Online Appendix

Information-Sharing Protocols and Reconciliation Risk in Two-Tier Fiscal Governance*

Campbell Leith†
University of Glasgow

Charles Nolan‡
University of Glasgow

Graeme Roy§
University of Glasgow

August 9, 2026

Abstract

This Online Appendix provides the formal derivations and auxiliary results supporting the main paper. It first states the key Gaussian updating results used throughout and then develops the baseline model, the results on calendar design, assumptions concordats and timing frictions, and the Scotland–UK reconciliation mapping. Subsequent Parts analyse forecast accuracy and reconciliation stability, calendar capacity and endogenous importing, repeated interaction, Bayesian learning and political narrative pressure.

*The usual disclaimer applies.

†Email: campbell.leith@glasgow.ac.uk
‡Email: charles.nolan@glasgow.ac.uk
§Email: graeme.roy@glasgow.ac.uk

Contents

1 Part I. Analytical Foundations 3
- 1.1 The Bayesian forecasting benchmark . . . 3
- 1.2 One signal and one unknown state . . . 4
- 1.3 Useful Gaussian properties . . . 5
- 1.4 Two posterior errors under a finite common prior . . . 6
- 1.5 Application to the OBR–SFC signal structure . . . 7
- 1.6 Optimality under quadratic loss . . . 8

2 Part II. Baseline Model 8
- 2.1 Private sharing of upstream precision-status information . . . 22

3 Part III. The Scotland–UK Reconciliation Mapping 26
- 3.1 Forecast accuracy and reconciliation stability . . . 26
- 3.2 Formal compatibility of the two objectives . . . 27
- 3.3 The Scottish reconciliation loading . . . 30

4 Part IV. Accuracy and Reconciliation Stability 32

5 Part V. Calendar Capacity and Endogenous Importing 37
- 5.1 Primitives . . . 37
- 5.2 Posterior forecasts and attainable precision . . . 38
- 5.3 Independent forecasting and import . . . 39
- 5.4 Optimal effort under independent forecasting . . . 39
- 5.5 The import decision . . . 40
- 5.6 Selective concordats as block-by-block interface design . . . 41

6 Part VI. Repeated Interaction and Interface Stability 45
- 6.1 Stage game . . . 45
- 6.2 Perfect monitoring . . . 46
- 6.3 Imperfect monitoring . . . 48

7 Part VII. Bayesian Learning and the Oversight Environment 49
- 7.1 Competence types and published forecasts . . . 50
- 7.2 Likelihood conditional on the realised state . . . 50
- 7.3 Error-only evaluation as a restricted special case . . . 52
- 7.4 The oversight environment . . . 53
- 7.5 Identification, limitations and design implications . . . 53

8 Part VIII. Political Narrative Pressure 54

1 Part I. Analytical Foundations

This Part collects the notation and Gaussian updating results used throughout the Appendix. It establishes the Bayesian benchmark against which the informational and incentive effects of alternative coordination protocols are evaluated. Readers familiar with Gaussian signal extraction may skip this and return to the formulas below as required.

1.1 The Bayesian forecasting benchmark

Uncertainty about fiscal fundamentals is represented by commonly known prior distributions. In the baseline application,

$$\theta \sim \mathcal{N}(0, V_\theta), \qquad \sigma \sim \mathcal{N}(0, V_\sigma),$$

where θ denotes a UK-wide fiscal-relevant component and σ a Scotland-specific component. Institutions observe noisy signals and update their beliefs using Bayes' rule. Under quadratic loss, the accuracy-minimising forecast is the conditional expectation,

$$\hat{\theta} = \mathbb{E}[\theta \mid \mathcal{I}], \qquad \hat{y} = \mathbb{E}[y \mid \mathcal{I}],$$

for the relevant information set $\mathcal{I}$.

This benchmark serves two purposes. First, it provides a parsimonious representation of forecast formation from prior information and new evidence. Second, it separates pure learning from incentive-driven alignment. If an institution acquires additional information, the Bayesian benchmark incorporates that information through the posterior mean. Movement toward another institution's published baseline beyond that warranted by the posterior can therefore be attributed to the reputational, political or institutional incentives introduced later in the analysis.

Equivalent signal-extraction or rational-expectations formulations could represent the same substantive objects. The Gaussian Bayesian formulation

is used because it makes the timing and content of information explicit and delivers closed-form expressions for forecasts and forecast-error variances.

1.2 One signal and one unknown state

Let θ be an unknown scalar state with prior

$$\theta \sim \mathcal{N}(\mu_0, V_\theta).$$

A forecaster observes

$$z = \theta + \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, v),$$

where ε is independent of θ. Equivalently,

$$z \mid \theta \sim \mathcal{N}(\theta, v).$$

The prior variance V_θ measures ex ante uncertainty about the state; the signal-noise variance v measures the reliability of the new observation.

Bayes' rule gives

$$p(\theta \mid z) \propto p(z \mid \theta) p(\theta).$$

Ignoring constants that do not depend on θ,

$$p(\theta) \propto \exp\left\{-\frac{(\theta - \mu_0)^2}{2V_\theta}\right\}, \qquad p(z \mid \theta) \propto \exp\left\{-\frac{(z - \theta)^2}{2v}\right\}.$$

Hence

$$p(\theta \mid z) \propto \exp\left\{-\frac{1}{2}\left[\frac{(\theta - \mu_0)^2}{V_\theta} + \frac{(z - \theta)^2}{v}\right]\right\}.$$

Expanding the quadratic terms and retaining those that depend on θ gives

$$\frac{(\theta - \mu_0)^2}{V_\theta} + \frac{(z - \theta)^2}{v} = \left(\frac{1}{V_\theta} + \frac{1}{v}\right)\theta^2 - 2\left(\frac{\mu_0}{V_\theta} + \frac{z}{v}\right)\theta + \text{constant}.$$

Define posterior precision by

$$\Lambda \equiv \frac{1}{V_\theta} + \frac{1}{v}.$$

Completing the square yields

$$\mu_1 = \frac{\frac{\mu_0}{V_\theta} + \frac{z}{v}}{\frac{1}{V_\theta} + \frac{1}{v}},$$

and therefore

$$\theta \mid z \sim \mathcal{N}(\mu_1, V_1),$$

where

$$V_1 = \left(\frac{1}{V_\theta} + \frac{1}{v}\right)^{-1}, \qquad \mu_1 = \frac{\frac{\mu_0}{V_\theta} + \frac{z}{v}}{\frac{1}{V_\theta} + \frac{1}{v}}.$$

The posterior mean can equivalently be written as

$$\mu_1 = \omega z + (1-\omega)\mu_0, \qquad \omega \equiv \frac{V_\theta}{V_\theta + v} \in (0,1).$$

Thus the relative weight on the signal increases with prior uncertainty and decreases with signal noise. Posterior precision is additive,

$$\frac{1}{V_1} = \frac{1}{V_\theta} + \frac{1}{v},$$

so that

$$V_1 < \min\{V_\theta, v\}.$$

The limiting cases provide immediate checks: as $v \to 0$, $\mu_1 \to z$ and $V_1 \to 0$; as $v \to \infty$, $\mu_1 \to \mu_0$ and $V_1 \to V_\theta$.

The main model sets $\mu_0 = 0$. In that case,

$$\mu_1 = \frac{V_\theta}{V_\theta + v} z, \qquad V_1 = \left(\frac{1}{V_\theta} + \frac{1}{v}\right)^{-1},$$

or, in forecast notation,

$$\hat{\theta} \equiv \mathbb{E}[\theta \mid z] = \frac{V_\theta}{V_\theta + v} z, \qquad \mathrm{Var}(\theta - \hat{\theta}) = \left(\frac{1}{V_\theta} + \frac{1}{v}\right)^{-1}.$$

1.3 Useful Gaussian properties

The derivations below repeatedly use three standard properties, together with the finite-prior covariance lemma established in the next subsection.

First, linear transformations preserve normality: if $X \sim \mathcal{N}(\mu, \sigma^2)$ and a, b are constants, then

$$aX + b \sim \mathcal{N}(a\mu + b, a^2\sigma^2).$$

Second, sums of independent normal random variables are normal: if $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$ and $Y \sim \mathcal{N}(\mu_Y, \sigma_Y^2)$ are independent, then

$$X + Y \sim \mathcal{N}(\mu_X + \mu_Y, \sigma_X^2 + \sigma_Y^2).$$

Third, conditional expectations are affine under joint normality. In particular, if (θ, z) are jointly normal, then $\mathbb{E}[\theta \mid z]$ is affine in z.

1.4 Two posterior errors under a finite common prior

The independence of signal noises does not imply independence of posterior forecast errors when two forecasters estimate the same latent state from a finite common prior. This distinction is used repeatedly below.

Lemma (finite-prior covariance of posterior errors). Let

$$\theta \sim \mathcal{N}(0, V_\theta), \qquad z_i = \theta + \varepsilon_i, \qquad \varepsilon_i \sim \mathcal{N}(0, v_i), \quad i \in \{O, S\},$$

where ε_O, ε_S and θ are mutually independent. Define

$$\hat{\theta}_i \equiv \mathbb{E}[\theta \mid z_i], \qquad e_i \equiv \theta - \hat{\theta}_i, \qquad V_i \equiv \mathrm{Var}(e_i) = \left(\frac{1}{V_\theta} + \frac{1}{v_i}\right)^{-1}.$$

Then

$$\mathrm{Cov}(e_O, e_S) = \frac{V_O V_S}{V_\theta} > 0. \tag{1}$$

In particular, independent signal noises generate positively correlated posterior errors unless the prior is diffuse in the limiting sense $V_\theta \to \infty$.

Proof. Write

$$\kappa_i \equiv \frac{V_\theta}{V_\theta + v_i}, \qquad \hat{\theta}_i = \kappa_i(\theta + \varepsilon_i).$$

Hence

$$e_i = (1 - \kappa_i)\theta - \kappa_i \varepsilon_i.$$

Because the signal noises are mutually independent and independent of θ,

$$\mathrm{Cov}(e_O, e_S) = (1 - \kappa_O)(1 - \kappa_S)V_\theta.$$

Using

$$1 - \kappa_i = \frac{v_i}{V_\theta + v_i} = \frac{V_i}{V_\theta}$$

gives (1). □

This says that the common prior leaves a residual component of uncertainty about θ in both forecasts. A concordat or a sequential publication rule can increase this common exposure, but it does not create covariance from a zero baseline. The relevant design object is therefore the *incremental* covariance generated by the interface.

1.5 Application to the OBR–SFC signal structure

The one-signal calculation applies directly to each signal in the baseline model. For the OBR signal,

$$\hat{\theta}_O = \frac{V_\theta}{V_\theta + v_O} z_O, \qquad V_O = \left(\frac{1}{V_\theta} + \frac{1}{v_O}\right)^{-1}.$$

For the SFC signal about the UK-wide component,

$$\hat{\theta}_S = \frac{V_\theta}{V_\theta + v_S} z_{S\theta}, \qquad V_S = \left(\frac{1}{V_\theta} + \frac{1}{v_S}\right)^{-1}.$$

For the SFC signal about the Scotland-specific component,

$$\hat{\sigma}_S = \frac{V_\sigma}{V_\sigma + w_S} z_{S\sigma}, \qquad V_\Sigma = \left(\frac{1}{V_\sigma} + \frac{1}{w_S}\right)^{-1}.$$

These posterior means are the 'pure-learning' benchmarks used in the subsequent analysis of publication order, assumptions concordats and institutional incentives.

1.6 Optimality under quadratic loss

Let an institution choose a forecast a to minimise expected squared error conditional on its information set $\mathcal{I}$:

$$\min_a \ \mathbb{E}\big[(\theta - a)^2 \mid \mathcal{I}\big].$$

Differentiating with respect to a gives

$$-2\mathbb{E}[\theta - a \mid \mathcal{I}] = 0,$$

and hence

$$a = \mathbb{E}[\theta \mid \mathcal{I}].$$

The posterior mean is therefore the unique accuracy-minimising forecast under quadratic loss.

2 Part II. Baseline Model

This Part states the benchmark environment and proves Propositions 1–3. Part I provides the Gaussian updating identities used below. The benchmark combines a common UK-wide component, a Scotland-specific component, institution-specific signals, and a protocol governing publication and the use of shared assumptions.

Let

$$\theta \sim \mathcal{N}(0, V_\theta), \qquad \sigma \sim \mathcal{N}(0, V_\sigma), \qquad \theta \perp \sigma,$$

and define

$$y \equiv \theta + \sigma.$$

The OBR observes

$$z_O = \theta + \varepsilon_O, \qquad \varepsilon_O \sim \mathcal{N}(0, v_O),$$

while the SFC observes

$$z_{S\theta} = \theta + \varepsilon_{S\theta}, \qquad \varepsilon_{S\theta} \sim \mathcal{N}(0, v_S),$$

and

$$z_{S\sigma} = \sigma + \varepsilon_{S\sigma}, \qquad \varepsilon_{S\sigma} \sim \mathcal{N}(0, w_S),$$

with all noise terms mutually independent and independent of (θ, σ). Under quadratic loss, the relevant posterior means and error variances are

$$\hat{\theta}_O \equiv \mathbb{E}[\theta \mid z_O] = \frac{V_\theta}{V_\theta + v_O}\, z_O, \qquad \mathrm{Var}(\theta - \hat{\theta}_O) = \left(\frac{1}{V_\theta} + \frac{1}{v_O}\right)^{-1} \equiv V_O,$$

$$\hat{\theta}_S \equiv \mathbb{E}[\theta \mid z_{S\theta}] = \frac{V_\theta}{V_\theta + v_S}\, z_{S\theta}, \qquad \mathrm{Var}(\theta - \hat{\theta}_S) = \left(\frac{1}{V_\theta} + \frac{1}{v_S}\right)^{-1} \equiv V_S,$$

and

$$\hat{\sigma}_S \equiv \mathbb{E}[\sigma \mid z_{S\sigma}] = \frac{V_\sigma}{V_\sigma + w_S}\, z_{S\sigma}, \qquad \mathrm{Var}(\sigma - \hat{\sigma}_S) = \left(\frac{1}{V_\sigma} + \frac{1}{w_S}\right)^{-1} \equiv V_\Sigma.$$

Proposition 1 considers a sequential calendar, $O \to S$, under which the SFC observes $\hat{\theta}_O$ before publishing. Proposition 2 compares a no-concordat regime, in which the SFC uses $\hat{\theta}_S$ for the common component, with a concordat regime, in which it must use $\hat{\theta}_O$. In both cases the reconciliation amount is

$$R \equiv (y - \hat{y}_S) + \lambda(\theta - \hat{\theta}_O), \qquad \lambda \in \mathbb{R}, \tag{2}$$

where $\hat{y}_S$ is the SFC's published forecast of y.

Proposition 1 (calendar and herding). Assume the calendar is $O \to S$, so SFC observes $\hat{\theta}_O$ before publishing. Suppose SFC chooses the *θ-component* of its published baseline, denoted $\tilde{\theta}$, to minimise

$$\mathcal{J}(\tilde{\theta}) \equiv \mathbb{E}\Big[(\theta - \tilde{\theta})^2 \mid z_{S\theta}, z_O\Big] \;+\; \eta(\tilde{\theta} - \hat{\theta}_O)^2, \qquad \eta \geq 0,$$

where η measures the strength of a "disagreement penalty" (perceived-independence concerns expressed as a cost of visibly diverging from the OBR baseline). Then the optimal $\tilde{\theta}^*$ satisfies

$$\tilde{\theta}^* = \frac{\hat{\theta}_B + \eta\, \hat{\theta}_O}{1 + \eta}, \tag{3}$$

where

$$\hat{\theta}_B \equiv \mathbb{E}[\theta \mid z_{S\theta}, z_O]$$

is the Bayesian posterior mean of θ using *both* signals. In particular, for any $\eta > 0$,

$$\tilde{\theta}^* \neq \hat{\theta}_B \quad \text{unless} \quad \hat{\theta}_B = \hat{\theta}_O,$$

and $\tilde{\theta}^*$ is pulled toward the published OBR number $\hat{\theta}_O$ ("herding").

Proof. Using the conditional mean–variance decomposition,

$$\mathbb{E}\Big[(\theta - \tilde{\theta})^2 \mid z_{S\theta}, z_O\Big] = \mathrm{Var}(\theta \mid z_{S\theta}, z_O) + \Big(\hat{\theta}_B - \tilde{\theta}\Big)^2.$$

The conditional variance is independent of $\tilde{\theta}$, so the problem reduces to minimising

$$\Big(\hat{\theta}_B - \tilde{\theta}\Big)^2 + \eta(\tilde{\theta} - \hat{\theta}_O)^2.$$

The first-order condition is

$$-2(\hat{\theta}_B - \tilde{\theta}) + 2\eta(\tilde{\theta} - \hat{\theta}_O) = 0,$$

which gives

$$(1 + \eta)\tilde{\theta} = \hat{\theta}_B + \eta\hat{\theta}_O$$

and hence equation (3). Moreover,

$$\tilde{\theta}^* - \hat{\theta}_B = \frac{\eta}{1 + \eta}\,(\hat{\theta}_O - \hat{\theta}_B).$$

For $\eta > 0$, the published baseline therefore lies between the Bayesian posterior mean and the already published OBR forecast, with strict displacement whenever the two differ. □

The result is specific to a sequential calendar: the realised OBR forecast is available both as information and as the reference point in the disagreement penalty. Under simultaneous publication there is no already published number to play this anchoring role.

Proposition 2 (assumptions concordat: accuracy and common-mode exposure). The SFC publishes a forecast of

$$y = \theta + \sigma$$

of the form

$$\hat{y}_S \equiv \tilde{\theta} + \hat{\sigma}_S.$$

Under no concordat,

$$\tilde{\theta} = \hat{\theta}_S, \qquad \hat{y}_S^N = \hat{\theta}_S + \hat{\sigma}_S,$$

whereas under a concordat,

$$\tilde{\theta} = \hat{\theta}_O, \qquad \hat{y}_S^C = \hat{\theta}_O + \hat{\sigma}_S.$$

Maintain independent signal noises and define

$$C_{OS} \equiv \mathrm{Cov}(\theta - \hat{\theta}_O, \theta - \hat{\theta}_S) = \frac{V_O V_S}{V_\theta}$$

by Lemma 1.4 (See the formulas in Part I).

(i) Accuracy. If $V_O \leq V_S$, then

$$\mathbb{E}[(y - \hat{y}_S^C)^2] \leq \mathbb{E}[(y - \hat{y}_S^N)^2].$$

(ii) Additional common-mode exposure. Under no concordat, the two shared-component forecast errors have covariance $C_{OS} > 0$. Under the concordat they coincide and their covariance is V_O. The concordat therefore raises error covariance by

$$V_O - C_{OS} = V_O \left(1 - \frac{V_S}{V_\theta}\right) > 0.$$

(iii) Reconciliation volatility. Let

$$R^N \equiv (y - \hat{y}_S^N) + \lambda(\theta - \hat{\theta}_O), \qquad R^C \equiv (y - \hat{y}_S^C) + \lambda(\theta - \hat{\theta}_O).$$

Then

$$\mathrm{Var}(R^N) = V_S + V_\Sigma + \lambda^2 V_O + 2\lambda C_{OS},$$

and

$$\mathrm{Var}(R^C) = (1+\lambda)^2 V_O + V_\Sigma.$$

Therefore the concordat reduces reconciliation variance if and only if

$$V_O + 2\lambda(V_O - C_{OS}) \le V_S. \tag{4}$$

Proof. Under regime N,

$$y - \hat{y}_S^N = (\theta - \hat{\theta}_S) + (\sigma - \hat{\sigma}_S),$$

whereas under regime C,

$$y - \hat{y}_S^C = (\theta - \hat{\theta}_O) + (\sigma - \hat{\sigma}_S).$$

The θ and σ blocks are independent, so

$$\mathbb{E}[(y - \hat{y}_S^N)^2] = V_S + V_\Sigma, \qquad \mathbb{E}[(y - \hat{y}_S^C)^2] = V_O + V_\Sigma.$$

This proves part (i).

Part (ii) follows directly from Lemma 1.4 under regime N and from the identity of the two shared-component errors under regime C.

Finally,

$$R^N = (\theta - \hat{\theta}_S) + (\sigma - \hat{\sigma}_S) + \lambda(\theta - \hat{\theta}_O).$$

Using

$$\mathrm{Cov}(\theta - \hat{\theta}_S, \theta - \hat{\theta}_O) = C_{OS}$$

and the independence of the σ block gives

$$\mathrm{Var}(R^N) = V_S + V_\Sigma + \lambda^2 V_O + 2\lambda C_{OS}.$$

Under the concordat,

$$R^C = (1+\lambda)(\theta - \hat{\theta}_O) + (\sigma - \hat{\sigma}_S),$$

so

$$\mathrm{Var}(R^C) = (1+\lambda)^2 V_O + V_\Sigma.$$

The condition $\mathrm{Var}(R^C) \leq \mathrm{Var}(R^N)$ is equivalent to

$$(1+2\lambda)V_O - 2\lambda C_{OS} \leq V_S,$$

which is (4). □

Propositions 1 and 2 turn coordination into an explicit protocol-design problem. Sequential publication can anchor the second mover and induce herding. A concordat can improve average accuracy when the shared baseline is more precise, but it also raises common-mode exposure above the positive covariance already generated by the finite common prior. Equation (4) states the resulting system-level trade-off between precision and incremental correlated exposure.

Proposition 3 (calendar under late data: correlation with little additional information). The proposition compares reconciliation variance under simultaneous publication with reconciliation variance under sequential publication, $O \to S$, when the SFC observes the OBR baseline before publishing and faces a disagreement penalty. Retain the benchmark state structure, so that $\theta \sim \mathcal{N}(0, V_\theta)$, $\sigma \sim \mathcal{N}(0, V_\sigma)$, $\theta \perp \sigma$, where the Scottish outcome is $y \equiv \theta + \sigma$. The SFC observes a signal about σ:

$$z_{S\sigma} = \sigma + \varepsilon_{S\sigma}, \qquad \varepsilon_{S\sigma} \sim \mathcal{N}(0, w_S),$$

which is independent of everything else. Let

$$\hat{\sigma}_S \equiv \mathbb{E}[\sigma \mid z_{S\sigma}], \qquad V_\Sigma \equiv \mathrm{Var}(\sigma - \hat{\sigma}_S) = \left(\frac{1}{V_\sigma} + \frac{1}{w_S}\right)^{-1}.$$

The σ component is identical across calendar regimes and therefore drops out of their comparison.

Let the data-arrival indicator be $A \in \{H, L\}$ with

$$\Pr(A = H) = \pi, \qquad \Pr(A = L) = 1 - \pi.$$

Here $A = H$ means key information arrives on time (high-quality signals); $A = L$ means key information is late (low-quality signals). Conditional on $A = a \in \{H, L\}$, the institutions observe independent noisy signals:

$$z_O^a = \theta + \varepsilon_O^a, \qquad \varepsilon_O^a \sim \mathcal{N}(0, v_O^a),$$

$$z_S^a = \theta + \varepsilon_S^a, \qquad \varepsilon_S^a \sim \mathcal{N}(0, v_S^a),$$

with all noise terms independent of each other and of (θ, σ). Assume signal quality is higher when data arrives on time,

$$0 < v_O^H < v_O^L, \qquad 0 < v_S^H < v_S^L.$$

The OBR forecast of θ (posterior mean) is,

$$\hat{\theta}_O^a \equiv \mathbb{E}[\theta \mid z_O^a] = \frac{V_\theta}{V_\theta + v_O^a} z_O^a,$$

with forecast error and variance

$$e_O^a \equiv \theta - \hat{\theta}_O^a, \qquad V_O^a \equiv \mathrm{Var}(e_O^a) = \left(\frac{1}{V_\theta} + \frac{1}{v_O^a}\right)^{-1}.$$

On the other hand, the SFC forecast of y is of the form

$$\hat{y}_S^a \equiv \tilde{\theta}^a + \hat{\sigma}_S,$$

where $\tilde{\theta}^a$ is the published baseline for the θ component. Define the reconciliation amount:

$$R^a \equiv (y - \hat{y}_S^a) + \lambda(\theta - \hat{\theta}_O^a), \qquad \lambda \in \mathbb{R}.$$

Since $y = \theta + \sigma$ and $\hat{y}_S^a = \tilde{\theta}^a + \hat{\sigma}_S$, we can write

$$R^a = (\theta - \tilde{\theta}^a) + (\sigma - \hat{\sigma}_S) + \lambda(\theta - \hat{\theta}_O^a).$$

Under simultaneous publication (regime Sim), the SFC publishes without observing $\hat{\theta}_O^a$ and uses its own θ signal:

$$\tilde{\theta}_{Sim}^a \equiv \hat{\theta}_S^a \equiv \mathbb{E}[\theta \mid z_S^a] = \frac{V_\theta}{V_\theta + v_S^a} z_S^a.$$

Define the SFC θ-error under simultaneity:

$$e_S^a \equiv \theta - \hat{\theta}_S^a, \qquad V_S^a \equiv \mathrm{Var}(e_S^a) = \left(\frac{1}{V_\theta} + \frac{1}{v_S^a}\right)^{-1}.$$

By Lemma 1.4, the two one-signal posterior errors have covariance

$$C_{OS}^a \equiv \mathrm{Cov}(e_O^a, e_S^a) = \frac{V_O^a V_S^a}{V_\theta} > 0.$$

Under sequential publication (regime Seq), the OBR publishes first. The SFC then observes $\hat{\theta}_O^a$ and chooses $\tilde{\theta}^a$ to minimise

$$\mathcal{J}(\tilde{\theta}^a) = \mathbb{E}\Big[(\theta - \tilde{\theta}^a)^2 \mid z_S^a, \hat{\theta}_O^a\Big] + \eta(\tilde{\theta}^a - \hat{\theta}_O^a)^2, \qquad \eta \geq 0.$$

Here η is the reduced-form disagreement penalty.

Three preliminary lemmas collect the Gaussian identities used in the proof.

Lemma 1. Conditional on $A = a$, let SFC observe both signals (z_S^a, z_O^a). Define the Bayesian two-signal posterior mean

$$\hat{\theta}_B^a \equiv \mathbb{E}[\theta \mid z_S^a, z_O^a].$$

Then the posterior variance is

$$V_B^a \equiv \mathrm{Var}(\theta - \hat{\theta}_B^a) = \left(\frac{1}{V_\theta} + \frac{1}{v_S^a} + \frac{1}{v_O^a}\right)^{-1}.$$

Proof. Update sequentially. Start from prior precision $1/V_\theta$. After seeing z_S^a, posterior precision is $1/V_\theta + 1/v_S^a$. After also seeing z_O^a, posterior precision is $1/V_\theta + 1/v_S^a + 1/v_O^a$. The posterior variance is the inverse of posterior precision. □

Lemma 2. Conditional on $A = a$, $\hat{\theta}_O^a$ is an affine (linear) transformation of z_O^a:

$$\hat{\theta}_O^a = \kappa_O^a z_O^a, \qquad \kappa_O^a \equiv \frac{V_\theta}{V_\theta + v_O^a} \in (0, 1).$$

Therefore knowing $\hat{\theta}_O^a$ is informationally equivalent to knowing z_O^a.

Proof. Immediate from the normal–normal updating formula for the posterior mean. Since $\kappa_O^a \neq 0$, one can recover $z_O^a = \hat{\theta}_O^a / \kappa_O^a$. □

Define the Bayesian two-signal forecast error

$$e_B^a \equiv \theta - \hat{\theta}_B^a.$$

Lemma 3. Conditional on $A = a$:

$$\mathrm{Var}(e_B^a) = V_B^a, \qquad \mathrm{Var}(e_O^a) = V_O^a, \qquad \mathrm{Cov}(e_B^a, e_O^a) = V_B^a.$$

Proof. The first two equalities are definitions of V_B^a and V_O^a.

For the covariance:

$$\mathrm{Cov}(e_B^a, e_O^a) = \mathrm{Cov}(\theta - \hat{\theta}_B^a,\ \theta - \hat{\theta}_O^a) = \mathrm{Cov}(\theta - \hat{\theta}_B^a, \theta) - \mathrm{Cov}(\theta - \hat{\theta}_B^a, \hat{\theta}_O^a).$$

We evaluate the two terms.

Step 1. $\mathrm{Cov}(\theta - \hat{\theta}_B^a, \theta) = \mathrm{Var}(\theta - \hat{\theta}_B^a) = V_B^a$. To see this, write $\theta = \hat{\theta}_B^a + e_B^a$, so

$$\mathrm{Cov}(e_B^a, \theta) = \mathrm{Cov}(e_B^a, \hat{\theta}_B^a + e_B^a) = \mathrm{Cov}(e_B^a, \hat{\theta}_B^a) + \mathrm{Var}(e_B^a).$$

But $\hat{\theta}_B^a = \mathbb{E}[\theta \mid z_S^a, z_O^a]$ is the L^2 projection of θ on the information set generated by (z_S^a, z_O^a). Therefore the projection error e_B^a is orthogonal to any function of (z_S^a, z_O^a), including $\hat{\theta}_B^a$. Hence $\mathrm{Cov}(e_B^a, \hat{\theta}_B^a) = 0$, giving

$$\mathrm{Cov}(e_B^a, \theta) = \mathrm{Var}(e_B^a) = V_B^a.$$

Step 2. $\mathrm{Cov}(\theta - \hat{\theta}_B^a, \hat{\theta}_O^a) = 0$. This holds because $\hat{\theta}_O^a$ is a function of z_O^a only (Lemma 2), and z_O^a is contained in the information set (z_S^a, z_O^a). By the same orthogonality property, e_B^a is orthogonal to any function of (z_S^a, z_O^a), and hence orthogonal to $\hat{\theta}_O^a$.

Combining Step 1 and Step 2 yields

$$\mathrm{Cov}(e_B^a, e_O^a) = V_B^a.$$

□

Proof of Proposition 3. Under regime Sim,

$$\tilde{\theta}^a_{Sim} = \hat{\theta}^a_S, \qquad R^a_{Sim} = e^a_S + (\sigma - \hat{\sigma}_S) + \lambda e^a_O.$$

The σ error is independent of the shared-component errors. The two one-signal posterior errors are not independent, however: by the finite-prior covariance lemma,

$$\mathrm{Cov}(e^a_S, e^a_O) = C^a_{OS}.$$

Therefore

$$\mathrm{Var}(R^a_{Sim}) = V^a_S + \lambda^2 V^a_O + 2\lambda C^a_{OS} + V_\Sigma.$$

Under regime Seq, the SFC minimises

$$\mathcal{J}(\tilde{\theta}^a) = \mathbb{E}\Big[(\theta - \tilde{\theta}^a)^2 \mid z^a_S, \hat{\theta}^a_O\Big] + \eta(\tilde{\theta}^a - \hat{\theta}^a_O)^2.$$

By Lemma 2, conditioning on $\hat{\theta}^a_O$ is equivalent to conditioning on z^a_O. Define the Bayesian benchmark

$$\hat{\theta}^a_B \equiv \mathbb{E}[\theta \mid z^a_S, z^a_O].$$

Now apply the same steps as in Proposition 1 of the baseline case:

$$\mathbb{E}\Big[(\theta - \tilde{\theta}^a)^2 \mid z^a_S, z^a_O\Big] = \mathrm{Var}(\theta \mid z^a_S, z^a_O) + (\hat{\theta}^a_B - \tilde{\theta}^a)^2.$$

The conditional variance term does not depend on $\tilde{\theta}^a$. Hence minimising $\mathcal{J}$ is equivalent to minimising

$$(\hat{\theta}^a_B - \tilde{\theta}^a)^2 + \eta(\tilde{\theta}^a - \hat{\theta}^a_O)^2.$$

Differentiate and set to zero:

$$-2(\hat{\theta}^a_B - \tilde{\theta}^a) + 2\eta(\tilde{\theta}^a - \hat{\theta}^a_O) = 0,$$

which implies

$$(1+\eta)\tilde{\theta}^a = \hat{\theta}^a_B + \eta\hat{\theta}^a_O, \qquad \Longrightarrow \qquad \tilde{\theta}^a_{Seq} = \frac{\hat{\theta}^a_B + \eta\hat{\theta}^a_O}{1+\eta}.$$

To express the sequential-regime θ error in terms of e_B^a and e_O^a, define $\alpha \equiv \eta/(1+\eta) \in [0,1)$. Then

$$\tilde{\theta}_{Seq}^a = (1-\alpha)\hat{\theta}_B^a + \alpha\hat{\theta}_O^a.$$

So the SFC θ-error under Seq is

$$e_{Seq}^a \equiv \theta - \tilde{\theta}_{Seq}^a = (1-\alpha)(\theta - \hat{\theta}_B^a) + \alpha(\theta - \hat{\theta}_O^a) = (1-\alpha)e_B^a + \alpha e_O^a.$$

Under Seq, the reconciliation can be written as

$$R_{Seq}^a = e_{Seq}^a + (\sigma - \hat{\sigma}_S) + \lambda e_O^a = (1-\alpha)e_B^a + (\alpha+\lambda)e_O^a + (\sigma - \hat{\sigma}_S).$$

The σ term is independent and adds V_Σ to variance, so

$$\mathrm{Var}(R_{Seq}^a) = \mathrm{Var}\big((1-\alpha)e_B^a + (\alpha+\lambda)e_O^a\big) + V_\Sigma.$$

Expand the variance:

$$\begin{aligned}\mathrm{Var}(R_{Seq}^a) = (1-\alpha)^2\mathrm{Var}(e_B^a) + (\alpha+\lambda)^2\mathrm{Var}(e_O^a) \\ + 2(1-\alpha)(\alpha+\lambda)\mathrm{Cov}(e_B^a, e_O^a) + V_\Sigma.\end{aligned}$$

Now substitute Lemma 3:

$$\mathrm{Var}(R_{Seq}^a) = (1-\alpha)^2 V_B^a + (\alpha+\lambda)^2 V_O^a + 2(1-\alpha)(\alpha+\lambda)V_B^a + V_\Sigma.$$

Collect the V_B^a terms:

$$\begin{aligned}&(1-\alpha)^2 V_B^a + 2(1-\alpha)(\alpha+\lambda)V_B^a \\ &\quad = (1-\alpha)\big[(1-\alpha) + 2(\alpha+\lambda)\big]V_B^a \\ &\quad = (1-\alpha)(1+\alpha+2\lambda)V_B^a.\end{aligned}$$

Since $\alpha = \eta/(1+\eta)$, we have $1-\alpha = 1/(1+\eta)$ and $1+\alpha = (1+2\eta)/(1+\eta)$. Therefore

$$(1-\alpha)(1+\alpha+2\lambda) = \frac{1}{1+\eta}\left(\frac{1+2\eta}{1+\eta} + 2\lambda\right) = \frac{1+2\eta+2\lambda(1+\eta)}{(1+\eta)^2}.$$

Hence the reconciliation variance under Seq is

$$\mathrm{Var}(R^a_{Seq}) = V_\Sigma + (\lambda+\alpha)^2 V^a_O + \frac{1+2\eta+2\lambda(1+\eta)}{(1+\eta)^2}\, V^a_B.$$

This is the explicit formula we will compare to Sim.

Scottish exact-cancellation benchmark. Set $\lambda = -1$. Since $\lambda+\alpha = -1/(1+\eta)$ and $c(-1,\eta) = -1/(1+\eta)^2$,

$$\mathrm{Var}(R^a_{Seq}) = V_\Sigma + \frac{V^a_O - V^a_B}{(1+\eta)^2}. \tag{5}$$

Under simultaneity,

$$\mathrm{Var}(R^a_{Sim}) = V_\Sigma + V^a_S + V^a_O - 2C^a_{OS}. \tag{6}$$

To compare them, write prior precision as $p \equiv 1/V_\theta$, downstream signal precision as $s^a \equiv 1/v^a_S$, and upstream signal precision as $o^a \equiv 1/v^a_O$. Then

$$V^a_S + V^a_B - 2C^a_{OS} = \frac{o^a(o^a+p+2s^a)}{(o^a+p)(p+s^a)(o^a+p+s^a)} > 0.$$

Moreover, $V^a_O - V^a_B > 0$ and $(1+\eta)^{-2} \le 1$. Hence

$$\begin{aligned}&\mathrm{Var}(R^a_{Sim}) - \mathrm{Var}(R^a_{Seq})\\ &\quad = \left(V^a_S + V^a_B - 2C^a_{OS}\right) + \left(1 - \frac{1}{(1+\eta)^2}\right)(V^a_O - V^a_B) > 0\end{aligned}$$

for non-degenerate signals. Thus sequencing always lowers reconciliation variance when $\lambda = -1$. The state-dependent reversal in Proposition 3 is therefore a general-interface possibility, not a feature of the exact Scottish cash-arithmetic benchmark. In Scotland the relevant trade-off is between this cash stabilisation and the amount of additional information preserved.

Define the conditional variance difference

$$\Delta^a \equiv \mathrm{Var}(R^a_{Seq}) - \mathrm{Var}(R^a_{Sim}).$$

Using the formulas above:

$$\Delta^a = \left[(\lambda+\alpha)^2 - \lambda^2\right] V^a_O + \frac{1+2\eta+2\lambda(1+\eta)}{(1+\eta)^2} V^a_B - V^a_S - 2\lambda C^a_{OS}.$$

Since

$$(\lambda+\alpha)^2-\lambda^2=2\lambda\alpha+\alpha^2, \qquad \alpha=\frac{\eta}{1+\eta},$$

we obtain

$$\Delta^a=(2\lambda\alpha+\alpha^2)V_O^a+c(\lambda,\eta)V_B^a-V_S^a-2\lambda C_{OS}^a,$$

where

$$c(\lambda,\eta)\equiv\frac{1+2\eta+2\lambda(1+\eta)}{(1+\eta)^2}.$$

The final term reflects the positive covariance already induced by the finite common prior. The wedge therefore measures the *change* in common exposure generated by sequencing and the disagreement penalty, rather than comparing sequential covariance with a zero-covariance benchmark.

The ex ante variance difference is

$$\Delta(\pi)\equiv\mathbb{E}\big[\mathrm{Var}(R_{Seq}^A)-\mathrm{Var}(R_{Sim}^A)\big]=\pi\Delta^H+(1-\pi)\Delta^L.$$

This is an affine function of π.

Claim (existence of a threshold). If $\Delta^L>0$ and $\Delta^H<0$, then there exists a unique $\bar{\pi}\in(0,1)$ such that:

$$\Delta(\pi)>0 \text{ for } \pi<\bar{\pi}, \qquad \Delta(\pi)<0 \text{ for } \pi>\bar{\pi}.$$

Moreover,

$$\bar{\pi}=\frac{\Delta^L}{\Delta^L-\Delta^H}.$$

Proof. Solve $\Delta(\pi)=0$:

$$\pi\Delta^H+(1-\pi)\Delta^L=0 \quad\Longleftrightarrow\quad \pi(\Delta^H-\Delta^L)=-\Delta^L \quad\Longleftrightarrow\quad \pi=\frac{\Delta^L}{\Delta^L-\Delta^H}.$$

Since $\Delta^L>0$ and $\Delta^H<0$, the denominator $\Delta^L-\Delta^H>0$ and the fraction lies in $(0,1)$. Because $\Delta(\pi)$ is linear in π, the sign pattern follows. □

The sign configuration is not imposed mechanically. The following primitive condition shows that it arises on a non-empty parameter region. Fix $\eta > 0$ and let $\alpha = \eta/(1+\eta)$. Suppose

$$\lambda > -\frac{\alpha}{2} = -\frac{\eta}{2(1+\eta)}, \tag{7}$$

and let both downstream signals be informative, so $0 < v_S^H < v_S^L < \infty$. Holding those downstream variances fixed, consider the upstream signal variances. As $v_O^H \downarrow 0$,

$$V_O^H \to 0, \qquad V_B^H \to 0, \qquad C_{OS}^H \to 0,$$

and hence

$$\Delta^H \to -V_S^H < 0.$$

As $v_O^L \uparrow \infty$,

$$V_O^L \to V_\theta, \qquad V_B^L \to V_S^L, \qquad C_{OS}^L \to V_S^L.$$

Substitution into the wedge gives

$$\Delta^L \longrightarrow \frac{\eta\,[\eta + 2\lambda(1+\eta)]}{(1+\eta)^2}\big(V_\theta - V_S^L\big) > 0, \tag{8}$$

where the inequality follows from (7) and $V_S^L < V_\theta$. By continuity, sufficiently precise upstream information in the high-readiness state and sufficiently imprecise upstream information in the low-readiness state imply $\Delta^H < 0 < \Delta^L$. Economically, when the upstream signal is very precise, its informational contribution dominates; when it becomes nearly uninformative, a positive disagreement penalty can still pull the downstream publication toward it and raise reconciliation variance. This sufficient condition is not necessary, but it establishes that the threshold result is supported by a transparent set of primitives rather than by the sign assumptions alone.

When $\Delta^L > 0$ and $\Delta^H < 0$, simultaneity therefore dominates for sufficiently small π. To isolate the informational channel, define the pure Bayesian gain from adding z_O^a to the SFC information set:

$$G^a \equiv V_S^a - V_B^a.$$

Let

$$A^a \equiv \frac{1}{V_\theta} + \frac{1}{v_S^a}, \qquad C^a \equiv \frac{1}{v_O^a}.$$

Then

$$G^a = \frac{1}{A^a} - \frac{1}{A^a + C^a} = \frac{C^a}{A^a(A^a + C^a)}.$$

The gain is increasing in upstream precision C^a:

$$\frac{\partial G}{\partial C} = \frac{1}{(A + C)^2} > 0,$$

but decreasing in the precision already available to the SFC:

$$\frac{\partial G}{\partial A} = -\frac{C(2A + C)}{A^2(A + C)^2} < 0.$$

Because both institutions' signals are more precise in the high-readiness state, the ordering $G^H > G^L$ is not automatic. It holds if and only if

$$\frac{C^H}{A^H(A^H + C^H)} > \frac{C^L}{A^L(A^L + C^L)}. \tag{9}$$

Thus late data reduce the informational value of sequencing when the upstream loss of precision is sufficiently large relative to the loss of precision in the SFC's own signal. The threshold result above does not require (9); it requires the state-specific signs $\Delta^L > 0$ and $\Delta^H < 0$.

The proof separates the two effects of sequencing. It allows the SFC to learn from the upstream forecast, but it also changes the covariance of the two errors, and a positive disagreement penalty pulls the downstream publication further toward the upstream baseline. The simultaneous benchmark already contains positive finite-prior covariance C_{OS}^a; sequencing adds or subtracts from that benchmark depending on the information and exposure terms in the wedge. Simultaneous publication can therefore minimise reconciliation volatility when the state-specific inequalities in the proposition are satisfied.

2.1 Private sharing of upstream precision-status information

In the baseline Gaussian model with known variances, the upstream point forecast is an affine transformation of its signal and therefore reveals the

upstream mean information. The sharing rule ρ has additional content when the reliability of that baseline varies across forecast rounds and the downstream institution cannot infer the relevant precision status from the point forecast alone. Examples include data-vintage flags, costing-readiness information and nowcast diagnostics.

Let $s \in \{G, B\}$ denote the upstream precision status, with

$$0 < v_O(G) < v_O(B).$$

The common state and signals are

$$\theta \sim \mathcal{N}(0, V_\theta), \qquad z_O = \theta + \varepsilon_O, \qquad z_S = \theta + \varepsilon_S,$$

where

$$\varepsilon_O \mid s \sim \mathcal{N}(0, v_O(s)), \qquad \varepsilon_S \sim \mathcal{N}(0, v_S),$$

and the signal noises are independent conditional on s and independent of θ. Let

$$\hat{\theta}_O(s) \equiv \mathbb{E}[\theta \mid z_O, s], \qquad \hat{\theta}_S \equiv \mathbb{E}[\theta \mid z_S],$$

with posterior error variances

$$V_O(s) = \left(\frac{1}{V_\theta} + \frac{1}{v_O(s)}\right)^{-1}, \qquad V_S = \left(\frac{1}{V_\theta} + \frac{1}{v_S}\right)^{-1}.$$

The finite-prior covariance lemma (i.e., Lemma 1.4 in Part I) gives

$$C_{OS}(s) \equiv \mathrm{Cov}\Big(\theta - \hat{\theta}_O(s), \theta - \hat{\theta}_S \mid s\Big) = \frac{V_O(s) V_S}{V_\theta} > 0. \tag{10}$$

To obtain a transparent closed form, restrict attention to linear pooling:

$$\tilde{\theta}(\omega) = (1 - \omega)\hat{\theta}_S + \omega \hat{\theta}_O(s).$$

This restriction is not needed for the qualitative value-of-information result, but it makes clear how precision-status sharing changes the efficient reliance weight. Conditional on s, the mean-squared error is

$$L(\omega; s) = (1 - \omega)^2 V_S + \omega^2 V_O(s) + 2\omega(1 - \omega) C_{OS}(s). \tag{11}$$

The cross term is important–independent signal noises do not make the two posterior errors independent.

Define

$$D(s) \equiv V_S + V_O(s) - 2C_{OS}(s) > 0.$$

Then (11) can be written

$$L(\omega; s) = V_S + 2\omega\big(C_{OS}(s) - V_S\big) + \omega^2 D(s).$$

The state-contingent accuracy-minimising weight is therefore

$$\omega^*(s) = \frac{V_S - C_{OS}(s)}{V_S + V_O(s) - 2C_{OS}(s)}. \tag{12}$$

Because $C_{OS}(s) < \min\{V_S, V_O(s)\}$, this weight lies strictly between zero and one. The minimum conditional loss is

$$L^*(s) = V_S - \frac{\big(V_S - C_{OS}(s)\big)^2}{V_S + V_O(s) - 2C_{OS}(s)} = \frac{V_S V_O(s) - C_{OS}(s)^2}{V_S + V_O(s) - 2C_{OS}(s)}. \tag{13}$$

Under $\rho = 1$, the downstream institution observes s and chooses $\omega^*(s)$ state by state. Expected loss is

$$L^{\rho=1} = \mathbb{E}[L^*(s)].$$

Under $\rho = 0$, it must choose one constant weight. Let

$$\bar{V}_O \equiv \mathbb{E}[V_O(s)], \qquad \bar{C}_{OS} \equiv \mathbb{E}[C_{OS}(s)].$$

The optimal constant weight is

$$\bar{\omega} = \frac{V_S - \bar{C}_{OS}}{V_S + \bar{V}_O - 2\bar{C}_{OS}}, \tag{14}$$

with expected loss

$$L^{\rho=0} = V_S - \frac{(V_S - \bar{C}_{OS})^2}{V_S + \bar{V}_O - 2\bar{C}_{OS}}. \tag{15}$$

Proposition (value of precision-status sharing). Under the linear-pooling restriction,

$$L^{\rho=1} \leq L^{\rho=0}.$$

The inequality is strict whenever the optimal reliance weight differs across precision states with positive probability.

Proof. A constant weight is feasible when the institution can condition on s. Therefore, state by state,

$$\min_{\omega} L(\omega; s) \leq L(\bar{\omega}; s).$$

Taking expectations gives

$$\mathbb{E}[\min_{\omega} L(\omega; s)] \leq \mathbb{E}[L(\bar{\omega}; s)] = \min_{\omega} \mathbb{E}[L(\omega; s)].$$

Strict inequality obtains when the state-contingent minimisers are not almost surely equal. □

Interpretation. The value of ρ is not the value of transmitting a second copy of the upstream point forecast. It is the value of revealing how much confidence should be placed in that forecast in the current round. The formulas also sharpen the economic interpretation. Even without status sharing, the two posterior errors share finite-prior covariance. Sharing the status improves the downstream allocation of weight between two already correlated estimates.

If the downstream institution performs unrestricted Bayesian updating, observing the precision state likewise cannot reduce forecast quality: it enlarges the information set. Without the status flag, the posterior is generally a mixture and need not admit the simple linear formula above. The linear-pooling version is retained because it isolates the protocol margin in closed form while retaining the covariance implied by the finite common prior.

3 Part III. The Scotland–UK Reconciliation Mapping

This Part maps the reduced-form reconciliation equation used in the paper into the Scotland–UK fiscal framework. It first distinguishes the SFC's forecasting mandate from the cash-management objective generated by subsequent reconciliations. It then establishes two cases in which those objectives are compatible and derives the negative reconciliation loading, $\lambda = -1$, implied by the Block Grant Adjustment (BGA). The final corollary shows that a concordat on the shared component removes that component from the reconciliation arithmetic altogether.

3.1 Forecast accuracy and reconciliation stability

An instruction that the SFC produce the "best forecasts it can" naturally defines an accuracy objective. Conditional on policy and the information available to it, the SFC chooses forecasts that minimise expected loss—typically squared error—for Scottish receipts and the related aggregates within its remit. The mandate is for the SFC to produce the most accurate and informative baseline available to Parliament and the public.

Reconciliation volatility is a distinct object. For a devolved tax, the adjustment to the Scottish budget can be written as

$$R = (T_S - \widehat{T}_S) - (BGA - \widehat{BGA}).$$

It depends on two forecast errors and, crucially, on their covariance. Minimising the forecast error in T_S therefore need not minimise the volatility of R.

The fiscal framework asks the system to achieve both objectives. It requires an independent institution to produce a credible Scottish baseline, while also exposing the Scottish budget to the net cash adjustment generated

when Scottish receipts and the BGA are reconciled. In the decomposition

$$y = \theta + \sigma, \qquad BGA = \theta,$$

θ is the shared component and σ is Scotland-specific. Forecast accuracy directs the SFC to use its best estimate of both components. Cash stability instead favours alignment with the OBR on θ, because common treatment of that component makes it cancel from the reconciliation.

The two objectives can nevertheless coincide. They do so when the OBR provides at least as precise an estimate of the shared component as the SFC's alternative source, or when an assumptions concordat defines the OBR baseline as part of the conditioning information on which the SFC is required to base its forecast. In the latter case, alignment is not an ex post distortion of the SFC forecast; it is an explicit feature of the forecasting problem.

3.2 Formal compatibility of the two objectives

Let

$$y \equiv \theta + \sigma, \qquad B \equiv \theta,$$

where y denotes Scottish receipts (or the devolved tax base), θ is the shared component that drives the BGA, and σ is Scotland-specific. The SFC publishes

$$\widehat{y}_S \equiv \widehat{\theta} + \widehat{\sigma}_S,$$

where $\widehat{\sigma}_S$ is its estimate of σ. The OBR forecast of the BGA object is

$$\widehat{B} \equiv \widehat{\theta}_O.$$

Define

$$e_S \equiv y - \widehat{y}_S, \qquad e_B \equiv B - \widehat{B} = \theta - \widehat{\theta}_O,$$

and the reconciliation of the net position as

$$R \equiv e_S - e_B = (y - \widehat{y}_S) - (\theta - \widehat{\theta}_O).$$

Under quadratic loss, the best forecast of any random variable X given information $\mathcal{I}$ is $\mathbb{E}[X \mid \mathcal{I}]$.

Consider two objectives. The SFC's accuracy mandate chooses $\widehat{\theta}$ to minimise

$$\mathbb{E}\big[(y - \widehat{y}_S)^2\big] = \mathbb{E}\big[(\theta - \widehat{\theta})^2\big] + \mathbb{E}\big[(\sigma - \widehat{\sigma}_S)^2\big],$$

where the second term is fixed with respect to $\widehat{\theta}$. The cash-management objective chooses $\widehat{\theta}$ to minimise $\mathrm{Var}(R)$, equivalently $\mathbb{E}[R^2]$ when $\mathbb{E}[R] = 0$.

Proposition (compatibility of objectives). In the setup above, the accuracy mandate and the cash-management objective select the same $\widehat{\theta}$ in either of the following cases.

1. *Precision dominance.* The SFC observes the OBR information about θ and must choose between posterior means based separately on

   $$z_O = \theta + \varepsilon_O, \quad \varepsilon_O \sim \mathcal{N}(0, v_O), \qquad z_S = \theta + \varepsilon_S, \quad \varepsilon_S \sim \mathcal{N}(0, v_S),$$

   under the common prior $\theta \sim \mathcal{N}(0, V_\theta)$:

   $$\widehat{\theta} \in \{\mathbb{E}[\theta \mid z_O], \mathbb{E}[\theta \mid z_S]\}.$$

   If $v_O \leq v_S$, an accuracy-minimising choice is

   $$\widehat{\theta} = \mathbb{E}[\theta \mid z_O] = \widehat{\theta}_O,$$

   which is also the reconciliation-stabilising choice. The choice is unique when $v_O < v_S$.

2. *Concordat as an agreed conditioning set.* The concordat forms part of the conditioning information for the SFC forecast: the permitted information set contains z_O for the shared component together with the SFC's information $\mathcal{I}_\sigma$ about σ. The SFC's best forecast is then

   $$\widehat{y}_S^C = \mathbb{E}[y \mid z_O, \mathcal{I}_\sigma] = \mathbb{E}[\theta \mid z_O] + \mathbb{E}[\sigma \mid \mathcal{I}_\sigma] = \widehat{\theta}_O + \widehat{\sigma}_S.$$

   Alignment on θ is therefore part of the stated forecasting problem rather than an ex post distortion.

Proof

Using $y = \theta + \sigma$ and $\widehat{y}_S = \widehat{\theta} + \widehat{\sigma}_S$ gives

$$\begin{aligned} R &= (y - \widehat{y}_S) - (\theta - \widehat{\theta}_O) \\ &= (\sigma - \widehat{\sigma}_S) + (\widehat{\theta}_O - \widehat{\theta}). \end{aligned}$$

Hence

$$\begin{aligned} \mathrm{Var}(R) = \mathrm{Var}(\sigma - \widehat{\sigma}_S) + \mathrm{Var}(\widehat{\theta}_O - \widehat{\theta}) \\ + 2\,\mathrm{Cov}\Big(\sigma - \widehat{\sigma}_S, \widehat{\theta}_O - \widehat{\theta}\Big). \end{aligned}$$

Under the maintained decomposition, the Scotland-specific forecast error is independent of the shared-component information used to form $\widehat{\theta}$. The covariance term is therefore zero, and the cash-management objective is minimised by $\widehat{\theta} = \widehat{\theta}_O$.

For the precision-dominance case, the minimum mean squared error attainable from a single Gaussian signal with noise variance v is its posterior variance:

$$\min_{a(z)} \mathbb{E}\big[(\theta - a(z))^2\big] = \mathrm{Var}(\theta \mid z) = \left(\frac{1}{V_\theta} + \frac{1}{v}\right)^{-1}.$$

Writing

$$V(v) \equiv \left(\frac{1}{V_\theta} + \frac{1}{v}\right)^{-1},$$

$V(v)$ is increasing in v. Thus $v_O \leq v_S$ implies $V(v_O) \leq V(v_S)$, so the OBR-based posterior mean is accuracy-minimising among the two feasible estimators. It is also the cash-stabilising choice established above.

For the conditioning-set case, let $\mathcal{I}^C$ consist of z_O for the shared component and $\mathcal{I}_\sigma$ for the Scotland-specific component. Quadratic loss implies

$$\widehat{y}_S^C = \mathbb{E}[\theta + \sigma \mid z_O, \mathcal{I}_\sigma].$$

Under independence of θ and σ and separability of the information sets,

$$\mathbb{E}[\theta + \sigma \mid z_O, \mathcal{I}_\sigma] = \mathbb{E}[\theta \mid z_O] + \mathbb{E}[\sigma \mid \mathcal{I}_\sigma] = \widehat{\theta}_O + \widehat{\sigma}_S.$$

The agreed conditioning set therefore makes alignment on the shared component consistent with the SFC's accuracy mandate. □

Outside these two cases, the objectives need not coincide. In particular, if $v_O > v_S$, importing the OBR shared component can worsen the accuracy of $\widehat{y}_S$ while stabilising R by construction. A governance trade-off then arises unless the concordat is stated transparently as part of the conditioning information rather than produced through informal pressure on published outputs.

3.3 The Scottish reconciliation loading

Retain the decomposition

$$y \equiv \theta + \sigma, \qquad B \equiv \theta,$$

and let the institutions publish

$$\widehat{y}_S = \widehat{\theta}_S + \widehat{\sigma}_S, \qquad \widehat{B} = \widehat{\theta}_O.$$

The corresponding forecast errors are

$$e_S \equiv y - \widehat{y}_S, \qquad e_B \equiv B - \widehat{B}.$$

For devolved taxes, reconciliation is outturn net of the BGA minus its forecast counterpart:

$$R \equiv (y - \widehat{y}_S) - (B - \widehat{B}) = e_S - e_B. \tag{16}$$

Thus the generic reduced form $R = e_S + \lambda e_B$ has

$$\lambda = -1.$$

The negative loading means that positive comovement between the two forecast errors stabilises the net Scottish budget position. The following corollary gives the stronger result generated by a shared-component concordat.

Corollary (shared-component cancellation under $\lambda = -1$). Assume that the SFC and OBR forecasts are formed in the Gaussian linear environment described in the paper. Consider an assumptions concordat that imposes

$$\widehat{\theta}_S \equiv \widehat{\theta}_O. \tag{17}$$

Then:

1. the reconciliation reduces to the Scotland-specific forecast error,

$$R = \sigma - \widehat{\sigma}_S,$$

   so the shared component θ drops out entirely;

2. relative to any regime in which $\widehat{\theta}_S \neq \widehat{\theta}_O$ with positive probability, the concordat weakly reduces reconciliation variance,

$$\mathrm{Var}(R) \text{ under concordat } \leq \mathrm{Var}(R) \text{ under no concordat},$$

   with strict inequality whenever $\widehat{\theta}_O - \widehat{\theta}_S$ is non-degenerate and uncorrelated with $\sigma - \widehat{\sigma}_S$.

Proof

Using $y = \theta + \sigma$, $B = \theta$, $\widehat{y}_S = \widehat{\theta}_S + \widehat{\sigma}_S$ and $\widehat{B} = \widehat{\theta}_O$,

$$\begin{aligned} R &= (y - \widehat{y}_S) - (B - \widehat{B}) \\ &= \big((\theta + \sigma) - (\widehat{\theta}_S + \widehat{\sigma}_S)\big) - \big(\theta - \widehat{\theta}_O\big) \\ &= (\sigma - \widehat{\sigma}_S) + (\widehat{\theta}_O - \widehat{\theta}_S). \end{aligned}$$

Under the concordat, the second term vanishes, proving part 1.

Let R^C and R^N denote reconciliation under the concordat and without it. Then

$$R^C = \sigma - \widehat{\sigma}_S, \qquad R^N = (\sigma - \widehat{\sigma}_S) + (\widehat{\theta}_O - \widehat{\theta}_S).$$

It follows that

$$\begin{aligned}\mathrm{Var}(R^N) - \mathrm{Var}(R^C) = \mathrm{Var}(\widehat{\theta}_O - \widehat{\theta}_S) \\ + 2\,\mathrm{Cov}\Big(\sigma - \widehat{\sigma}_S, \widehat{\theta}_O - \widehat{\theta}_S\Big).\end{aligned}$$

In the benchmark decomposition, the Scotland-specific forecast error is independent of the shared-component signal errors, so the covariance term is zero. Therefore

$$\mathrm{Var}(R^N) - \mathrm{Var}(R^C) = \mathrm{Var}(\widehat{\theta}_O - \widehat{\theta}_S) \geq 0,$$

with strict inequality when the difference between the two shared-component forecasts is non-degenerate. □

The compatibility proposition and the cancellation corollary establish different claims. Compatibility concerns whether the SFC's accuracy mandate itself selects the OBR shared-component forecast. Cancellation concerns the reconciliation arithmetic once both institutions use the same shared-component forecast. The second result does not require the first: alignment always removes θ from the reconciliation, although it need not be the SFC's unconstrained accuracy-minimising choice. This distinction is central to the governance argument in the paper.

4 Part IV. Accuracy and Reconciliation Stability

This Part distinguishes two objectives, mentioned earlier, and that are often conflated in discussions of devolved fiscal forecasting: minimising the error in an institution's own forecast and minimising the volatility of the subsequent cash reconciliation. Once forecast errors feed into fiscal adjustments, the relevant object is the joint distribution of errors rather than either error considered in isolation.

Recall the reduced-form reconciliation object

$$R_t \equiv (y_t - \hat{y}_{S,t}) + \lambda(\theta_t - \hat{\theta}_{O,t}).$$

Define

$$e_{S,t} \equiv y_t - \hat{y}_{S,t}, \qquad e_{O,t} \equiv \theta_t - \hat{\theta}_{O,t},$$

so that

$$R_t = e_{S,t} + \lambda e_{O,t} = (y_t - \hat{y}_{S,t}) + \lambda(\theta_t - \hat{\theta}_{O,t}). \tag{18}$$

Its variance is

$$\mathrm{Var}(R_t) = \mathrm{Var}(e_{S,t}) + \lambda^2 \mathrm{Var}(e_{O,t}) + 2\lambda\,\mathrm{Cov}(e_{S,t}, e_{O,t}).$$

Publication calendars, assumptions concordats and information-sharing arrangements can therefore affect reconciliation risk through both marginal forecast variances and the covariance of forecast errors.

We distinguish two conditional loss functions. The devolved IFI's forecasting mandate selects

$$\hat{y}_{S,t}^{acc} \in \arg\min_{\hat{y} \in \mathcal{I}_{S,t}} \mathbb{E}\big[(y_t - \hat{y})^2 \mid \mathcal{I}_{S,t}\big]. \tag{19}$$

By contrast, a government concerned with cash management would select

$$\hat{y}_{S,t}^{stab} \in \arg\min_{\hat{y} \in \mathcal{I}_{S,t}} \mathbb{E}\Big[R_t^2 \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big], \tag{20}$$

where the central forecast is observed at the time of the devolved publication. The first objective depends only on the devolved forecast error. The second depends on both forecast errors and their covariance.

Forecasting as a hedge

Lemma 1 (Stabilisation uses the forecast as a hedge). If $\hat{\theta}_{O,t}$ is observed by the devolved forecaster when choosing $\hat{y}_{S,t}$, then under (20) the optimal publication satisfies

$$\hat{y}_{S,t}^{stab} + \lambda\,\hat{\theta}_{O,t} = \mathbb{E}\Big[y_t + \lambda\theta_t \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big]. \tag{21}$$

Equivalently,

$$R_t^{stab} = \big(y_t + \lambda\theta_t\big) - \mathbb{E}\Big[y_t + \lambda\theta_t \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big], \tag{22}$$

so the reconciliation is the conditional-mean-zero residual for the net object $y_t + \lambda\theta_t$.

Proof. Equation (18) can be written as

$$R_t = (y_t + \lambda\theta_t) - (\hat{y}_{S,t} + \lambda\hat{\theta}_{O,t}).$$

Conditional on $\hat{\theta}_{O,t}$, choosing $\hat{y}_{S,t}$ is equivalent to choosing the second term in parentheses. Under quadratic loss, the optimal value is the conditional expectation of $y_t + \lambda\theta_t$. This gives (21); substitution yields (22). □

The lemma makes the wedge between the two objectives transparent. Under an accuracy mandate, the devolved forecast targets y_t. Under a stability objective, it also absorbs part of the central forecast error because the relevant target is the net fiscal object $y_t + \lambda\theta_t$. The devolved publication therefore acts partly as a hedge.

Gaussian leaning toward the central baseline

To make the hedging channel explicit, suppose

$$\begin{aligned}
\theta_t &\sim \mathcal{N}(0, V_\theta), \\
z_{S,t} &= \theta_t + \varepsilon_{S,t}, \qquad \varepsilon_{S,t} \sim \mathcal{N}(0, v_S), \\
z_{O,t} &= \theta_t + \varepsilon_{O,t}, \qquad \varepsilon_{O,t} \sim \mathcal{N}(0, v_O), \\
\sigma_t &\perp \theta_t,\ z_{S,t},\ z_{O,t},
\end{aligned} \tag{23}$$

with independence across t. The central institution publishes

$$\hat{\theta}_{O,t} = \mathbb{E}[\theta_t \mid z_{O,t}] = \frac{V_\theta}{V_\theta + v_O}\, z_{O,t}. \tag{24}$$

The devolved forecaster observes $(z_{S,t}, \hat{\theta}_{O,t})$ before publishing and uses the conditional-mean forecast of σ_t.

Lemma 2 (Linear-combination form and the leaning weight). Under (23)–(24), the stabilisation-optimal shared-component publication can be written as

$$\hat{\theta}_{S,t}^{stab} = (1+\lambda)\, w_S\, \hat{\theta}_{S,t}^{acc} \;+\; \omega(\lambda; V_\theta, v_S, v_O)\, \hat{\theta}_{O,t}, \tag{25}$$

where

$$w_S \equiv \frac{v_O\,(V_\theta + v_S)}{v_S v_O + V_\theta(v_S + v_O)}, \tag{26}$$

$$\omega(\lambda; V_\theta, v_S, v_O) \equiv (1+\lambda)\,\frac{v_S\,(V_\theta + v_O)}{v_S v_O + V_\theta(v_S + v_O)} - \lambda. \tag{27}$$

and

$$\hat{\theta}_{S,t}^{acc} \equiv \mathbb{E}[\theta_t \mid z_{S,t}] = \frac{V_\theta}{V_\theta + v_S} z_{S,t}.$$

Moreover,

$$\frac{\partial \omega}{\partial \lambda} = \frac{v_S\,(V_\theta + v_O)}{v_S v_O + V_\theta(v_S + v_O)} - 1 < 0. \tag{28}$$

Thus the weight on the central baseline rises as λ becomes more negative. The two coefficients in (25) sum to one. For $-1 \le \lambda \le 0$, including the Scottish benchmark, they are non-negative, so the expression is a genuine weighted average. For unrestricted $\lambda \in \mathbb{R}$, it is more generally a linear combination. A negative reconciliation loading makes positive error covariance stabilising and induces greater leaning toward the central forecast; a positive loading makes that covariance destabilising and induces less leaning.

Proof. Under $y_t = \theta_t + \sigma_t$, with $\hat{\sigma}_{S,t}$ set to its conditional mean, the stability problem reduces to choosing $\hat{\theta}_{S,t}$ to minimise

$$\mathbb{E}\left[\left((1+\lambda)\theta_t - \hat{\theta}_{S,t} - \lambda\hat{\theta}_{O,t}\right)^2 \mid z_{S,t}, \hat{\theta}_{O,t}\right].$$

Hence

$$\hat{\theta}_{S,t}^{stab} + \lambda\hat{\theta}_{O,t} = (1+\lambda)\,\mathbb{E}[\theta_t \mid z_{S,t}, \hat{\theta}_{O,t}].$$

Because (24) is one-to-one in $z_{O,t}$, observing $\hat{\theta}_{O,t}$ is equivalent to observing $z_{O,t}$. The Gaussian posterior mean is

$$\mathbb{E}[\theta_t \mid z_{S,t}, z_{O,t}] = \frac{\frac{1}{v_S} z_{S,t} + \frac{1}{v_O} z_{O,t}}{\frac{1}{V_\theta} + \frac{1}{v_S} + \frac{1}{v_O}}.$$

Using

$$z_{S,t} = \frac{V_\theta + v_S}{V_\theta}\hat{\theta}^{acc}_{S,t}, \qquad z_{O,t} = \frac{V_\theta + v_O}{V_\theta}\hat{\theta}_{O,t},$$

and substituting into the stabilisation condition gives (25) with the coefficients in (26). Differentiating ω yields (28). The strict inequality follows because

$$v_S(V_\theta + v_O) < v_S v_O + V_\theta(v_S + v_O).$$

□

The two lemmas together deliver the main implication. Forecast accuracy and reconciliation stability are distinct objectives even under common quadratic loss. The stability objective generally induces a deliberate covariance between devolved and central forecast errors. Whether that covariance reduces or increases fiscal volatility depends on the sign of λ; in the Scotland–UK mapping of Part III, $\lambda = -1$, so alignment on the shared component is stabilising in cash terms.

5 Part V. Calendar Capacity and Endogenous Importing

This Part introduces a second route to forecast alignment. In Proposition 3, late-arriving information limits the informational gain from sequencing. Here the information is available, but a compressed interval between final policy costings and publication restricts the devolved institution's capacity to process it. The institution must then allocate scarce analytical effort across the shared component, θ, and the devolved-specific component, σ. It may be optimal to import the central forecast of θ and concentrate its own resources on σ. Alignment is therefore an endogenous response to a binding capacity constraint rather than reputational herding. The associated governance instruments are calendar spacing, earlier costings and adequate resourcing.

5.1 Primitives

Each period contains two independent components,

$$\theta_t \sim \mathcal{N}(0, V_\theta), \qquad \sigma_t \sim \mathcal{N}(0, V_\sigma), \qquad \theta_t \perp \sigma_t,$$

and the devolved forecast target is

$$y_t \equiv \theta_t + \sigma_t.$$

The central institution O publishes a forecast of the shared component,

$$\hat{\theta}_{O,t} = \theta_t - u_{O,t}, \qquad u_{O,t} \sim \mathcal{N}(0, V_O),$$

where $V_O > 0$ is the reduced-form mean-squared error of the central forecast of θ_t.[1]

[1] We take V_O as an exogenous summary of the central institution's precision on the shared component. Endogenising it through the signal structure below would not change the mechanism.

The devolved institution S receives private signals

$$z_{S\theta,t} = \theta_t + \varepsilon_{S\theta,t}, \qquad \varepsilon_{S\theta,t} \sim \mathcal{N}\big(0, v_\theta(e_{\theta,t}, \tau_t)\big),$$
$$z_{S\sigma,t} = \sigma_t + \varepsilon_{S\sigma,t}, \qquad \varepsilon_{S\sigma,t} \sim \mathcal{N}\big(0, w_\sigma(e_{\sigma,t}, \tau_t)\big),$$

independent of each other and of (θ_t, σ_t). Here $\tau_t \geq 0$ is the time available between receipt of final Scottish policy costings and the devolved forecast deadline. The effort choices

$$e_{\theta,t} \geq 0, \qquad e_{\sigma,t} \geq 0$$

represent deployable analytical resources and satisfy

$$e_{\theta,t} + e_{\sigma,t} \leq \bar{E}.$$

Signal noise falls with both effort and processing time:

$$\frac{\partial v_\theta}{\partial e_\theta} < 0, \quad \frac{\partial v_\theta}{\partial \tau} < 0, \qquad \frac{\partial w_\sigma}{\partial e_\sigma} < 0, \quad \frac{\partial w_\sigma}{\partial \tau} < 0.$$

For explicit calculations, we use

$$v_\theta(e_\theta, \tau) = \frac{\bar{v}_\theta}{1 + \alpha_\theta e_\theta \tau}, \qquad w_\sigma(e_\sigma, \tau) = \frac{\bar{w}_\sigma}{1 + \alpha_\sigma e_\sigma \tau}, \qquad \alpha_\theta, \alpha_\sigma > 0. \tag{29}$$

5.2 Posterior forecasts and attainable precision

Conditional on $(e_{\theta,t}, \tau_t)$, the posterior-mean forecast of θ_t is

$$\hat{\theta}_{S,t} \equiv \mathbb{E}[\theta_t \mid z_{S\theta,t}] = \frac{V_\theta}{V_\theta + v_\theta(e_{\theta,t}, \tau_t)} z_{S\theta,t},$$

with posterior variance

$$V_\theta^S(e_{\theta,t}, \tau_t) \equiv \mathrm{Var}(\theta_t - \hat{\theta}_{S,t} \mid z_{S\theta,t}) = \left(\frac{1}{V_\theta} + \frac{1}{v_\theta(e_{\theta,t}, \tau_t)}\right)^{-1}.$$

Similarly,

$$\hat{\sigma}_{S,t} \equiv \mathbb{E}[\sigma_t \mid z_{S\sigma,t}] = \frac{V_\sigma}{V_\sigma + w_\sigma(e_{\sigma,t}, \tau_t)} z_{S\sigma,t},$$
$$V_\sigma^S(e_{\sigma,t}, \tau_t) \equiv \mathrm{Var}(\sigma_t - \hat{\sigma}_{S,t} \mid z_{S\sigma,t}) = \left(\frac{1}{V_\sigma} + \frac{1}{w_\sigma(e_{\sigma,t}, \tau_t)}\right)^{-1}.$$

Under quadratic loss, expected mean-squared forecast error equals expected posterior variance, so these variances are the relevant loss components.

5.3 Independent forecasting and import

Under the independent architecture,

$$\hat{y}_{S,t}^{\mathrm{ind}} \equiv \hat{\theta}_{S,t} + \hat{\sigma}_{S,t}.$$

Independence across the two components implies

$$\mathbb{E}\big[(y_t - \hat{y}_{S,t}^{\mathrm{ind}})^2\big] = V_\theta^S(e_{\theta,t}, \tau_t) + V_\sigma^S(e_{\sigma,t}, \tau_t). \tag{30}$$

Under the import architecture,

$$\hat{y}_{S,t}^{\mathrm{imp}} \equiv \hat{\theta}_{O,t} + \hat{\sigma}_{S,t},$$

so

$$\mathbb{E}\Big[(y_t - \hat{y}_{S,t}^{\mathrm{imp}})^2\Big] = V_O + V_\sigma^S(e_{\sigma,t}, \tau_t). \tag{31}$$

Because effort devoted to θ cannot improve the imported forecast, import entails $e_{\theta,t} = 0$ and $e_{\sigma,t} = \bar{E}$.

5.4 Optimal effort under independent forecasting

Under a best-forecast mandate, S solves

$$\begin{aligned} \min_{e_{\theta,t}, e_{\sigma,t} \geq 0} \quad & V_\theta^S(e_{\theta,t}, \tau_t) + V_\sigma^S(e_{\sigma,t}, \tau_t) \\ \text{s.t.} \quad & e_{\theta,t} + e_{\sigma,t} \leq \bar{E}. \end{aligned}$$

Because both posterior variances fall with effort, the capacity constraint binds. Suppressing time subscripts, define

$$L^{\text{ind}}(e_\theta;\tau) \equiv V_\theta^S(e_\theta,\tau) + V_\sigma^S(\bar{E}-e_\theta,\tau), \qquad e_\theta \in [0,\bar{E}]. \tag{32}$$

Using

$$V_j^S(e_j,\tau) = \left[\frac{1}{V_j} + \frac{1}{v_j(e_j,\tau)}\right]^{-1}, \qquad j \in \{\theta,\sigma\},$$

and differentiating with respect to the corresponding effort choice, (29) gives

$$\frac{\partial V_\theta^S(e_\theta,\tau)}{\partial e_\theta} = -\left(V_\theta^S(e_\theta,\tau)\right)^2 \frac{\alpha_\theta \tau}{\bar{v}_\theta}, \tag{33}$$

$$\frac{\partial V_\sigma^S(e_\sigma,\tau)}{\partial e_\sigma} = -\left(V_\sigma^S(e_\sigma,\tau)\right)^2 \frac{\alpha_\sigma \tau}{\bar{w}_\sigma}. \tag{34}$$

The marginal return to effort is therefore smaller when the processing window is compressed and larger when residual uncertainty is high.

Any interior optimum $e_\theta^*(\tau) \in (0,\bar{E})$ satisfies

$$\left(V_\theta^S(e_\theta^*,\tau)\right)^2 \frac{\alpha_\theta}{\bar{v}_\theta} = \left(V_\sigma^S(e_\sigma^*,\tau)\right)^2 \frac{\alpha_\sigma}{\bar{w}_\sigma}, \qquad e_\sigma^* = \bar{E} - e_\theta^*. \tag{35}$$

Thus the independent architecture equalises marginal reductions in posterior variance across the two components, adjusted for the effectiveness of effort in each forecasting technology. If this condition has no interior solution, the optimum lies at $e_\theta^* = 0$ or $e_\theta^* = \bar{E}$.

5.5 The import decision

Under import, expected loss is

$$L^{\text{imp}}(\tau) \equiv V_O + V_\sigma^S(\bar{E},\tau). \tag{36}$$

Under independent forecasting, optimised loss is

$$L^{\text{ind}*}(\tau) \equiv \min_{e_\theta \in [0,\bar{E}]} \left\{ V_\theta^S(e_\theta,\tau) + V_\sigma^S(\bar{E}-e_\theta,\tau) \right\}. \tag{37}$$

Import is optimal if and only if

$$L^{\text{imp}}(\tau) \leq L^{\text{ind}*}(\tau), \qquad \text{that is,} \qquad V_O + V_\sigma^S(\bar{E}, \tau) \leq L^{\text{ind}*}(\tau). \tag{38}$$

Since $V_\sigma^S(\bar{E}, \tau) \leq V_\sigma^S(\bar{E} - e_\theta, \tau)$ for every $e_\theta \geq 0$,

$$L^{\text{ind}*}(\tau) \geq \min_{e_\theta \in [0, \bar{E}]} V_\theta^S(e_\theta, \tau) + V_\sigma^S(\bar{E}, \tau).$$

A sufficient condition for import is therefore

$$V_O \leq \min_{e_\theta \in [0, \bar{E}]} V_\theta^S(e_\theta, \tau). \tag{39}$$

If even the best shared-component forecast that S can produce is no more precise than the central forecast, importing $\hat{\theta}_O$ is optimal under the best-forecast objective.

Under (29), the minimum is attained at $e_\theta = \bar{E}$:

$$\min_{e_\theta \in [0, \bar{E}]} V_\theta^S(e_\theta, \tau) = V_\theta^S(\bar{E}, \tau) \tag{40}$$

$$= \left[\frac{1}{V_\theta} + \frac{1}{\bar{v}_\theta} + \frac{\alpha_\theta \tau}{\bar{v}_\theta} \bar{E} \right]^{-1}. \tag{41}$$

Moreover,

$$\frac{\partial V_\theta^S(\bar{E}, \tau)}{\partial \tau} = - \left(V_\theta^S(\bar{E}, \tau) \right)^2 \frac{\alpha_\theta \bar{E}}{\bar{v}_\theta} < 0. \tag{42}$$

A shorter processing window therefore worsens the best attainable independent precision on θ and makes the import condition easier to satisfy. The resulting alignment is efficient conditional on the capacity constraint: S imports the central shared-component baseline and directs its scarce effort toward the devolved-specific component.

5.6 Selective concordats as block-by-block interface design

The shared component may contain several forecast blocks,

$$\theta \equiv (\theta^{(1)}, \ldots, \theta^{(K)}) \in \mathbb{R}^K,$$

and devolved receipts may load differently on those blocks:

$$y \equiv b'\theta + \sigma, \qquad b \in \mathbb{R}^K. \tag{43}$$

A block may therefore be relatively unimportant for the upstream institution's aggregate forecast while carrying a large loading for devolved receipts.

Assume the state blocks are mutually independent. Within each block k, both institutions start from the same finite prior and observe independent signal noises. Let

$$e_O^{(k)} \equiv \theta^{(k)} - \hat{\theta}_O^{(k)}, \qquad e_S^{(k)} \equiv \theta^{(k)} - \hat{\theta}_S^{(k)},$$

with

$$\mathrm{Var}(e_O^{(k)}) = V_O^{(k)}, \qquad \mathrm{Var}(e_S^{(k)}) = V_S^{(k)}.$$

The finite-prior covariance within block k is

$$C_k \equiv \mathrm{Cov}(e_O^{(k)}, e_S^{(k)}) = \frac{V_O^{(k)} V_S^{(k)}}{V_\theta^{(k)}} > 0. \tag{44}$$

Across distinct blocks, forecast errors are orthogonal:

$$\mathrm{Cov}(e_i^{(k)}, e_j^{(\ell)}) = 0 \qquad (k \neq \ell;\ i, j \in \{O, S\}).$$

Cross-block dependence could be admitted with covariance matrices; the present block-independent version isolates the effect of selecting which blocks are shared.

Let

$$A \subseteq \{1, \ldots, K\}$$

be the set of blocks treated as common inputs. Define diagonal selection matrices

$$(M_A)_{kk} \equiv \mathbf{1}\{k \in A\}, \qquad M_{A^c} \equiv I - M_A.$$

The published shared-component baseline is

$$\tilde{\theta} = M_A \hat{\theta}_O + M_{A^c} \hat{\theta}_S, \tag{45}$$

and the devolved forecast is

$$\hat{y}_S \equiv b'\tilde{\theta} + \hat{\sigma}_S.$$

If A contains every block, this is a full concordat; if A is empty, each institution retains its separately formed estimate throughout.

Proposition (accuracy and incremental common-mode exposure under a selective concordat).

1. *Accuracy.* The devolved mean-squared forecast error is

$$\mathbb{E}[(y - \hat{y}_S)^2] = V_\Sigma + \sum_{k \in A} b_k^2 V_O^{(k)} + \sum_{k \notin A} b_k^2 V_S^{(k)}. \tag{46}$$

   Relative to no concordat, adding block k to A weakly improves accuracy if and only if $V_O^{(k)} \leq V_S^{(k)}$.

2. *Common-mode exposure.* Let the upstream aggregate shared-state error be

$$E_O \equiv b'e_O,$$

   and the downstream shared-state error under selection set A be

$$E_S(A) \equiv b'(M_A e_O + M_{A^c} e_S).$$

   Their covariance is

$$\mathrm{Cov}(E_O, E_S(A)) = \sum_{k \in A} b_k^2 V_O^{(k)} + \sum_{k \notin A} b_k^2 C_k. \tag{47}$$

   Under no concordat, the covariance is already

$$\sum_{k=1}^{K} b_k^2 C_k > 0.$$

   Adding block k to the concordat raises common-mode exposure by

$$b_k^2 \big(V_O^{(k)} - C_k\big) > 0. \tag{48}$$

3. *Design implication.* A full concordat is not generally optimal unless the upstream institution is weakly more precise in every relevant block and the designer places no independent value on preserving additional information or limiting incremental common-mode exposure. With heterogeneous loadings and remit-driven precision, the coherent interface is selective.

Proof. The devolved forecast error is

$$y - \hat{y}_S = b'(M_A e_O + M_{A^c} e_S) + (\sigma - \hat{\sigma}_S).$$

Only one institution's error enters each block of the downstream forecast. Cross-block orthogonality and independence of the σ block therefore give (46).

For the covariance result,

$$\begin{aligned}\mathrm{Cov}(E_O, E_S(A)) &= \mathrm{Cov}\left(\sum_k b_k e_O^{(k)}, \sum_{k\in A} b_k e_O^{(k)} + \sum_{k\notin A} b_k e_S^{(k)}\right) \\ &= \sum_{k\in A} b_k^2 V_O^{(k)} + \sum_{k\notin A} b_k^2 C_k,\end{aligned}$$

where cross-block terms vanish. Replacing the separately formed downstream error in block k with the upstream error changes that block's covariance contribution from $b_k^2 C_k$ to $b_k^2 V_O^{(k)}$, proving (48). □

The design problem is therefore not whether forecasts should be wholly common or wholly separate. It is which inputs should be shared by design. The finite common prior means that even separately formed forecasts contain baseline comovement. A selective concordat adds to that comovement only in the chosen blocks. Public documentation should therefore identify both the blocks shared by design and those deliberately left to separate judgement, together with the data vintages and assumptions governing each.

6 Part VI. Repeated Interaction and Interface Stability

The baseline model is static. This Part asks when repeated interaction can sustain mutually beneficial sharing of agreed inputs even though sharing is not a one-shot equilibrium. The three propositions in this Part are local to the repeated-game extension and are distinct from Propositions 1–3 in Part II. Here reputation refers to compliance with a cooperative protocol, not to a reputation for producing any particular forecast. Its central implication is that stable close working depends on interface design. Cooperation is easier to sustain when compliance is clearly defined, observable and verifiable.

6.1 Stage game

Fiscal cycles are indexed by $t = 1, 2, \ldots$, and both institutions discount future payoffs by $\delta \in (0, 1)$. In each period,

$$\theta_t \sim \mathcal{N}(0, V_\theta),$$

independently across t. Institution $i \in \{O, S\}$ observes

$$z_{i,t} = \theta_t + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim \mathcal{N}(0, v_i),$$

with signal errors independent across institutions and periods and independent of θ_t.

At the start of the period, each institution chooses whether to share its private signal:

$$s_{i,t} \in \{0, 1\},$$

where $s_{i,t} = 1$ denotes sharing. Sharing costs the supplying institution $c > 0$. After any exchange, each institution publishes its posterior-mean forecast of θ_t. Under quadratic loss, expected squared forecast error equals posterior variance.

Let

$$V_i^0 \equiv \left(\frac{1}{V_\theta} + \frac{1}{v_i} \right)^{-1}$$

be institution i's posterior variance when it observes only its own signal, and let

$$V^B \equiv \left(\frac{1}{V_\theta} + \frac{1}{v_O} + \frac{1}{v_S} \right)^{-1}$$

be the posterior variance when both signals are available. The value to institution i of receiving the other institution's signal is

$$\Delta V_i \equiv V_i^0 - V^B > 0.$$

Its expected one-period payoff is

$$u_i(s_{i,t}, s_{-i,t}) \equiv -\operatorname{Var}(\theta_t - \hat{\theta}_{i,t}) - c s_{i,t}.$$

The four action profiles therefore yield

	$s_{-i,t} = 1$	$s_{-i,t} = 0$
$s_{i,t} = 1$	$-V^B - c$	$-V_i^0 - c$
$s_{i,t} = 0$	$-V^B$	$-V_i^0$.

Proposition 1 (one-shot equilibrium). In the one-shot sharing game, not sharing is a strictly dominant action for each institution, and $(s_O, s_S) = (0, 0)$ is the unique Nash equilibrium.

Proof. If the other institution shares, withholding saves c without changing the information available to i. If the other institution does not share, withholding again saves c without changing i's information. Hence not sharing is strictly dominant. □

6.2 Perfect monitoring

Suppose actions are perfectly observed at the end of each period. Consider grim-trigger strategies: both institutions share while no deviation has occurred; following any deviation, neither shares thereafter.

On the cooperative path, institution i obtains

$$V_i^C = \frac{-V^B - c}{1-\delta}.$$

A one-period deviation yields $-V^B$ immediately and the permanent non-sharing continuation value

$$V_i^N = \frac{-V_i^0}{1-\delta}.$$

The deviation value is therefore

$$V_i^D = -V^B + \delta V_i^N.$$

Proposition 2 (interface stability with perfect monitoring). Under grim-trigger strategies and perfect monitoring, mutual sharing can be sustained as a subgame-perfect equilibrium if

$$\delta \geq \frac{c}{\Delta V_i}, \qquad \text{for each } i \in \{O, S\}.$$

Proof. The no-deviation condition $V_i^C \geq V_i^D$ is

$$\frac{-V^B - c}{1-\delta} \geq -V^B + \delta \frac{-V_i^0}{1-\delta},$$

which is equivalent to

$$c \leq \delta(V_i^0 - V^B) = \delta \Delta V_i.$$

Rearranging gives the stated condition. □

The condition has the standard relational-contract interpretation. Cooperation is easier to sustain when the informational gain ΔV_i is large, the administrative cost c is small, or the institutions place greater weight on future interaction.

6.3 Imperfect monitoring

In practice, compliance may not be perfectly observable. Let a deviation be detected with probability $q \in [0,1]$. A higher q represents a more auditable interface: agreed calendars, request logs, data-vintage records and documented delivery protocols make failures to comply easier to verify.

Proposition 3 (interface stability with imperfect monitoring). Under grim-trigger strategies and detection probability q, mutual sharing can be sustained if

$$\delta \geq \frac{c}{q\Delta V_i + (1-q)c}, \qquad \text{for each } i \in \{O, S\}.$$

When $q = 1$, this reduces to the perfect-monitoring condition above. When $q = 0$, the condition becomes $\delta \geq 1$, so cooperation cannot be sustained.

Proof. A deviation yields $-V^B$ in the current period. From the next period, the deviation is detected with probability q, in which case play switches permanently to non-sharing, and is undetected with probability $1 - q$, in which case cooperation continues. Hence

$$V_i^D = -V^B + \delta \left[q\frac{-V_i^0}{1-\delta} + (1-q)\frac{-V^B - c}{1-\delta} \right].$$

Imposing $V_i^C \geq V_i^D$ gives

$$c \leq \delta\big[q\Delta V_i + (1-q)c\big],$$

which rearranges to the stated condition. □

These very simple results may be given an institutional interpretation. The repeated-game layer isolates three channels through which the political and administrative environment affects close working. First, additional clearance requirements, politicised scrutiny of routine contacts and ad hoc process changes raise the cost c of sharing. Second, clear public boundaries

raise q by making compliance verifiable. Third, institutional disruption can shorten the effective horizon. If the cooperative relationship survives to the next fiscal cycle with probability $\rho \in (0, 1]$, the effective discount factor is $\tilde{\delta} = \rho\delta$.

Political interference therefore need not take the form of directing an institution's published forecast. It can operate by raising c, lowering q, or reducing ρ, each of which makes defensive non-cooperation more likely. Conversely, standardised data pipelines and templates lower c; auditable calendars and protocols raise q; and stable remit boundaries increase the effective horizon.

Publication calendars and assumptions concordats enter this extension as elements of interface design. They can increase ΔV_i by ensuring that exchanged inputs are timely and decision-relevant, reduce c by routinising exchange, and raise q by making compliance documentable. The repeated-game result therefore reinforces the distinction developed in the main analysis: stable coordination concerns a monitored process for sharing inputs under clear public boundaries. It does not require convergence of published outputs, and those boundaries help routine cooperation coexist with demonstrable analytical independence.

7 Part VII. Bayesian Learning and the Oversight Environment

This Part adds a fixed, unobserved competence type for each institution and asks what outsiders can learn from realised forecast performance. The extension follows the standard Bayesian-learning logic in which repeated outcomes update beliefs about a persistent hidden type (Baley and Veldkamp, 2021). It is not needed for the main results. Its purpose is as a preliminary attempt to formalise a governance channel. An oversight environment can improve learning by making information constraints

observable, or impede learning by treating disagreement itself as evidence of unreliability.

7.1 Competence types and published forecasts

Fiscal cycles are again indexed by $t = 1, 2, \ldots$, with

$$\theta_t \sim \mathcal{N}(0, V_\theta)$$

independently across t. Institution $i \in \{O, S\}$ has a fixed type

$$\kappa_i \in \{H, L\}, \qquad p_{i,1} \equiv \Pr(\kappa_i = H) \in (0, 1),$$

where high competence means a more precise private signal:

$$0 < v_i(H) < v_i(L).$$

In period t,

$$z_{i,t} = \theta_t + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} \mid \kappa_i \sim \mathcal{N}(0, v_i(\kappa_i)).$$

Signal noises are independent across institutions and periods conditional on types and are independent of the state.

The institution knows its type and publishes the posterior mean

$$a_{i,t} = q_i(\kappa_i) z_{i,t}, \qquad q_i(\kappa) \equiv \frac{V_\theta}{V_\theta + v_i(\kappa)}.$$

Define the realised forecast error

$$e_{i,t} \equiv \theta_t - a_{i,t}.$$

The evaluator observes the published forecast and the realised state and therefore observes both $(a_{i,t}, \theta_t)$, equivalently $(e_{i,t}, \theta_t)$.

7.2 Likelihood conditional on the realised state

Conditional on the realised state and a candidate competence type,

$$a_{i,t} \mid \theta_t, \kappa \sim \mathcal{N}\big(q_i(\kappa)\theta_t,\ q_i(\kappa)^2 v_i(\kappa)\big)\,.$$

Equivalently,

$$e_{i,t} \mid \theta_t, \kappa \sim \mathcal{N}\big(\mu_i(\theta_t, \kappa),\, s_i^2(\kappa)\big)\,, \tag{49}$$

where

$$\mu_i(\theta, \kappa) \equiv \big(1 - q_i(\kappa)\big)\theta, \qquad s_i^2(\kappa) \equiv q_i(\kappa)^2 v_i(\kappa).$$

The conditional mean is generally non-zero because a posterior-mean forecast shrinks the private signal toward the common prior. Consequently, a rational evaluator who observes the realised state should not use a zero-mean marginal error density as the full likelihood.

Let

$$\mathcal{L}_{i,t}(\kappa) \equiv \varphi\big(e_{i,t}; \mu_i(\theta_t, \kappa), s_i^2(\kappa)\big)\,,$$

where $\varphi(\cdot; \mu, V)$ is the normal density. Bayes' rule gives

$$p_{i,t+1} = \frac{p_{i,t}\mathcal{L}_{i,t}(H)}{p_{i,t}\mathcal{L}_{i,t}(H) + (1 - p_{i,t})\mathcal{L}_{i,t}(L)}. \tag{50}$$

In log-odds form,

$$\ell_{i,t+1} = \ell_{i,t} + \log \frac{\mathcal{L}_{i,t}(H)}{\mathcal{L}_{i,t}(L)}, \qquad \ell_{i,t} \equiv \log \frac{p_{i,t}}{1 - p_{i,t}}.$$

The one-period log likelihood ratio is

$$\log \frac{\mathcal{L}_{i,t}(H)}{\mathcal{L}_{i,t}(L)} = \frac{1}{2} \log \frac{s_i^2(L)}{s_i^2(H)} - \frac{\big(e_{i,t} - \mu_i(\theta_t, H)\big)^2}{2s_i^2(H)} + \frac{\big(e_{i,t} - \mu_i(\theta_t, L)\big)^2}{2s_i^2(L)}. \tag{51}$$

Unlike a simple variance-comparison rule, the information in a realised error depends on both its size and the realised state. There is therefore no general state-independent threshold of the form $|e| < e^*$.

Proposition (competence learning with full public information). Suppose $v_i(H) \neq v_i(L)$. Then the evaluator updates according to (50). Posterior belief in high competence rises in period t if and only if the right-hand side of (51) is positive. Across repeated independent fiscal cycles, competence is identifiable whenever the two conditional forecast distributions differ on a set of realised states of positive probability.

Sketch of Proof. The update follows from Bayes' rule and (49). Distinct signal variances imply distinct conditional distributions for the published forecast except in degenerate parameter cases. Standard likelihood-ratio consistency then gives identification from repeated independent observations. □

7.3 Error-only evaluation as a restricted special case

If the evaluator deliberately discards the realised state level and conditions only on the scalar forecast error, integrating over θ_t gives

$$e_{i,t} \mid \kappa \sim \mathcal{N}(0, V_i(\kappa)), \qquad V_i(\kappa) = \left(\frac{1}{V_\theta} + \frac{1}{v_i(\kappa)} \right)^{-1}.$$

Under this restricted information set, the marginal error likelihood is valid:

$$\widetilde{\mathcal{L}}_{i,t}(\kappa) = \varphi(e_{i,t}; 0, V_i(\kappa)).$$

The posterior belief in high competence then rises if and only if

$$|e_{i,t}| < e_i^*,$$

where

$$(e_i^*)^2 = \frac{\log(V_i(L)/V_i(H))}{1/V_i(H) - 1/V_i(L)}.$$

This symmetric threshold is therefore a result for *error-only evaluation*, not for an evaluator who uses the full publicly observed pair $(a_{i,t}, \theta_t)$.

This has implications for cross-institution dependence. Unconditionally, the two institutions' posterior forecast errors in a given period share finite-prior covariance. Conditional on the realised state and competence types, however, their signal noises are independent, so the joint likelihood of the two published forecasts factors across institutions. An evaluator learning jointly about both types should therefore condition on the realised state rather than multiply unconditional marginal error densities.

7.4 The oversight environment

To represent the consequences of external evaluation, suppose the institution's reduced-form period payoff includes

$$\gamma\,\mathbb{E}[p_{i,t+1} \mid \mathcal{I}_{i,t}], \qquad \gamma \geq 0.$$

The parameter γ captures the fact that perceived competence can affect credibility, scrutiny, access to information and the stability of operating protocols. This does not require the institution to value reputation intrinsically.

Two governance regimes should be distinguished. Under *supportive governance*, the evaluator observes the relevant conditioning set and constraints, including data vintages, costing readiness and declared assumptions. Better attribution helps distinguish underlying capability from an unusually difficult information environment. Under *interfering governance*, divergence from a salient baseline is itself treated as evidence of unreliability. That creates pressure for output convergence even when separately produced forecasts contain useful information.

7.5 Identification, limitations and design implications

This likelihood reinforces three design implications. First, forecast reviews should record the state and information set against which each forecast was made; the same numerical error can carry different evidence about competence in different states. Second, performance should be evaluated over sufficiently long windows because any one error is noisy evidence. Third, enforced alignment pools published outcomes and reduces the independent variation from which capability and model differences might otherwise be inferred.

Calendar rules and assumptions concordats can support learning when they make information sets and responsibilities observable. They undermine

it when they are used to erase analytically meaningful differences in published outputs. The competence-learning layer therefore complements the baseline protocol analysis, but it does not alter the paper's main results.

8 Part VIII. Political Narrative Pressure

This Part introduces a political narrative anchor into the baseline sequential-publication problem. The symmetric specification provides a linear benchmark in which both institutions are pulled toward a common anchor. The one-sided specification captures the empirically more relevant case in which pressure is activated only when an institution's forecast lies below the preferred narrative. Throughout, the publication order is $O \to S$. Let

$$\hat{\theta}_{O,t}^{acc} \equiv \mathbb{E}[\theta_t \mid z_{O,t}], \qquad \hat{\theta}_{B,t} \equiv \mathbb{E}[\theta_t \mid z_{S\theta,t}, z_{O,t}],$$

denote, respectively, the upstream posterior mean based on its own signal and the two-signal posterior mean available to the downstream institution. In the Gaussian environment, observing the upstream publication is informationally equivalent to observing $z_{O,t}$.

A symmetric quadratic narrative wedge

Let g_t denote the narrative anchor for the shared component. The parameters $\pi_O \geq 0$ and $\pi_S \geq 0$ measure narrative pressure on the upstream and downstream institutions, while $\eta \geq 0$ measures the downstream cost of diverging from the upstream publication.

Upstream publication. Suppose O chooses $\hat{\theta}_{O,t}$ to minimise

$$\mathcal{J}_{O,t}(\hat{\theta}_{O,t}) \equiv \mathbb{E}\Big[(\theta_t - \hat{\theta}_{O,t})^2 \mid z_{O,t}\Big] + \pi_O(\hat{\theta}_{O,t} - g_t)^2. \tag{52}$$

Using the standard quadratic-loss decomposition, this is equivalent, up to an additive constant, to minimising

$$(a - \hat{\theta}_{O,t}^{acc})^2 + \pi_O(a - g_t)^2.$$

Lemma A.1 (upstream symmetric narrative pressure). The unique minimiser of (52) is

$$\hat{\theta}^*_{O,t} = \frac{\hat{\theta}^{acc}_{O,t} + \pi_O g_t}{1 + \pi_O}. \tag{53}$$

Proof. The first-order condition is $2(a - \hat{\theta}^{acc}_{O,t}) + 2\pi_O(a - g_t) = 0$, which gives (53). Strict convexity gives uniqueness. □

Downstream publication. After observing $\hat{\theta}^*_{O,t}$, suppose S chooses $\tilde{\theta}_t$ to minimise

$$\mathcal{J}_{S,t}(\tilde{\theta}_t) \equiv \mathbb{E}\Big[(\theta_t - \tilde{\theta}_t)^2 \mid z_{S\theta,t}, z_{O,t}\Big] + \eta(\tilde{\theta}_t - \hat{\theta}^*_{O,t})^2 + \pi_S(\tilde{\theta}_t - g_t)^2. \tag{54}$$

Up to an additive constant, the objective is

$$(\tilde{\theta}_t - \hat{\theta}_{B,t})^2 + \eta(\tilde{\theta}_t - \hat{\theta}^*_{O,t})^2 + \pi_S(\tilde{\theta}_t - g_t)^2.$$

Proposition A.1 (downstream symmetric narrative pressure). The unique minimiser of (54) is

$$\tilde{\theta}^*_t = \frac{\hat{\theta}_{B,t} + \eta\,\hat{\theta}^*_{O,t} + \pi_S g_t}{1 + \eta + \pi_S}. \tag{55}$$

Substituting (53) gives

$$\tilde{\theta}^*_t = \frac{\hat{\theta}_{B,t} + \dfrac{\eta}{1+\pi_O}\hat{\theta}^{acc}_{O,t} + \left(\pi_S + \dfrac{\eta\pi_O}{1+\pi_O}\right) g_t}{1 + \eta + \pi_S}. \tag{56}$$

Proof. The first-order condition is

$$2(\tilde{\theta}_t - \hat{\theta}_{B,t}) + 2\eta(\tilde{\theta}_t - \hat{\theta}^*_{O,t}) + 2\pi_S(\tilde{\theta}_t - g_t) = 0.$$

Solving gives (55); substitution of (53) gives (56). Strict convexity gives uniqueness. □

The benchmark nests the earlier cases. If $\pi_O = \pi_S = 0$, the downstream publication is the weighted average generated by the calendar-and-divergence

problem. If $\eta = 0$, only the downstream narrative wedge remains. Under a strict concordat fixing $\tilde{\theta}_t \equiv \hat{\theta}_{O,t}$, the downstream choice problem disappears and narrative pressure operates through the upstream publication alone. The symmetric specification therefore separates three anchors: the two-signal estimate, the upstream publication and the political narrative.

One-sided narrative pressure

Political pressure is often asymmetric: an institution is criticised for publishing a forecast below the preferred narrative, but not for exceeding it. The one-sided penalty $(g_t - a)_+^2$, where $(u)_+ \equiv \max\{u, 0\}$, captures this feature.

Upstream publication. Let $\mu_{O,t} \equiv \hat{\theta}_{O,t}^{acc}$. Suppose O minimises

$$(\hat{\theta}_{O,t} - \mu_{O,t})^2 + \pi_O \, (g_t - \hat{\theta}_{O,t})_+^2, \qquad \pi_O \geq 0. \tag{57}$$

Lemma A.2 (upstream one-sided wedge). The unique minimiser of (57) is

$$\hat{\theta}_{O,t}^* = \begin{cases} \mu_{O,t}, & \mu_{O,t} \geq g_t, \\ \dfrac{\mu_{O,t} + \pi_O g_t}{1 + \pi_O}, & \mu_{O,t} < g_t. \end{cases} \tag{58}$$

Proof. If $\hat{\theta}_{O,t} \geq g_t$, the penalty is zero and the minimiser is $\mu_{O,t}$, which is feasible if and only if $\mu_{O,t} \geq g_t$. If $\mu_{O,t} < g_t$, the minimiser lies below g_t, where the objective is quadratic and the first-order condition gives the second branch of (58). The resulting value is indeed below g_t, and strict convexity gives uniqueness. □

Downstream publication. Given $\hat{\theta}_{O,t}^*$, suppose S minimises

$$(\tilde{\theta}_t - \hat{\theta}_{B,t})^2 + \eta(\tilde{\theta}_t - \hat{\theta}_{O,t}^*)^2 + \pi_S \, (g_t - \tilde{\theta}_t)_+^2, \qquad \eta, \pi_S \geq 0. \tag{59}$$

Define the downstream publication in the absence of narrative pressure by

$$\tilde{\theta}_t^0 \equiv \frac{\hat{\theta}_{B,t} + \eta \hat{\theta}^*_{O,t}}{1+\eta}. \tag{60}$$

Proposition A.2 (downstream one-sided wedge). The unique minimiser of (59) is

$$\tilde{\theta}^*_t = \begin{cases} \tilde{\theta}_t^0, & \tilde{\theta}_t^0 \geq g_t, \\ \dfrac{\hat{\theta}_{B,t} + \eta\, \hat{\theta}^*_{O,t} + \pi_S g_t}{1+\eta+\pi_S}, & \tilde{\theta}_t^0 < g_t. \end{cases} \tag{61}$$

Proof. If $\tilde{\theta}_t \geq g_t$, the penalty is zero and the minimiser is $\tilde{\theta}_t^0$, feasible if and only if $\tilde{\theta}_t^0 \geq g_t$. Otherwise the minimiser lies below g_t, where the objective is quadratic and the first-order condition gives the second branch of (61). Again, strict convexity gives uniqueness. □

Transmission through the interface. Consider the asymmetric case $\pi_O > 0$ and $\pi_S = 0$. Whenever $\eta > 0$, an active upstream wedge is transmitted into the downstream publication through the cost of divergence. From (58)–(60),

$$\begin{aligned} \frac{\partial \hat{\theta}^*_{O,t}}{\partial g_t} &= \mathbf{1}\{\mu_{O,t} < g_t\} \frac{\pi_O}{1+\pi_O}, \\ \frac{\partial \tilde{\theta}^*_t}{\partial g_t} &= \frac{\eta}{1+\eta} \frac{\partial \hat{\theta}^*_{O,t}}{\partial g_t} \\ &= \mathbf{1}\{\mu_{O,t} < g_t\} \frac{\eta}{1+\eta} \frac{\pi_O}{1+\pi_O}. \end{aligned}$$

Thus political optimism can propagate downstream even without a formal concordat and even when the downstream institution faces no direct narrative penalty.

Covariance and fiscal-risk implications

The symmetric model also yields a useful diagnostic implication. With a finite common prior, separately formed posterior errors already comove even

when signal noises are independent. In this extension, both publications are additionally weighted averages of statistical estimates and the common anchor g_t, so variation in that anchor can add a further common component. Schematically,

$$\hat{\theta}^*_{O,t} = \mu_{O,t} + a_O(g_t - \mu_{O,t}), \qquad \tilde{\theta}^*_t = \hat{\theta}_{B,t} + a_S(g_t - \hat{\theta}_{B,t}) + b_S(\hat{\theta}^*_{O,t} - \hat{\theta}_{B,t}),$$

for non-negative weights a_O, a_S and b_S. Variation in the common anchor therefore adds to common-mode forecast error above the Bayesian covariance already implied by the finite prior. Correlated forecast failures are not, by themselves, evidence of either a concordat or informational dependence.

The fiscal consequences depend on the reconciliation mapping. In Scotland-style arithmetic, the shared component enters the net reconciliation with the opposite sign. Greater comovement can therefore compress short-run reconciliation volatility through cancellation, even when the common forecast is biased upward. That short-run cash smoothing does not eliminate the underlying error: it may instead reappear through later baseline revisions and reconciliations. In transfer systems where the shared component enters with the same sign, the same common wedge raises common-mode risk and can increase reconciliation volatility.

The governance implication is that transparency and block-level attribution complement formal independence. Auditable conditioning sets help distinguish political pressure from genuine information differences and reveal whether the interface is transmitting a common narrative through the cost of divergence.

References

Baley, I. and Veldkamp, L. (2021). Bayesian learning. CEPR Discussion Paper DP16377, Centre for Economic Policy Research (CEPR). Also available via SSRN. Accessed 2025-12-24.