and $K_i = P z_i / F_i$. The posterior from the first release becomes the prior for the second. Do not advance the state merely because the second row arrived on a later calendar date.

Detailed solution

Both rows measure the same reference-period state α_r . Their release dates differ. We therefore process the first row at release date $r_1 = r$, then process the second row at release date $r_2 > r_1$, without inserting an economic time transition between them.

For the first row, $z_1 = 1$, $H_1 = 1$, and $y_1 = 1$. Starting from $a_0 = 0$ and $P_0 = 1$,

$$v_1 = y_1 - z_1 a_0 = 1 - 1(0) = 1,$$

$$F_1 = z_1^2 P_0 + H_1 = 1^2(1) + 1 = 2,$$

$$K_1 = \frac{P_0 z_1}{F_1} = \frac{1}{2}.$$

Thus

$$a_1 = a_0 + K_1 v_1 = \frac{1}{2}, \quad P_1 = (1 - K_1 z_1) P_0 = \frac{1}{2}.$$

At the next vintage, the second row has $z_2 = 2$, $H_2 = 4$, and $y_2 = -1$. Its model expectation is $z_2 a_1 = 2(1/2) = 1$, so

$$v_2 = y_2 - z_2 a_1 = -1 - 1 = -2.$$

The innovation variance and gain are

$$F_2 = z_2^2 P_1 + H_2 = 2^2 \left(\frac{1}{2} \right) + 4 = 6,$$

$$K_2 = \frac{P_1 z_2}{F_2} = \frac{(1/2)2}{6} = \frac{1}{6}.$$

The second update gives

$$a_2 = a_1 + K_2 v_2 = \frac{1}{2} + \frac{1}{6}(-2) = \frac{1}{6},$$

$$P_2 = (1 - K_2 z_2)P_1 = \left(1 - \frac{1}{3}\right) \frac{1}{2} = \frac{1}{3}.$$

There is a useful precision check. Dividing the second measurement by two gives $y_2/2 = \alpha + \varepsilon_2/2$, whose error variance is $4/2^2 = 1$. The first measurement also has error variance 1. The prior and the two transformed observations therefore contribute equal precision, and their average is

$$\frac{0 + 1 + (-1/2)}{3} = \frac{1}{6}.$$

The bookkeeping should label the cells as row 1, reference period τ , released at r_1 , and row 2, reference period τ , released at r_2 . Confusing r_2 with a new reference period would apply a state transition that the economic timing does not justify.

Exercise 6: Why time alone does not narrow a fixed-target band

Problem. [core, pencil] Consider the local level model and a target h periods ahead. Show that its variance is $P_{\tau|\tau} + hQ$. Advance one period with no observation and show that the variance for the same fixed target is unchanged.

Hint

Write the future state as the current filtered state plus the sum of the next h independent state shocks. After one period passes, the current-state variance gains one Q , while the remaining horizon loses one Q .

Detailed solution

In the local level model,

$$\alpha_{\tau+1} = \alpha_{\tau} + \eta_{\tau}, \quad \text{Var}(\eta_{\tau}) = Q.$$

Iterating the state equation h times gives

$$\alpha_{\tau+h} = \alpha_{\tau} + \eta_{\tau} + \eta_{\tau+1} + \cdots + \eta_{\tau+h-1}.$$

Conditional on information through τ , the filtered state has variance $P_{\tau|\tau}$. The future shocks are independent of that state and of one another, so variances add:

$$\text{Var}(\alpha_{\tau+h} \mid Y_{\tau}) = P_{\tau|\tau} + \underbrace{Q + \cdots + Q}_{h \text{ shocks}} = P_{\tau|\tau} + hQ.$$

Now let one period pass without an observation. Because there is no measurement update,

$$P_{\tau+1|\tau+1} = P_{\tau+1|\tau} = P_{\tau|\tau} + Q.$$

The fixed target is still $\alpha_{\tau+h}$, but from date $\tau+1$ it is only $h-1$ steps ahead. Its variance is

$$P_{\tau+1|\tau+1} + (h-1)Q = (P_{\tau|\tau} + Q) + (h-1)Q = P_{\tau|\tau} + hQ.$$

Nothing has narrowed. One fewer future shock remains, but the shock that may have occurred during the elapsed period is now part of uncertainty about the current state. The two changes cancel exactly. A fixed-target band narrows when new information reduces state uncertainty, not merely because the calendar moves closer to the target date.

Exercise 7: Derive the quarterly weights

Problem. [core, pencil] Let a quarterly flow be the sum of three monthly levels. Use a first-order log approximation to derive the five monthly-growth weights used to map a monthly latent process into quarter-on-quarter growth. Explain why an end-of-quarter stock variable would require a different aggregation rule.

Hint

Pair the three current-quarter log levels with the three preceding-quarter log levels:

$$(\ell_{\tau} - \ell_{\tau-3}) + (\ell_{\tau-1} - \ell_{\tau-4}) + (\ell_{\tau-2} - \ell_{\tau-5}).$$

Expand each difference as a sum of monthly growth rates and count how often each growth rate appears.

Detailed solution

Let L_τ be the monthly flow level, $\ell_\tau = \log L_\tau$, and $g_\tau = \ell_\tau - \ell_{\tau-1}$. A first-order expansion around three equal monthly levels gives

$$\log(L_\tau + L_{\tau-1} + L_{\tau-2}) \approx \log 3 + \frac{1}{3}(\ell_\tau + \ell_{\tau-1} + \ell_{\tau-2}).$$

Subtract the same approximation for the preceding quarter. The constants cancel:

$$y_\tau^Q \approx \frac{1}{3}[(\ell_\tau - \ell_{\tau-3}) + (\ell_{\tau-1} - \ell_{\tau-4}) + (\ell_{\tau-2} - \ell_{\tau-5})].$$

Expand the three differences:

$$\ell_\tau - \ell_{\tau-3} = g_\tau + g_{\tau-1} + g_{\tau-2},$$

$$\ell_{\tau-1} - \ell_{\tau-4} = g_{\tau-1} + g_{\tau-2} + g_{\tau-3},$$

$$\ell_{\tau-2} - \ell_{\tau-5} = g_{\tau-2} + g_{\tau-3} + g_{\tau-4}.$$

Collecting like terms gives

$$y_\tau^Q \approx \frac{1}{3}(g_\tau + 2g_{\tau-1} + 3g_{\tau-2} + 2g_{\tau-3} + g_{\tau-4}),$$

or

$$y_\tau^Q \approx \frac{1}{3}g_\tau + \frac{2}{3}g_{\tau-1} + g_{\tau-2} + \frac{2}{3}g_{\tau-3} + \frac{1}{3}g_{\tau-4}.$$

The tent shape $(1, 2, 3, 2, 1)/3$ is a counting result. The middle growth rate appears in all three paths connecting monthly levels across the two quarters; the rates at the ends appear in only one path.

A stock variable is not summed or averaged through the quarter. If the stock is measured at the quarter's end, quarter-on-quarter log growth is

$$\ell_\tau - \ell_{\tau-3} = g_\tau + g_{\tau-1} + g_{\tau-2}.$$

The corresponding weights are $(1, 1, 1)$, not the five-point tent. More generally, the aggregation row must reflect whether the published variable is a flow accumulated through the period, a period average, or a stock measured at a particular date.

Exercise 8: Decompose the news

Problem. [core] Using the numbers in Exercise 2, compute the revision as Kv . Then suppose a parameter is re-estimated after the release and moves the reported nowcast by another 0.1 percentage point. Write a small revision table that separates data news from parameter re-estimation.

Hint

First hold the parameters fixed and move from the old nowcast to $a^- + Kv$. Only after that step should you add the re-estimation effect. Record the sign of the 0.1-point parameter contribution explicitly.

Detailed solution

From Exercise 2,

$$v = -1.0, \quad K = 0.36.$$

The fixed-parameter data-news contribution is

$$Kv = 0.36(-1.0) = -0.36.$$

Starting from 1.50, the fixed-parameter updated nowcast is therefore

$$1.50 - 0.36 = 1.14.$$

Suppose “another 0.1 percentage point” means that re-estimation raises the reported nowcast by +0.10. A transparent revision table is:

Stage or contribution	Calculation	Contribution	Nowcast level
Before the release			1.50
Data news, old parameters	$Kv = 0.36(-1.0)$	-0.36	1.14
Parameter re-estimation	stated separately	+0.10	1.24
Total revision	$-0.36 + 0.10$	-0.26	1.24

The news decomposition answers a counterfactual question: how much would the nowcast have moved if the model parameters had remained fixed and only the new observation had changed? Re-estimation changes the mapping from data to the reported target, so it belongs in a separate row. The sign must not be left implicit. If re-estimation instead lowers the nowcast by 0.1 point, replace +0.10 with -0.10 ; the final level is then 1.04 and the total revision is -0.46. Either table is mechanically correct only if it states which direction the parameter effect takes.

Exercise 9: Evaluate 2008Q4

Problem. [core, data] Retrieve the GDP vintages corresponding to the advance, final, and current estimates of 2008Q4 real GDP growth. Document the retrieval dates and units. For a January 2009 nowcast of -4.0 percent, compute errors against the first release and the later vintage. State which target answers which economic question.

Hint

Keep three dates separate: the quarter being measured, the historical release date, and the date on which you retrieved the archived record. Choose and state an error convention before doing the subtraction.

Detailed solution

The source record used in these notes is:

	Historical release		
Target vintage	date	2008Q4 real GDP growth	Units/source record
Advance estimate	January 30, 2009	-3.8%	Quarter-on-quarter percent change at a seasonally adjusted annual rate; BEA advance release (U.S. Bureau of Economic Analysis, 2009a)
Final estimate	March 26, 2009	-6.3%	Same units; BEA final release (U.S. Bureau of Economic Analysis, 2009b)
Current series vintage	FRED record retrieved August 2, 2026	-8.5%	Same units; BEA series A191RL1Q225SBEA distributed by FRED (Federal Reserve Bank of St. Louis, 2026)

For the two historical BEA releases, a reproducible submission should add the date on which the student accessed each archived page. That retrieval date is not the same as January 30 or March 26, 2009: those are the dates on which the values first entered the public information set.

Define the signed error as

$$e = \text{nowcast} - \text{target}.$$

Against the advance estimate,

$$e_{\text{advance}} = -4.0 - (-3.8) = -0.2.$$

The absolute error is 0.2 percentage point. The negative sign says that the nowcast predicted a slightly larger contraction than the first release.

Against the current later vintage,

$$e_{\text{current}} = -4.0 - (-8.5) = 4.5.$$

The absolute error is 4.5 percentage points. The positive sign says that the nowcast was much less negative than the economy is now estimated to have been. For reference, the error against the March 2009 final estimate is $-4.0 - (-6.3) = 2.3$ points.

The first-release target asks, “Could the model predict the number that a real-time user was about to see?” It is appropriate for evaluating a forecast of the initial publication. The later-vintage target asks, “How close was the real-time estimate to a more mature retrospective measure of economic activity?” It is appropriate when the substantive objective is the underlying economy rather than the first print.

Neither target changes what the January forecaster knew. The information set must still stop in January 2009. Only the value used later to score that fixed nowcast changes.

Exercise 10: Further state-space practice

Problem. [extra] Put an ARMA(1,1) process in state-space form and verify the dimensions of T, R, Z, Q, H . Then derive the Joseph covariance form from the shorter covariance update when K is the optimal gain.

Hint

For the ARMA representation, store both the current observation and current innovation in the state. For the Joseph identity, expand the two quadratic terms and use $F = ZPZ' + H$ together with $K = PZ'F^{-1}$.

Detailed solution

Use the convention

$$y_t = \phi y_{t-1} + u_t + \theta u_{t-1}, \quad u_t \sim \mathcal{N}(0, \sigma_u^2).$$

Define the two-dimensional state

$$\alpha_t = \begin{pmatrix} y_t \\ u_t \end{pmatrix}.$$

Then

$$\begin{pmatrix} y_{t+1} \\ u_{t+1} \end{pmatrix} = \begin{pmatrix} \phi & \theta \\ 0 & 0 \end{pmatrix} \begin{pmatrix} y_t \\ u_t \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \end{pmatrix} u_{t+1}.$$

The measurement equation is exact:

$$y_t = \begin{pmatrix} 1 & 0 \end{pmatrix} \alpha_t + \varepsilon_t, \quad H = \text{Var}(\varepsilon_t) = 0.$$

Thus one valid state-space representation is

$$T = \begin{pmatrix} \phi & \theta \\ 0 & 0 \end{pmatrix}, \quad R = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad Q = \begin{pmatrix} \sigma_u^2 \end{pmatrix}, \quad H = \begin{pmatrix} 0 \end{pmatrix}.$$

The dimensions are

Object	Dimension
α_t	2×1
T	2×2
R	2×1
Q	1×1

$$\begin{array}{ll} Z & 1 \times 2 \\ H & 1 \times 1 \end{array}$$

Substituting the first row of the transition equation gives $y_{t+1} = \phi y_t + \theta u_t + u_{t+1}$, which verifies the ARMA(1,1) recursion. The same new innovation enters both state components; that shared shock is why $R = (1, 1)'$ rather than two unrelated disturbances.

For the covariance result, let P denote the predicted state covariance and let

$$F = ZPZ' + H, \quad K = PZ'F^{-1}.$$

The Joseph form is

$$P^+ = (I - KZ)P(I - KZ)' + KHK'.$$

Expand it:

$$\begin{aligned} P^+ &= P - KZP - PZ'K' + KZPZ'K' + KHK' \\ &= P - KZP - PZ'K' + K(ZPZ' + H)K' \\ &= P - KZP - PZ'K' + KFK'. \end{aligned}$$

Because $K = PZ'F^{-1}$,

$$KF = PZ', \quad KFK' = PZ'K'.$$

Substitution cancels the last two terms:

$$P^+ = P - KZP = (I - KZ)P.$$

The Joseph and short forms are therefore identical in exact arithmetic when K is the optimal gain. The Joseph form is often safer in floating-point computation because it is written as a sum of two positive-semidefinite matrices. That structure makes symmetry and nonnegative variances more robust to rounding error.

References

FEDERAL RESERVE BANK OF ST. LOUIS. (2026). *Real gross domestic product [A191RL1Q225SBEA]*.

U.S. BUREAU OF ECONOMIC ANALYSIS. (2009a). *Gross domestic product, fourth quarter 2008 (advance)*.

U.S. BUREAU OF ECONOMIC ANALYSIS. (2009b). *Gross domestic product, fourth quarter 2008 (final) and corporate profits*.