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 13, 2026

Abstract

Independent fiscal institutions often produce forecasts that jointly determine public budgets, with errors later triggering cash adjustments. Fiscal risk therefore depends not only on each forecast's accuracy, but also on whether their errors move together. Publication timing, common assumptions, information sharing and deadlines can affect both forecast quality and later budget corrections. Publishing second may encourage one institution to follow the other too closely, while tight deadlines may make reliance on another forecast rational even when independent analysis remains valuable. In Scotland, errors in forecasts of devolved revenues and the Block Grant Adjustment affect the budget in opposite directions. Similar assumptions can therefore reduce short-run cash volatility while also reducing the independent information contained in the two forecasts. Independence and coordination must therefore be designed jointly. Transparent, selective coordination can limit avoidable fiscal risk while preserving useful differences in judgement.

JEL Classification: D83; H68; H77

Keywords: independent fiscal institutions; fiscal federalism; forecast coordination; reconciliation risk; information sharing

*The usual disclaimer applies. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising from this submission.

†Email: campbell.leith@glasgow.ac.uk

‡Email: charles.nolan@glasgow.ac.uk

§Email: graeme.roy@glasgow.ac.uk

1 Introduction

Independent fiscal institutions (IFIs) — fiscal councils and parliamentary budget offices in various national guises — have become a standard element of modern fiscal frameworks. Their spread is especially striking in multi-level settings: federations, decentralised states, and devolved systems in which fiscal decisions are made at more than one tier of government. International organisations have increasingly treated IFIs not as optional "add-ons", but as part of the core institutional infrastructure that supports credible fiscal policy. The OECD Principles for Independent Fiscal Institutions, and the subsequent programme of OECD IFI reviews, emphasise operational independence, access to information, transparency, and clear accountability (OECD, 2014, 2019, 2025). Within the EU, reforms to the fiscal governance framework have further strengthened the role of IFIs, including provisions on their tasks, independence, and "comply-or-explain" treatment of their assessments (European Commission, 2025a,b). The IMF likewise emphasises fiscal councils as a pillar of rules-based frameworks and, more broadly, as a tool for improving the quality and credibility of fiscal policy (Alonso-Albarran et al., 2025).

Yet this institutional trend immediately raises a design problem that becomes sharper in devolved and federal environments. IFIs are typically 'independent' with a mandate to forecast and scrutinise 'as well as possible.' At the same time, the fiscal system typically requires that multiple institutions (and multiple executive departments) interact under hard calendar constraints. Forecasts must be produced to a timetable, with inputs sometimes delivered late or revised, with policy costings prepared under time pressure, and with a need to separate what is "shared by design" (common baselines, common data, common assumptions) from what is genuinely institution-specific. In practice, then, independence is sustained not by isolation but by managed interdependence; that is, a set

of protocols that determine what information is shared, when it is shared, how assumptions are aligned (or allowed to differ), and how disagreements are communicated. These protocols are not merely procedural. They may significantly shape the joint distribution of forecast errors across institutions and therefore affect both accountability and budget outcomes.

The paper draws on four important ideas. Forecast-combination theory shows why joint accuracy depends on the covariance of forecast errors, not only their individual variances (Bates and Granger, 1969; Wang et al., 2023). Reputational concerns can induce agents to follow visible predecessors too closely (Scharfstein and Stein, 1990), while public information can improve common understanding yet reduce the value of independent signals when coordination motives are strong (Morris and Shin, 2002). Fiscal councils, meanwhile, influence policy through information, scrutiny and accountability rather than direct control of fiscal instruments (Beetsma and Debrun, 2016). Our contribution is to combine these insights in a two-tier fiscal setting where joint forecast errors trigger later cash reconciliations.

The salience of this "interface" problem is easy to see in decentralised fiscal systems where transfers, settlements, and reconciliations ('true-ups') are material.[1]

A simple way to frame this interface–and the fiscal devolution and federalism that produces it–is to start from the fact that sub-national budgets are typically set on the basis of forecasts. In its most basic form, one can think of the local budget constraint as being balanced *ex ante* using forecast numbers:

$$\begin{aligned}(\text{forecast local taxes}) - (\text{forecast local spending}) \\ + (\text{forecast net transfers from the centre}) \quad = 0.\end{aligned}$$

This is not a modelling assumption so much as an operational description; budgets are set in real time using predicted revenues, predicted spending

[1]Appendix A describes several such real-world systems and provides the broad institutional motivation for the reduced-form set-up.

pressures, and predicted transfers/settlements.

When outturns arrive, the fiscal framework may then generate a *reconciliation* (a cash adjustment, settlement, or accounting correction) reflecting the difference between what was forecast and what actually happened. At a high level, the reconciliation is driven by two ingredients. First, there is the forecast error in the local net position (local taxes relative to local spending). Second, in some systems there is an additional adjustment tied to forecast errors in net transfers from the centre. In words,

$$\begin{aligned}(\text{net reconciliation}) \;&= \text{forecast error in (local taxes} - \text{local spending)} \\ &\quad + \text{(a transfer-related forecast-error term)}.\end{aligned}$$

The transfer-related component is not universal. Some systems contain such explicit true-ups. Others largely treat transfers as fixed within-year and let any drift feed into later baselines. This distinction matters because it governs how strongly transfer forecast errors translate into contemporaneous cash adjustments, and therefore how important transfers are for the volatility of the realised local/devolved budget.

A special case: Scotland within the UK. Scotland fits naturally into the same logic, but with an important twist that we treat separately in the paper. The Scottish budget is set using a forecast-balanced constraint that includes devolved revenues and a block-grant-based funding settlement.[2] In words:

$$\begin{aligned}&(\text{forecast Scottish taxes}) - (\text{forecast Scottish spending}) \\ &\qquad + (\text{forecast block grant}) \\ &\qquad - (\text{forecast block grant adjustment}) \;=\; 0.\end{aligned}$$

The block grant is the transfer from the UK Government to Scotland, with changes continuing to be determined through the Barnett formula

[2]Scotland has some limited borrowing capacity, but we set this aside as it is of limited significance at present.

(Scottish Government and United Kingdom Government, 2023). The Block Grant Adjustments (BGAs) capture the hypothetical net revenues that the rUK loses as a result of devolving specific fiscal powers to Scotland. When outturns arrive, the mechanical reconciliations that matter for Scottish cash management are primarily tied to (i) forecast errors in devolved/assigned revenues and related local-budget components, and (ii) forecast errors in the BGA. By contrast, there is not, in the same mechanical end-year sense, an adjustment for within-year over- or under-spending in the UK budget that the Barnett formula was based on (even though such outturn drift can affect baselines for subsequent budgets). Put differently, the block grant is central for funding, but the systematic reconciliation mechanism in the Scottish framework operates mainly through devolved receipts and the BGA. This makes the reconciliation risks faced by the Scottish budget somewhat distinct from those typical in other devolved systems.

The Scotland–UK framework is the paper's main application, but the underlying problem is broader. Appendix A sketches several multi-level fiscal systems and shows how different combinations of own-source revenues, transfers, common forecasts and later settlements motivate the reduced-form reconciliation set-up used below.

1.1 Contribution and results

We study the interface between two independent fiscal institutions operating at different tiers of government. Their forecasts are used to set budgets and, because forecast errors later generate cash reconciliations, the economically relevant object is the joint distribution of errors rather than the accuracy of either institution considered in isolation. Coordination is therefore unavoidable, but its form matters: the same arrangements that reduce duplication or cash volatility may also weaken the information conveyed by independent forecasts.

We present coordination as a protocol fixed in advance. The protocol has four elements: a publication calendar; an assumptions concordat specifying which inputs are common and which may differ; a pre-publication sharing rule governing additional data, nowcasts, and technical assumptions; and a transparency regime determining what outsiders can verify. The protocol shapes publication-time information sets, feasible precision, forecast covariance, and the ability of outside observers to distinguish legitimate agreement from pressure-induced convergence.

Our first result shows that sequential publication can induce anchoring. When the second mover is penalised for visible disagreement, it is pulled toward the first mover's baseline even when doing so reduces the independent information contained in its forecast. Our second result shows that assumptions concordats affect more than comparability–by governing the shared component of forecasts, they change error covariance and hence reconciliation volatility. Our third result identifies when late information erodes the value of sequencing. If upstream precision deteriorates sufficiently relative to the downstream institution's own information, the second mover gains little usable information while the system can still inherit additional common exposure and herding risk. These results turn calendar design, assumptions governance, and data-readiness disclosure into first-order economic choices.

The Scotland–UK application sharpens the argument. The Block Grant Adjustment enters the Scottish reconciliation with the opposite sign to devolved revenues. Common treatment of the shared component can therefore produce mechanical cancellation in cash terms. But cancellation is not the same as informational efficiency since alignment may stabilise the budget while reducing the independent information supplied by the SFC. The application thus reveals clearly why coordination and independence cannot be assessed separately.

The later analysis develops three further implications. First, once

reconciliation volatility is itself costly, the forecast that minimises expected squared error need not minimise fiscal risk; a budget-setting authority may prefer a forecast that leans against the expected reconciliation. Second, time compression and limited analytical capacity can make importing the upstream forecast privately optimal, even when independent production would be more informative absent the deadline. Third, selective concordats and political-pressure extensions show that neither divergence nor convergence is easily interpreted: divergence may reflect efficient differences in remit and precision, while convergence may reflect common pressure rather than shared information.

The paper proceeds as follows. Section 2 sets out the model and the coordination protocol. Section 3 derives the results on calendar sequencing, assumptions concordats, and late information. Section 4 maps the framework into the SFC–OBR setting and derives the Scottish corollary. Section 5 distinguishes forecast accuracy from reconciliation stability, establishes the forecast-as-hedging result, and explains the value and limits of an independent devolved IFI. Section 6 studies late costings, capacity allocation, and endogenous importing of the upstream baseline. Section 7 draws together the governance implications and concludes. The Online Appendix contains the proofs and the further extensions on precision-status sharing, selective concordats, repeated interaction, Bayesian learning, and political narrative pressure.[3]

[3]Specifically, the Online Appendix, available on request or at https://charles-nolan.github.io, is structured as follows. Part I describes some basic technical results in Bayesian analysis used extensively. Part II contains the proofs of Propositions 1–3 and the extension in which the sharing rule (defined below by ρ) transmits precision-status information beyond the published point forecast. Part III derives the Scotland–UK reconciliation mapping and the Scottish corollary. Part IV develops the forecast-as-hedging and covariance results. Part V contains the calendar–capacity model and the block-by-block selective-concordat extension. Parts VI and VII develop repeated interaction and Bayesian learning under external evaluation. Part VIII analyses symmetric and one-sided political narrative pressure and shows how common pressure wedges can add to the baseline covariance of forecast errors even without a concordat.

2 Model and coordination protocol

2.1 Environment and timing

The model is very stylised. A fiscal cycle has three dates, $t \in \{0, 1, 2\}$, two independent fiscal institutions (IFIs), $i \in \{O, S\}$, and a devolved budget-setting authority, G. O is the 'upstream' or central IFI and S is the 'downstream' or devolved institution. The uncertain state is (θ, σ), with a commonly known prior. We interpret θ as a shared or central fiscal component and σ as a devolved-specific component. The devolved fiscal-relevant outcome is

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

Thus, the upstream institution forecasts the shared component, while the downstream institution's target also depends on information specific to the devolved tier.[4]

At $t = 0$, a cycle-specific coordination protocol Π is fixed by rule, memorandum of understanding, or agreed operating procedure. Each IFI chooses information-production effort $e_i \geq 0$ and, where the protocol permits, exchanges specified inputs. At $t = 1$, the institutions publish forecasts—simultaneously or sequentially—and G sets its budget using those forecasts. At $t = 2$, the state is realised, forecast errors become observable, reconciliations occur, and an evaluator P updates beliefs about the institutions.

2.2 The coordination protocol

Given Π, institution i observes a private signal

$$z_i = m_i(\theta, \sigma) + \varepsilon_i, \qquad \mathbb{E}[\varepsilon_i] = 0,$$

[4]For most of the analysis θ is scalar. The block generalisation used for selective concordats is developed in Online Appendix Part V.

where $m_i(\cdot)$ is known. A convenient benchmark is $m_O(\theta, \sigma) = \theta$ and $m_S(\theta, \sigma) = y$. Information-production effort increases usable precision:

$$\mathrm{Var}(\varepsilon_i \mid e_i) = \nu_i(e_i), \qquad \nu_i'(e_i) < 0, \qquad c_i'(e_i) > 0, \quad c_i''(e_i) > 0.$$

Later sections allow timing pressure and late costings to reduce the precision attainable from a given amount of effort.

We represent "close working" by a protocol fixed in advance,

$$\Pi \equiv (\kappa, A, \rho, \mathcal{T}).$$

Its four elements govern different margins of the interface. The calendar κ fixes publication order and, under sequencing, the lag:

$$\kappa \in \{\text{simultaneous},\ O \to S,\ S \to O\}, \qquad \Delta \geq 0.$$

The assumptions concordat A specifies which inputs belong to a common conditioning set and which may differ. The sharing rule ρ specifies what may be exchanged privately before publication beyond what sequencing alone reveals.[5] Transparency $\mathcal{T}$ determines what outsiders can observe and verify about the protocol, including rules of the road, data vintages, costing status and compliance.

The protocol matters through three induced objects: the institutions' publication-time information sets $\mathcal{I}_{i,t}(\Pi)$; the precision feasible within the cycle; and the public history $\mathcal{H}_t(\Pi)$ that outsiders can verify ex post. This distinction is useful because two protocols with different formal provisions are economically equivalent if they induce the same information, feasible precision and verifiable history.

[5] In the baseline Gaussian model with known variances, the upstream point forecast already summarises its signal about θ. Online Appendix Part II gives ρ operational content by allowing the institutions to share verifiable precision-status information—including data vintages, costing readiness and nowcast diagnostics—that is not contained in the point forecast. Without this information, the downstream institution must use a fixed reliance rule and can place too much weight on the upstream baseline precisely when that baseline is least reliable.

2.3 Forecasts, reconciliations, and institutional objectives

At publication, O produces $\hat{x}_O$ as a forecast of θ, while S produces $\hat{x}_S$ as a forecast of y. Define

$$\tilde{\theta} \equiv \theta - \hat{x}_O, \qquad \tilde{y} \equiv y - \hat{x}_S.$$

Forecast errors matter not only for measured accuracy but also because they can generate subsequent cash adjustments. We write the devolved reconciliation as

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

where λ captures how the shared component enters the devolved net position.[6] The devolved authority incurs a convex loss from reconciliation volatility, $L_G \equiv \mathbb{E}[v(R)]$, with $v(R) = \omega R^2$ as a convenient benchmark.

Each IFI values forecast accuracy and incurs effort costs. The protocol can also affect reputation and perceived independence because it changes what outsiders can infer from published outputs and realised errors. Let

$$\ell_O \equiv \mathbb{E}[\tilde{\theta}^2], \qquad \ell_S \equiv \mathbb{E}[\tilde{y}^2],$$

and let $\mathcal{R}_i$ denote a reputational payoff increasing in evaluator P's posterior assessment of competence. Let $\mathcal{D}_i$ denote a perceived-process or dependence index,

$$\mathcal{D}_i = d_i(\rho, A, \mathcal{T}; \hat{x}_O, \hat{x}_S), \qquad d_{i,\rho} \geq 0, \quad d_{i,\mathcal{T}} \leq 0 \quad \text{(natural benchmark)}.$$

IFI payoffs are

$$U_O \equiv -\ell_O - c_O(e_O) + \beta_O\mathbb{E}[\mathcal{R}_O] - \chi_O\mathbb{E}[\mathcal{D}_O],$$
$$U_S \equiv -\ell_S - c_S(e_S) + \beta_S\mathbb{E}[\mathcal{R}_S] - \chi_S\mathbb{E}[\mathcal{D}_S].$$

[6] In Scotland, the Barnett-related block grant is not reconciled within-year, corresponding to a $\lambda = 0$ component, while the Block Grant Adjustment is a separately reconciled object entering with the opposite sign. Section 4 develops this mapping explicitly.

with $\beta_i, \chi_i \geq 0$.

Equilibrium is Perfect Bayesian Equilibrium. Evaluator P observes the published forecasts and their timing, the errors realised at $t = 2$, and the elements of Π disclosed under $\mathcal{T}$. Beliefs are updated by Bayes' rule wherever possible, and each IFI chooses effort and forecasting behaviour optimally given the protocol and anticipated evaluation. The propositions below isolate particular margins of this general environment by holding the remaining margins fixed.

3 Three propositions on protocol design

Proposition 1 (Calendar: learning versus herding under demonstrable-independence concerns)

Environment. Let $\theta \sim \mathcal{N}(0, V_\theta)$ and $\sigma \sim \mathcal{N}(0, V_\sigma)$ be independent, with $y \equiv \theta + \sigma$. Institution O observes

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

while institution S observes

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

with all noise terms mutually independent and independent of (θ, σ).

Under quadratic loss, the accuracy-minimising forecast is the relevant conditional mean. Define

$$\hat{\theta}_O \equiv \mathbb{E}[\theta \mid z_O], \qquad \hat{\theta}_S \equiv \mathbb{E}[\theta \mid z_{S\theta}], \qquad \hat{\sigma}_S \equiv \mathbb{E}[\sigma \mid z_{S\sigma}],$$

and let $\hat{\theta}_B \equiv \mathbb{E}[\theta \mid z_{S\theta}, z_O]$ denote the posterior mean using both signals about θ. Standard normal–normal updating gives

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

Similarly,

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

and

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

The underlying Gaussian derivations are collected in Online Appendix Part I.

Statement. Suppose the calendar is sequential, $\kappa = O \to S$, so that S observes $\hat{\theta}_O$ before choosing the θ-component of its own published baseline, denoted $\tilde{\theta}$. Institution S minimises

$$\mathcal{J}(\tilde{\theta}) \equiv \mathbb{E}\Big[(\theta - \tilde{\theta})^2 \mid z_{S\theta}, z_O\Big] + \eta(\tilde{\theta} - \hat{\theta}_O)^2, \qquad \eta \geq 0, \tag{1}$$

where η is a reduced-form cost of visibly diverging from the upstream baseline.[7] The unique minimiser is

$$\tilde{\theta}^* = \frac{\hat{\theta}_B + \eta \hat{\theta}_O}{1 + \eta}. \tag{2}$$

If $\eta > 0$ and $\hat{\theta}_B \neq \hat{\theta}_O$, then $\tilde{\theta}^* \neq \hat{\theta}_B$: the downstream publication is pulled toward the upstream baseline.

Proof. See Online Appendix Part II.

Interpretation and governance content. Sequential publication has two distinct effects. It allows the second mover to learn from the first publication, but it can also make visible disagreement costly. Proposition 1 isolates the second effect. The accuracy benchmark is $\hat{\theta}_B$, which uses both available signals. When $\eta > 0$, the published baseline instead compromises between that benchmark and the number already published by O. Observed

[7] Formally, this isolates the perceived-process term $\mathcal{D}_S = d_S(\cdot)$ by setting $\mathcal{D}_S = (\tilde{\theta} - \hat{\theta}_O)^2$ and holding the remaining components of U_S fixed.

convergence therefore need not represent common information or genuine agreement; it may partly reflect the institutional cost of being seen to disagree.

The governance implication is not that sequential calendars are always undesirable. Sequencing may be valuable when the first publication conveys information. The requirement is that the interface make justified divergence intelligible. This can be achieved either by publishing sufficiently close to simultaneously or, under sequencing, by disclosing conditioning sets, data vintages and scope boundaries clearly enough that differences are not automatically interpreted as institutional failure.

Proposition 2 (Assumptions concordat: accuracy, common-mode exposure, and reconciliation risk)[8]

Environment and regimes. Maintain the baseline decomposition $y \equiv \theta + \sigma$, with $\theta \perp \sigma$ and Gaussian signals as in Proposition 1. Let

$$\hat{\theta}_O \equiv \mathbb{E}[\theta \mid z_O], \qquad \hat{\theta}_S \equiv \mathbb{E}[\theta \mid z_{S\theta}], \qquad \hat{\sigma}_S \equiv \mathbb{E}[\sigma \mid z_{S\sigma}],$$

and write the corresponding posterior error variances as

$$V_O \equiv \mathrm{Var}(\theta - \hat{\theta}_O), \qquad V_S \equiv \mathrm{Var}(\theta - \hat{\theta}_S), \qquad V_\Sigma \equiv \mathrm{Var}(\sigma - \hat{\sigma}_S).$$

Institution S publishes $\hat{y}_S = \tilde{\theta} + \hat{\sigma}_S$. Under no concordat, $\tilde{\theta} = \hat{\theta}_S$; under a concordat, $\tilde{\theta} = \hat{\theta}_O$. Reconciliations are

$$R \equiv (y - \hat{y}_S) + \lambda(\theta - \hat{\theta}_O).$$

[8]Relative to the general payoff $U_S = -\ell_S - c_S(e_S) + \beta_S \mathbb{E}[\mathcal{R}_S] - \chi_S \mathbb{E}[\mathcal{D}_S]$, Proposition 2 holds effort and reputational incentives fixed and isolates how the protocol choice "concordat versus no concordat" shifts the statistical structure of forecast errors entering ℓ_S and the reconciliation object R. Transparency $\mathcal{T}$ matters through $\mathcal{D}_S$: when the concordat is publicly legible as a conditioning-set restriction, alignment in θ is less likely to be interpreted as undue dependence even though it mechanically increases error comovement.

The proposition asks how replacing a separately formed estimate of the shared component with the upstream baseline affects average accuracy, error covariance, and reconciliation volatility.

Statement. Assume the signal noises are independent across institutions and components. Define the finite-prior covariance between the two separately formed posterior errors by

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

Although the signal noises are independent, $C_{OS} > 0$ because both forecast errors retain a common component of prior uncertainty about θ.

(i) Accuracy. If $V_O \leq V_S$, then the concordat weakly lowers S's mean-squared forecast error for y:

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

(ii) Additional common-mode exposure. Under no concordat,

$$\mathrm{Cov}\Big(\theta - \hat{\theta}_O,\ \theta - \hat{\theta}_S\Big) = C_{OS} > 0.$$

Under the concordat the shared-component errors coincide, so the covariance becomes

$$\mathrm{Cov}\Big(\theta - \hat{\theta}_O,\ \theta - \hat{\theta}_O\Big) = V_O.$$

The concordat therefore raises shared-component 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

$$\mathrm{Var}(R^N) = V_S + V_\Sigma + \lambda^2 V_O + 2\lambda C_{OS}, \qquad \mathrm{Var}(R^C) = (1+\lambda)^2 V_O + V_\Sigma.$$

Hence the concordat reduces reconciliation variance if and only if

$$V_O + 2\lambda\big(V_O - C_{OS}\big) \leq V_S.$$

Proof. See Online Appendix Part II.

Interpretation and governance content. An assumptions concordat changes two objects at once. It can improve marginal forecast accuracy when the upstream estimate is sufficiently precise. It also raises common-mode exposure. That is, a finite common prior already makes separately formed posterior errors positively correlated, and the concordat increases that covariance from C_{OS} to V_O by replacing two partially distinct errors with the same error. The condition

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

identifies when the precision gain is large enough to offset the *additional* covariance generated by the concordat in the reconciliation object.

This is why concordats cannot be evaluated by comparing forecast errors institution by institution. Their fiscal effect depends on the joint distribution of errors and on the exposure parameter λ. A concordat can therefore be stabilising in one fiscal architecture and destabilising in another, even with the same underlying forecasts.

The governance implication is then that shared assumptions are most defensible when their scope is explicit, the upstream input is demonstrably strong, and the remaining margins of judgement stay genuinely independent. High transparency $\mathcal{T}$ makes that boundary observable and alignment on specified inputs can then be understood as coordination by design rather than evidence of capture.

Proposition 3 (Calendar under late data: correlation with little additional information)

Environment. Maintain $y \equiv \theta + \sigma$, with $\theta \sim \mathcal{N}(0, V_\theta)$, $\sigma \sim \mathcal{N}(0, V_\sigma)$, and $\theta \perp \sigma$. The devolved institution observes $z_{S\sigma} = \sigma + \varepsilon_{S\sigma}$, where

$\varepsilon_{S\sigma} \sim \mathcal{N}(0, w_S)$, and forms $\hat{\sigma}_S \equiv \mathbb{E}[\sigma \mid z_{S\sigma}]$. Its posterior error variance

$$V_\Sigma \equiv \mathrm{Var}(\sigma - \hat{\sigma}_S)$$

is common across calendar regimes and therefore enters reconciliation variance as an additive constant.

Let $A \in \{H, L\}$ denote data readiness, with $\Pr(A = H) = \pi$ and $\Pr(A = L) = 1 - \pi$. Conditional on $A = a$, the institutions observe independent signals

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

where $v_i^H < v_i^L$: on-time data are more informative. Define

$$\hat{\theta}_O^a \equiv \mathbb{E}[\theta \mid z_O^a], \qquad V_O^a \equiv \mathrm{Var}(\theta - \hat{\theta}_O^a),$$
$$\hat{\theta}_S^a \equiv \mathbb{E}[\theta \mid z_S^a], \qquad V_S^a \equiv \mathrm{Var}(\theta - \hat{\theta}_S^a),$$

and

$$\hat{\theta}_B^a \equiv \mathbb{E}[\theta \mid z_S^a, z_O^a], \qquad V_B^a \equiv \mathrm{Var}(\theta - \hat{\theta}_B^a) = \left(\frac{1}{V_\theta} + \frac{1}{v_S^a} + \frac{1}{v_O^a}\right)^{-1}.$$

The covariance between the two one-signal posterior errors is

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

Institution S publishes $\hat{y}_S^a = \tilde{\theta}^a + \hat{\sigma}_S$, and reconciliation is

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

Under simultaneous publication, S does not observe $\hat{\theta}_O^a$ and sets $\tilde{\theta}_{\mathrm{Sim}}^a = \hat{\theta}_S^a$. Under sequential publication, S observes $\hat{\theta}_O^a$ before publishing and minimises

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

Hence, as in Proposition 1,

$$\tilde{\theta}_{\mathrm{Seq}}^a = \frac{\hat{\theta}_B^a + \eta \hat{\theta}_O^a}{1 + \eta}.$$

Let $\alpha \equiv \eta/(1 + \eta) \in [0, 1)$.

Statement. For each $a \in \{H, L\}$, reconciliation variance under simultaneity is

$$\mathrm{Var}(R^a_{\mathrm{Sim}}) = V^a_S + \lambda^2 V^a_O + 2\lambda C^a_{OS} + V_\Sigma. \tag{3}$$

Under sequential publication,

$$\mathrm{Var}(R^a_{\mathrm{Seq}}) = V_\Sigma + (\lambda + \alpha)^2 V^a_O + c(\lambda, \eta) V^a_B, \qquad c(\lambda, \eta) \equiv \frac{1 + 2\eta + 2\lambda(1 + \eta)}{(1 + \eta)^2}. \tag{4}$$

The conditional variance wedge is therefore

$$\Delta^a \equiv \mathrm{Var}(R^a_{\mathrm{Seq}}) - \mathrm{Var}(R^a_{\mathrm{Sim}}) = (2\lambda\alpha + \alpha^2) V^a_O + c(\lambda, \eta) V^a_B - V^a_S - 2\lambda C^a_{OS}, \tag{5}$$

and the ex ante wedge is affine in π:

$$\Delta(\pi) \equiv \mathbb{E}[\Delta^A] = \pi \Delta^H + (1 - \pi)\Delta^L. \tag{6}$$

Threshold implication. If $\Delta^L > 0$ and $\Delta^H < 0$, there is a unique $\bar{\pi} \in (0, 1)$ such that

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

with

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

Under these sign conditions, simultaneity weakly dominates sequencing in expected reconciliation variance when the high-readiness state is sufficiently unlikely. Online Appendix Part II gives a primitive sufficient condition showing that this sign configuration arises on a non-empty parameter region.

Scottish benchmark. To anticipate our subsequent discussion we note in passing here that for the exact Scottish loading, when $\lambda = -1$ as we explain in more detail presently, the reversal described above does not arise. In each readiness state,

$$\mathrm{Var}(R^a_{\mathrm{Seq}}) = V_\Sigma + \frac{V^a_O - V^a_B}{(1 + \eta)^2} < V_\Sigma + V^a_S + V^a_O - 2C^a_{OS} = \mathrm{Var}(R^a_{\mathrm{Sim}}),$$

for non-degenerate signals. Sequencing therefore lowers reconciliation variance in the exact-cancellation benchmark. Proposition 3 is a therefore a general interface result. Thus, in the Scottish case the sharper tension is between cash stability and the amount of additional information preserved, not a reversal in the cash ranking of the two calendars.

Proof. See Online Appendix Part II.

Interpretation and governance content. Sequencing generates an informational benefit and changes the covariance structure of the reconciliation. The pure Bayesian gain from adding the O (OBR) signal to the S (SFC) information set is

$$G^a \equiv V_S^a - V_B^a = \frac{C^a}{A^a(A^a + C^a)}, \qquad A^a \equiv \frac{1}{V_\theta} + \frac{1}{v_S^a}, \quad C^a \equiv \frac{1}{v_O^a}.$$

This gain rises with upstream precision C^a but falls with the precision already available to the SFC, represented by A^a. Because both institutions' signals improve in the high-readiness state, $G^H > G^L$ is not automatic. The informational value of sequencing is higher under high readiness if and only if

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

Thus late data reduce the marginal informational value of sequencing when the upstream deterioration is sufficiently large relative to the deterioration in the SFC's own signal.

The wedge in (5) reflects that the simultaneous one-signal errors are already positively correlated through the finite common prior. Sequential publication changes that baseline covariance because the SFC observes the upstream signal and, when $\eta > 0$, is additionally pulled toward the upstream publication. The relevant case is therefore one in which additional alignment or common exposure survives when the incremental information conveyed

by sequencing is small. The simultaneous benchmark already contains the covariance implied by the common prior.

The design implication is to manage inputs and readiness rather than press for output convergence. Calendar discipline should reduce publication on incomplete information sets, while auditable disclosure of data vintages and costing status should make genuine informational differences legible. The threshold result remains conditional on the state-specific wedges satisfying $\Delta^L > 0$ and $\Delta^H < 0$.

Preliminary synthesis

Propositions 1–3 identify three distinct channels through which the interface shapes outcomes. Proposition 1 shows that sequencing can generate anchoring when visible disagreement is costly. Proposition 2 shows that concordats can improve marginal accuracy while adding to common-mode risk. Proposition 3 identifies the conditions under which deterioration in upstream information erodes the informational benefit of sequencing while sequencing and disagreement incentives may still add common exposure.

The common lesson is that independence and cooperation are jointly produced by protocol design. The interface determines the information available at publication, the covariance structure of errors, and what outsiders can verify about the process. Good design therefore aligns inputs that should genuinely be common, preserves independent judgement where it adds information, and makes the boundary between the two transparent.

The next section applies these mechanisms to a simplified version of the distinctive reconciliation arithmetic of the Scottish fiscal framework.

4 SFC and the OBR: the Scotland–UK mapping

The Scotland–UK fiscal framework gives the general reconciliation mechanism introduced above a particularly transparent form. Scotland's devolved budget contains three conceptually distinct elements: devolved own-source net revenues, the block grant (BG), and the block grant adjustment (BGA). The BG is determined mechanically from UK spending plans through the Barnett formula and is not reconciled against rest-of-UK revenue outturns. The BGA instead represents the counterfactual revenues that the UK government would have received had the relevant tax powers not been devolved. It is therefore forecast initially and subsequently reconciled against outturn.

The Scottish net position consequently depends on two forecast-sensitive objects: devolved revenues and the BGA. A positive error in the devolved-revenue forecast raises the eventual reconciliation, while an otherwise identical positive error in the BGA forecast lowers it because the BGA is deducted from the block grant. This opposite loading is the central feature of the Scottish case. In the reduced-form reconciliation equation, it implies $\lambda = -1$.

A corollary to Proposition 2: the Scottish BGA

Maintain the decomposition

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

where y denotes devolved Scottish net revenues, θ is a shared component that also drives the BGA benchmark, and σ is a devolved-specific component. Identify the BGA object with the shared component,

$$B \equiv \theta,$$

and let the relevant forecasts be

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

where $\widehat{\theta}_O$ is the upstream forecast used operationally for the BGA. Define forecast errors by $e_S \equiv y - \widehat{y}_S$ and $e_B \equiv B - \widehat{B}$. The reconciliation of the Scottish net position is therefore[9]

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

Corollary (BGA and mechanical cancellation under a concordat). Under the Gaussian linear environment above, consider an assumptions concordat on the shared component that imposes

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

Then:

(i) The reconciliation simplifies to the devolved-specific forecast error, $R = \sigma - \widehat{\sigma}_S$, so the shared component θ drops out of the reconciliation arithmetic entirely.

(ii) Relative to any regime in which $\widehat{\theta}_S \neq \widehat{\theta}_O$ with positive probability and $\widehat{\theta}_O - \widehat{\theta}_S$ is uncorrelated with $\sigma - \widehat{\sigma}_S$, the concordat weakly reduces reconciliation variance, with strict inequality whenever $\widehat{\theta}_O - \widehat{\theta}_S$ is non-degenerate:

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

Proof.

See Online Appendix Part III.

[9] To emphasise, in the generic reduced form $R = e_S + \lambda e_B$. Thus, the Scotland–UK mapping has $\lambda = -1$.

Interpretation. The corollary is a result about cash arithmetic and not that alignment is informationally efficient. Because the BGA enters the Scottish net position with the opposite sign to devolved revenues, common treatment of the shared component removes that component from the reconciliation mechanically. The immediate cash adjustment is then driven only by the devolved-specific error $\sigma - \widehat{\sigma}_S$.

That cancellation can be valuable as it reduces short-run reconciliation volatility and the associated pressure on fiscal management. But it does not establish that the common shared-component forecast is the most accurate forecast available to the SFC, nor that independent analysis of θ has no value. A concordat may therefore improve budget stability while reducing the independent information contained in the devolved forecast. The Scottish case makes especially clear the distinction developed in the next section between a best-forecast objective and a budget-stability objective.[10]

5 The costs and benefits of an independent fiscal institution

This section connects the protocol mechanisms in Propositions 1–3 to the institutional rationale for an independent devolved forecaster. The key point is that devolved systems create (at least) two distinct performance criteria. Independent fiscal institutions (IFIs) are typically mandated to "forecast well": to publish best estimates of revenues and related aggregates given stated policy and available information. Governments, by contrast, care as much, if not more, about the realised budget position and therefore about the volatility of cash adjustments when outturn differs from what was budgeted. When forecast errors turn into cash reconciliations, the cash-management (tax-smoothing) objective depends on the joint distribution of forecast errors, not on forecast accuracy in isolation.

[10] Online Appendix Part III provides the full mapping and derivation.

5.1 Reconciliation risk and competing objectives

Let the devolved fiscal-relevant outcome be y_t and let θ_t denote a shared/central component that also enters devolved cash arithmetic through the fiscal framework. Let O publish $\hat{\theta}_{O,t}$ and S publish $\hat{y}_{S,t}$. Define forecast errors

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

The reduced-form reconciliation object is

$$R_t \equiv e_{S,t} + \lambda e_{O,t} = (y_t - \hat{y}_{S,t}) + \lambda(\theta_t - \hat{\theta}_{O,t}), \tag{7}$$

where λ captures how the shared component loads into the devolved net position.[11] A simple identity makes clear why protocol design matters for fiscal management:

$$\mathrm{Var}(R_t) = \mathrm{Var}(e_{S,t}) + \lambda^2 \mathrm{Var}(e_{O,t}) + 2\lambda \,\mathrm{Cov}(e_{S,t}, e_{O,t}). \tag{8}$$

The covariance term is absent from evaluations of "forecast accuracy in isolation" but becomes first-order once reconciliations are cash-relevant.

Forecast accuracy and budget stability. We distinguish two natural objectives that can arise in devolved settings.

Forecasting objective (IFI mandate). A mandate to "forecast well" is naturally read as minimising expected squared forecast error for devolved outcomes conditional on the information available at publication time:

$$\hat{y}_{S,t}^{acc} \in \arg\min_{\hat{y} \in \mathcal{I}_{S,t}} \mathbb{E}\big[(y_t - \hat{y})^2 \mid \mathcal{I}_{S,t}\big]. \tag{9}$$

Budget-stability objective (government cash-management / tax smoothing). A government responsible for managing within-year cash outcomes has reason

[11] The sign and magnitude of λ are not "chosen" by policymakers; they are determined by fiscal architecture and by the direction in which R_t is defined (outturn minus forecast versus forecast minus outturn). So, as noted earlier, in the Scotland/BGA specialisation, the relevant shared component enters with the opposite sign in the net position, which corresponds to $\lambda < 0$ in (7) once θ_t is identified with the BGA benchmark.

to minimise the volatility of the cash adjustment itself:

$$\hat{y}_{S,t}^{stab} \in \arg\min_{\hat{y}\in\mathcal{I}_{S,t}} \mathbb{E}\Big[R_t^2 \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big], \tag{10}$$

where $\hat{\theta}_{O,t}$ is observed by the devolved forecaster at publication time under sequential publication or an input-sharing protocol.

The wedge is immediate from (7)–(8): the forecasting objective concerns $\mathrm{Var}(e_{S,t})$ alone, while the budget-stability objective depends on two errors and their covariance. This is the channel through which a forecaster can become a stabilisation instrument.

5.2 Forecast-as-hedging and the value of independence

Lemma (stabilisation targets the net cash object). If $\hat{\theta}_{O,t}$ is observed when S chooses $\hat{y}_{S,t}$, then under (10) the optimal publication satisfies

$$\hat{y}_{S,t}^{stab} + \lambda\hat{\theta}_{O,t} = \mathbb{E}\Big[y_t + \lambda\theta_t \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big]. \tag{11}$$

Equivalently,

$$R_t^{stab} = \big(y_t + \lambda\theta_t\big) - \mathbb{E}\Big[y_t + \lambda\theta_t \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}\Big], \tag{12}$$

so reconciliations are a conditional-mean-zero residual for the *net* object that enters the cash arithmetic.

Interpretation. Under cash-management motives, the relevant forecast target is not y_t per se but $y_t + \lambda\theta_t$. The central publication becomes a hedge object: the devolved publication loads mechanically against $\hat{\theta}_{O,t}$ through (11). A government forecaster can therefore optimally publish a number that is "biased" as a forecast of y_t relative to the posterior mean of y_t, because the publication is being used to smooth the cash adjustment rather than to maximise informational content.

Proof. See Online Appendix Part IV.

Why covariance is important.

Proposition 4: devolution makes covariance a first-order fiscal object. The reconciliation variance decomposes as in (8). Consequently:

1. Under the forecasting objective (9), the devolved publication minimises $\mathrm{Var}(e_{S,t})$ but does not in general control the covariance term in (8).

2. Under the budget-stability objective (10), the devolved publication is chosen to manage the joint distribution of errors, and in particular the covariance term $2\lambda \,\mathrm{Cov}(e_{S,t}, e_{O,t})$.

3. The wedge between the two objectives disappears on the shared component if either:

   (i) Informational coincidence: the central baseline is a sufficient statistic for θ_t given the devolved information at publication time, so that

   $$\mathbb{E}[\theta_t \mid \mathcal{I}_{S,t}, \hat{\theta}_{O,t}] = \hat{\theta}_{O,t}. \tag{13}$$

   (ii) Institutional coincidence: an explicit assumptions concordat defines the conditioning set by design (inputs shared by design), so that (13) holds by governance choice rather than as an empirical claim. In this case alignment on the shared component is not a distortion; it is transparent conditioning.

Proof. Parts (1)–(2) follow directly from (8) and the preceding lemma. Part (3) follows because (13) implies that the accuracy-optimal and stability-optimal conditional means coincide on θ_t given the permitted information set. □

The value and limits of an independent devolved IFI. The Lemma and Proposition 4 clarify the economic value of an independent devolved IFI. If the devolved forecaster is the government itself and it cares about tax smoothing, it has a structural incentive to use the published forecast as a stabilisation device: choose the publication to hedge against the central baseline and damp reconciliations, even at the expense of informational content. An independent IFI with an accuracy mandate removes this incentive to "hedge via publication" and therefore strengthens accountability and learning: the published number is a best estimate rather than an instrument.

The same results also explain why independence does not automatically deliver cash stability. Because the cash-adjustment objective depends on covariance, an accuracy-mandated IFI may publish a forecast with high informational content yet generate volatile reconciliations unless the interface is designed to manage the joint distribution of errors transparently. This is precisely where protocol design and fiscal buffers complement independence: shared conditioning sets where legitimate, calendar discipline and data-readiness disclosure to avoid publishing on thin information, and reserves/borrowing rules to absorb residual volatility.

Finally, the sign and interpretation of λ matters for how alignment should be read. In many transfer-based settings where shared components enter with the same sign, induced comovement can increase common-mode risk and raise reconciliation volatility. In the Scotland/BGA specialisation, where the shared component enters with the opposite sign, alignment on the shared component can compress reconciliation volatility through cancellation. This has a consequence for governance–stable reconciliations are not necessarily evidence of unbiased forecasts. If political narrative pressure or weak baselines bias the shared component, cancellation can mask the bias in the short run even while storing up adjustment risk later (for example, via subsequent revisions or policy-driven corrections). This is one reason

high transparency about conditioning sets and data readiness is a central ingredient of "coordination without capture."

The next subsection returns to the protocol mechanisms of Sections 2–3 and makes what we might call the "leaning" channel operational, yielding an explicit weight on the central baseline as a function of λ, signal precision, and timing constraints. By the leaning channel, we mean that the downstream forecast may be pulled towards the upstream forecast because doing so reduces expected cash reconciliations, not because the upstream forecast is necessarily more accurate.[12]

5.3 Making the 'leaning' channel explicit

The previous Lemma characterises the stabilisation objective in terms of the net cash object, but it does not by itself pin down how the devolved institution's *shared-component* publication moves relative to the central baseline. To make that comparative static operational, it is useful to impose a simple signal structure. In linear-Gaussian environments, the stabilisation-optimal choice of the shared-component forecast can be written in a transparent linear-combination form: the devolved publication is a linear combination of its own accuracy-optimal estimate and the observed central baseline,

$$\hat{\theta}_{S,t}^{stab} = (1-\omega)\,\hat{\theta}_{S,t}^{acc} + \omega\,\hat{\theta}_{O,t}. \tag{14}$$

Here $\hat{\theta}_{S,t}^{acc}$ denotes the devolved institution's accuracy-optimal estimate of the shared component based on its own signal alone, while $\omega = \omega(\lambda; V\theta, v_S, v_O)$ is the coefficient on the observed central baseline. It captures how the

[12]It is useful to distinguish four conceptually separate mechanisms. Under *learning*, the S moves towards the O because the latter's forecast contains useful information. Under *herding* or *anchoring*, it does so because visible disagreement is itself costly. Under the *leaning* channel, alignment is attractive because it hedges the subsequent reconciliation. Finally, under *capacity-driven reliance*, the S uses the O baseline because time or analytical constraints make the production of a fully independent forecast impractical. See the Online Appendix for the details.

reconciliation loading and the relative precision of the two signals shape the stabilisation-optimal publication. Timing does not enter separately in this Gaussian specialisation, but matters insofar as time compression or data constraints alter the feasible signal variances, particularly v_S. The coefficients sum to one. In the Gaussian specialisation below, for $-1 \leq \lambda \leq 0$, 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. The key comparative static is that ω increases as λ becomes more negative and when time compression or data constraints reduce the devolved institution's feasible precision on θ. In those cases, "importing" or leaning toward the central baseline becomes optimal even under a mandate-compliant forecasting rule, because the central publication functions as a hedge object in the cash arithmetic.

Gaussian specialisation: closed-form leaning weight. Assume the shared component follows

$$\begin{aligned} \theta_t &\sim \mathcal{N}(0, V_\theta), \\ z_{S,t} &= \theta_t + \varepsilon_{S,t}, \qquad \varepsilon_{S,t} \sim \mathcal{N}(0, v_S), \\ z_{O,t} &= \theta_t + \varepsilon_{O,t}, \qquad \varepsilon_{O,t} \sim \mathcal{N}(0, v_O), \\ \sigma_t &\perp \theta_t,\ z_{S,t},\ z_{O,t}, \end{aligned} \tag{15}$$

independently across t. The central institution publishes

$$\hat{\theta}_{O,t} = \mathbb{E}[\theta_t \mid z_{O,t}] = \frac{V_\theta}{V_\theta + v_O}\, z_{O,t}, \tag{16}$$

and the devolved forecaster observes $(z_{S,t}, \hat{\theta}_{O,t})$ before publishing. Let $\hat{\theta}^{acc}_{S,t} \equiv \mathbb{E}[\theta_t \mid z_{S,t}] = \frac{V_\theta}{V_\theta + v_S} z_{S,t}$ denote the accuracy-optimal shared-component estimate based on S's own signal.

Lemma: linear-combination form and comparative statics Under (15)–(16), the stabilisation-optimal shared-component publication satisfies

$$\hat{\theta}_{S,t}^{stab} = (1+\lambda)\, w_S\, \hat{\theta}_{S,t}^{acc} + \omega(\lambda; V_\theta, v_S, v_O)\, \hat{\theta}_{O,t}, \tag{17}$$

where

$$\begin{aligned} w_S &\equiv \frac{v_O\,(V_\theta + v_S)}{v_S v_O + V_\theta(v_S + v_O)}, \\ \omega(\lambda; V_\theta, v_S, v_O) &\equiv (1+\lambda)\,\frac{v_S\,(V_\theta + v_O)}{v_S v_O + V_\theta(v_S + v_O)} - \lambda. \end{aligned} \tag{18}$$

Moreover,

$$\frac{\partial \omega}{\partial \lambda} = \frac{v_S\,(V_\theta + v_O)}{v_S v_O + V_\theta(v_S + v_O)} - 1 < 0, \tag{19}$$

so the stabilisation motive places more weight on the central baseline as λ becomes more negative. The two coefficients in (17) sum to one. For $-1 \le \lambda \le 0$ they are both non-negative, so the linear combination is a genuine weighted average; outside that range one of the coefficients may lie outside $[0, 1]$.

Interpretation. Equation (17) captures "leaning": stabilisation induces a systematic shift in the weight on the central baseline, governed by $\omega(\lambda; \cdot)$. In a Scotland/BGA-style mapping, the relevant loading is negative, so $\lambda < 0$ raises the incentive to align on the shared component.

Proof. See Online Appendix Part IV.

Here $\hat{\theta}^{,acc}S, t$ denotes the devolved institution's accuracy-optimal estimate of the shared component based on its own signal alone, while $\omega = \omega(\lambda; V\theta, v_S, v_O)$ is the coefficient on the observed central baseline. It captures how the reconciliation loading and the relative precision of the two signals shape the stabilisation-optimal publication. Timing does not enter separately in this Gaussian specialisation, but matters insofar as time compression or data constraints alter the feasible signal variances, particularly v_S. Taken together, the results in this section imply that once forecast errors become cash adjustments, what matters for tax smoothing and cash management

is not only the marginal accuracy of each forecast but how the errors move together. A government forecaster therefore has an incentive to use the central baseline as a hedge object, so alignment can arise for stabilisation reasons even when it sacrifices informational content. An independent devolved IFI is most valuable precisely because it blocks that temptation to use the forecast as an instrument—forcing published numbers to remain best estimates—while stability is better achieved by engineering the interface (timing, shared conditioning sets, data readiness, and buffers) transparently rather than by informal pressure to converge on outputs.

6 Late costings and endogenous "import" of the shared component

The previous sections take the quality of the institutions' private information as given. In practice, however, forecast precision also depends on whether the institution has enough time and analytical capacity to process the information available to it. This is particularly important when devolved policy costings arrive close to the publication deadline. A shorter processing window lowers the return to analytical effort and may make it optimal for the devolved institution to import the upstream baseline for the shared component θ, while concentrating its scarce capacity on the devolved-specific component σ.

This mechanism is distinct from Proposition 3. There, late information may add little to the downstream forecast, while sequential publication can still induce convergence through reputational or interpretive pressure. Here the relevant information may exist, but the devolved institution cannot process it fully before publication. Alignment on θ is therefore not necessarily herding or capture. It may be the best forecast attainable under a binding time-and-capacity constraint. The associated governance margin is correspondingly concrete: calendar spacing, timely costings, and transparent

data-readiness rules affect whether independent forecasting remains feasible in practice.

6.1 Environment and forecasting architectures

Each period t there are 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 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 $V_O > 0$ summarises upstream precision on θ_t. [13]

The devolved institution observes two private signals,

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

independent of each other and of (θ_t, σ_t).

Here $\tau_t \geq 0$ is the effective time between receipt of final devolved costings or policy information and the publication deadline. The effort choices $e_{\theta,t} \geq 0$ and $e_{\sigma,t} \geq 0$ represent deployable analytical capacity, subject to

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

More time and more effort raise precision:

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

For explicit comparative statics, let

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

This specification captures a simple but important interaction; effort is less productive when there is little time in which to deploy it.

[13] One can endogenise V_O using the same signal structure as for S. Treating it as exogenous keeps the focus on the devolved institution's capacity constraint.

Posterior variances and expected loss. Under normal priors and signals, posterior-mean forecasts minimise mean-squared error under quadratic loss. The devolved posterior variances are [14]

$$V_\theta^S(e_{\theta,t}, \tau_t) = \left(\frac{1}{V_\theta} + \frac{1}{v_\theta(e_{\theta,t}, \tau_t)} \right)^{-1},$$
$$V_\sigma^S(e_{\sigma,t}, \tau_t) = \left(\frac{1}{V_\sigma} + \frac{1}{w_\sigma(e_{\sigma,t}, \tau_t)} \right)^{-1}.$$

Expected mean-squared forecast error therefore equals expected posterior variance. The capacity problem can be written directly as a choice over the two variance components.

Independent and import architectures. The devolved institution can either forecast θ_t itself or import the upstream baseline. These alternatives define two forecasting architectures.

Independent architecture.

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

Import architecture.

$$\hat{y}_{S,t}^{\text{imp}} \equiv \hat{\theta}_{O,t} + \hat{\sigma}_{S,t}, \qquad \mathbb{E}\Big[(y_t - \hat{y}_{S,t}^{\text{imp}})^2\Big] = V_O + V_\sigma^S(e_{\sigma,t}, \tau_t).$$

Once S imports $\hat{\theta}_{O,t}$, effort devoted to independently estimating θ_t no longer improves its published forecast. In this reduced form, import therefore implies $e_{\theta,t} = 0$ and $e_{\sigma,t} = \bar{E}$. Importing the shared component is not a withdrawal from forecasting altogether; it is a reallocation of scarce analytical capacity toward the component on which the devolved institution retains a distinctive informational role.

[14] Detailed derivations are given in Online Appendix Part V.

6.2 Effort allocation and endogenous importing

Under a best-forecast mandate, the independent architecture solves

$$\min_{e_{\theta,t}\in[0,\bar{E}]} L^{\text{ind}}(e_{\theta,t};\tau_t) \equiv V_\theta^S(e_{\theta,t},\tau_t) + V_\sigma^S(\bar{E}-e_{\theta,t},\tau_t).$$

Using the posterior-variance formulas above gives

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

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

$$-\frac{\partial V_\theta^S(e_\theta,\tau)}{\partial e_\theta}\Big|_{e_\theta=e_\theta^*(\tau)} = -\frac{\partial V_\sigma^S(e_\sigma,\tau)}{\partial e_\sigma}\Big|_{e_\sigma=\bar{E}-e_\theta^*(\tau)}.$$

The condition has a straightforward interpretation. Capacity is allocated until the marginal reduction in forecast variance is equal across the two components. When τ is small, the return to effort is weak on both margins, and even modest differences in relative productivity can produce a sharp reallocation toward the component on which effort remains most valuable. Corner solutions are therefore economically meaningful rather than pathological.

It is straightforward to characterise when importing optimal. Under import, expected loss is

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

Let

$$L^{\text{ind}\,*}(\tau) \equiv \min_{e_\theta\in[0,\bar{E}]} L^{\text{ind}}(e_\theta;\tau)$$

denote the minimum loss attainable under independent forecasting. Import is optimal if and only if

$$L^{\text{imp}}(\tau) \le L^{\text{ind}\,*}(\tau).$$

This comparison includes both sides of the capacity trade-off. Importing replaces S's own estimate of θ_t with the upstream estimate, but it also releases effort that can be used to improve the forecast of σ_t.

A useful sufficient condition is

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

If the upstream baseline is at least as precise as the best estimate of θ_t that S could produce even by devoting all feasible effort to that component, then importing is optimal under a best-forecast mandate. With the parametric technology,

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

and

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

A reduction in τ therefore worsens the best shared-component precision attainable by S and makes the import condition easier to satisfy. De facto alignment can emerge endogenously even though the formal mandate continues to require the institution's best forecast.

6.3 Time compression and the wider interface

Calendar time is therefore not an innocuous logistical detail. A short processing window lowers feasible precision and changes the optimal composition of forecasting activity. It can become efficient for the devolved institution to import the common component while concentrating scarce capacity on devolved-specific analysis. This reduces de facto informational independence on θ without requiring pressure, collusion, or a departure from the institution's statutory mandate.

The fiscal effect depends on how the two forecasts enter the reconciliation arithmetic. If time compression induces alignment on θ, it changes the covariance of forecast errors as well as their marginal variances. In the Scotland–BGA specialisation, alignment on the shared component can mechanically compress reconciliations through cancellation. In transfer systems where the relevant errors enter with the same sign, the same

alignment can instead increase common-mode risk. The calendar therefore affects not only forecast quality but the joint error distribution that governs cash management.

Because τ is partly determined by when governments settle policy decisions and provide costings, it is a genuine interface tool/lever. Either tier can affect the other institution's feasible information set without issuing any overt instruction about the published number. This is why minimum-notice norms for material measures, explicit costing cut-offs, provisional tagging of late measures, and public data-readiness logs are may be more than administrative housekeeping. They preserve the practical feasibility of independent analysis and allow outsiders to distinguish forecast errors caused by limited capacity from those caused by judgement or competence.[15]

Further interface extensions. Two further extensions sharpen the economic case for treating the interface itself as a policy instrument. First, when the shared component contains several blocks, remit makes precision endogenous. In other words, the downstream institution may rationally invest more heavily in blocks with high devolved-fiscal leverage, even when those blocks receive less attention upstream. Wholesale alignment can then discard valuable local information, whereas selective, block-by-block concordats can preserve independence where it is most useful while still reducing avoidable duplication and reconciliation risk. Online Appendix Part V formalises this trade-off.

Second, political narrative pressure can create misleading convergence by shifting one or both published baselines in the same direction. Correlated forecast errors may then reflect a common pressure wedge rather than shared information or genuine agreement, so observed alignment is not by itself

[15]In related work, we investigate compressed timetables in a somewhat richer environment when delays can also be beneficial in formulating policy at the central level. See "Who Controls the Clock? Fiscal Calendars and Informational Independence" (in preparation) C. Leith, C. Nolan and G. Roy.

evidence of successful coordination. Online Appendix Part VIII develops this result. Parts VI and VII examine complementary governance mechanisms: repeated interaction can sustain valuable information sharing when short-run incentives would otherwise prevent it, while Bayesian learning from realised forecast performance can strengthen accountability but may also create incentives to manage visible errors rather than maximise information.

7 Discussion and conclusions

This paper argues that the interface between independent fiscal institutions should be viewed as an economic object. Calendars, assumptions concordats, pre-publication sharing and disclosure rules may shape the information embedded in forecasts, the covariance of forecast errors and, where errors trigger later cash adjustments, the volatility of realised budget positions. The relevant chain sketched in our various scenarios is

$$\begin{aligned} \text{interface design} &\longrightarrow \text{joint error distribution} \\ &\longrightarrow \text{cash outcomes} \\ &\longrightarrow \text{incentives and governance trade-offs.} \end{aligned}$$

Our contribution is to make that chain explicit to identify which features of the interface move it.

The central implication is that forecast accuracy cannot be evaluated in isolation from the fiscal arithmetic into which forecasts enter. In systems with reconciliations, covariance matters alongside marginal accuracy. Sequential publication may compress independent information by encouraging the downstream institution to lean toward the upstream forecast, yet the same convergence may reduce reconciliation volatility. An assumptions concordat may improve precision by fixing a common conditioning set, but it may also add to common-mode risk. When upstream information deteriorates sufficiently relative to the downstream institution's own information,

sequencing may add little useful content while still increasing exposure to a shared error. We found that similarity is neither automatically evidence of sound coordination nor automatically evidence of capture, and divergence is neither automatically informative nor automatically dysfunctional.

This conditionality underpins our notion of coordination without capture. The objective is not to minimise contact between institutions. It is to align inputs transparently where common treatment is efficient—definitions, data vintages and agreed conditioning assumptions—while preserving independent judgement where it adds information. Transparency makes that boundary observable. It allows outsiders to distinguish agreement produced by an explicit protocol from convergence produced by reputational pressure, political influence or an unobserved loss of analytical capacity.

The capacity extension adds a distinct vulnerability. When costings arrive late or the publication window is compressed, a downstream institution may rationally import the upstream baseline for the shared component and redirect scarce resources toward devolved-specific analysis. That response need not indicate weak independence or poor judgement; it may be the constrained optimum. But it means that time-to-deadline is itself an interface instrument. Minimum-notice norms, costing cut-offs, provisional treatment of late measures and public data-readiness logs affect the feasible precision of the downstream forecast and the extent to which apparent agreement reflects genuine analysis rather than resource compression.

The Scotland–UK framework sharpens these general results because the Block Grant Adjustment enters the Scottish net position with the opposite sign to devolved receipts. Common treatment of the shared component can therefore generate mechanical cancellation in the reconciliation. This may stabilise cash in the short run even when it reduces independent information or allows a common bias to pass through both forecasts. The lesson is not that independence should be traded away for smoother cash flows. It is that the informational and cash consequences of coordination must

be assessed separately, and that any system-stability alignment should be explicit, auditable and confined to the components for which it is justified.

The wider extensions reinforce the same point. Remit-driven differences in analytical precision create a case for selective rather than wholesale concordats. Repeated interaction can sustain valuable information sharing that would not arise in a one-shot setting. Learning from realised forecast performance can strengthen accountability, while also creating incentives to manage visible errors. Common political narrative pressure can generate correlated forecast failures even without direct coordination. Together these results caution against reading institutional quality directly from the observed distance between two published numbers.

Finally, there are implications for the SFC and the wider IFI protocol agenda. Our analysis gives operational content to established principles on access to information, transparency, resourcing and evaluation of IFIs (OECD, 2014, 2019). In a high-stakes devolved setting with limited smoothing capacity, calendar discipline and data-readiness disclosure are fiscal risk-management tools, not merely administrative good practice (OECD, 2025). Concordats should be selective and block-specific: upstream comparative advantage may justify common treatment of some components, while the devolved institution may rationally invest more heavily in blocks with high devolved-fiscal leverage. The governance task is to make that boundary legible, so that remit-consistent disagreement is explained and mechanical agreement is not mistaken for independent confirmation.

Devolution turns forecast errors into cash. Once that is recognised, the protocol linking independent institutions becomes part of the fiscal architecture itself. Its design determines how much independent information survives, how common errors are amplified or contained, and who ultimately bears the resulting risk. That is the sense in which coordination is not institutional dressing but an economic policy instrument.

A Real-world fiscal systems and the scope of the framework

This appendix provides broader institutional motivation for the paper's set-up. It sketches several multi-level fiscal systems that combine own-source revenues, intergovernmental transfers, forecast-based budgets and, in some cases, later cash settlements. The cases differ in institutional detail, but they show why the joint behaviour of forecasts and transfer adjustments is a natural object of analysis. We pay attention to the sources of net revenues, the existence of reconciliation between forecasts and outturns, and the identity of the bodies producing the forecasts used to formulate the initial budget.

Germany: the Länder within a rules-based tax-sharing and equalisation system. For the German Länder budgets are set *ex ante* using forecasts of (i) own and shared-tax revenues, (ii) spending pressures, and (iii) net receipts/payments under the intergovernmental system (VAT distribution, equalisation and related federal grants). In practice, the revenue baseline for the large shared taxes is anchored in a joint federal–Länder forecasting process (the *Arbeitskreis Steuerschätzungen*, including a regionalisation procedure to apportion forecast tax revenues across Länder), so the "shared component" of the tax base is closer to a common object than in OBR–SFC settings (Federal Ministry of Finance (Germany), ndc). The binding background constraint is the constitutional debt brake (with the Länder required, in normal times, to balance budgets without structural net borrowing), which makes forecast surprises fiscally salient (Federal Ministry of Finance (Germany), nda; Stability Council (Germany), nd).

Relative to Scotland, the main cash-flow difference is that Länder shares of the big joint taxes largely move with realised collections under statutory sharing rules, so (outside the equalisation system) there is less need for

a Scotland-style "forecast-based transfer, then mechanical true-up" on the core tax-sharing flow. The natural German analogue of the transfer-related reconciliation term in the stylised rule is instead the explicit *Abrechnung* (settlement) built into the fiscal equalisation framework: equalisation/VAT shares are implemented on a provisional basis and subsequently corrected when final bases are determined, with differences between provisional payments and final entitlements settled via transfers (Federal Republic of Germany, 2005; Federal Ministry of Finance (Germany), ndb). In our notation, one can therefore interpret the German "net reconciliation" as the wedge between final and provisional equalisation-related transfers (plus any residual forecast error in the Länd's own net position that must be absorbed through within-year adjustment under the debt brake), while preserving the paper's deliberately stylised, timing-free representation.

Spain: the Comunidades Autónomas under a forecast-based "advance-and-settlement" funding system. Spain too fits the same operational template as in the stylised rule of the paper: regional budgets are set *ex ante* using forecasts of (i) regionally assigned/ceded taxes, (ii) spending pressures, and (iii) net resources under the regional financing system. For the 15 "common-regime" Comunidades Autónomas, the core financing framework (Ley 22/2009) combines tax capacity from ceded/shared taxes with equalisation-style transfers (notably the *Fondo de Garantía de Servicios Públicos Fundamentales* and the *Fondo de Suficiencia Global*, alongside convergence funds) (Ministerio de Hacienda (Spain), nd). A distinctive institutional feature—amenable to our reduced-form treatment—is that a large part of the system is implemented through *entregas a cuenta* (monthly advances) during the year, followed by an explicit *liquidación* (final settlement) once the legally-defined bases and actual collections are known (Gobierno de España, 2022; Ministerio de Hacienda (Spain), 2025; Cuenca,

2016).[16]

In our notation, Spain provides a particularly direct interpretation of the transfer-related reconciliation term:

$$R_{i,t}^{ES} \equiv \Big(\text{Liquidación}_{i,t} - \text{Entregas-a-cuenta}_{i,t}\Big) + \Big[(T_{i,t} - G_{i,t}) - (\widehat{T}_{i,t} - \widehat{G}_{i,t})\Big],$$

where the first bracket is a cash correction between final entitlements and within-year advances. That is, the true-up on the region's funding — positive if the region is owed extra money after the books are closed, negative if it was overpaid during the year. The second bracket is the region's own-budget forecast error, expressed in "net" form (Ministerio de Hacienda (Spain), 2025). This sits inside a broader rules-based fiscal framework applying to all tiers (including deficit/debt and expenditure-rule provisions) (Boletín Oficial del Estado, 2012). Spain's Independent Fiscal Institution, the *Autoridad Independiente de Responsabilidad Fiscal* (AIReF), fits the same broad pattern as our OBR–SFC discussion but with a different institutional structure: rather than two separate forecasters negotiating a shared component, AIReF is a single national IFI with an economy-wide remit (central, regional and local) that (i) reports on and can endorse/assess the macroeconomic forecasts underpinning public budgets, and (ii) monitors compliance with the fiscal rules and the operation of corrective mechanisms (Boletín Oficial del Estado, 2013; von Trapp et al., 2017). For our Bayesian implementation, we intentionally abstract from the (important) timing details of the Spanish settlement process and retain the reduced-form representation in which forecast errors feed into a stylised reconciliation term.

Australia: states, parliamentary budget offices, and transfer-determined resources. Australian state budgets are set *ex ante* on

[16]Two Comunidades operate under "foral" arrangements (Basque Country and Navarre), which follow different tax/transfer institutions; the mapping below is therefore best read as describing the common-regime system.

the basis of forecasts of own-source revenues (e.g. payroll taxes, property-related duties) and forecast spending pressures, but a material fraction of state resources is determined by intergovernmental transfers from the Commonwealth. The most prominent is the distribution of the national Goods and Services Tax (GST) pool as *general revenue assistance*, with state shares determined by population and GST "relativities" recommended by the Commonwealth Grants Commission (CGC) under the principle of horizontal fiscal equalisation (Commonwealth Grants Commission, nd, 2025; Australian Government, 2021). Several jurisdictions also now have parliamentary budget institutions whose credibility depends on statutory independence and access to executive information (for example the Victorian Parliamentary Budget Office and the NSW Parliamentary Budget Office, both of which have explicit legislative provisions supporting access to information held by government agencies) (Parliament of Victoria, 2017; Parliament of New South Wales, 2010).

This environment also reflects the paper's reduced-form "forecast budget + reconciliation" rule that we employ. Forecast errors in the state's own net position (own taxes relative to spending) matter for the realised balance, but a salient additional channel runs through transfer forecast errors because GST (and other Commonwealth payments) are large and are operationalised via within-year payments based on estimates with subsequent balancing adjustments. In particular, Commonwealth budget documentation explicitly reconciles GST revenue, GST entitlements and the advances paid to states, with prior-year balancing adjustments applied to later payments when entitlement differs from advances (Australian Government, 2025). This corresponds to a strictly positive transfer-related component in the net reconciliation term (a "settlement wedge" between final entitlements and advance payments).

Scotland: Block Grant Adjustment (BGA). As noted, the United Kingdom now provides a particularly interesting two-tier setting in which to study these issues. At the "central" tier, the Office for Budget Responsibility (OBR) produces the UK-wide macroeconomic and fiscal forecast used for the UK Budget, for assessing compliance with the UK's fiscal rules, and other variables of relevance to the devolved forecasters. The OBR also produce the forecast that underpin the Treasury's decision on the BGA element of the Scottish budget (OECD, 2020). At the devolved tier, the Scottish Fiscal Commission (SFC) is the official forecaster for the Scottish Budget and operates within a framework where forecast errors translate into later cash adjustments. The OECD's 2025 review describes Scotland as a high-stakes forecasting environment in which errors have direct budget consequences. It also highlights late information and compressed timescales, recommending stronger adherence to agreed deadlines and continued public reporting where deadlines are missed (OECD, 2025, p. 8).

Scotland's planned bond programme also gives institutional quality a market-facing role. The Scottish Government plans its first issuance in 2026–27 as part of a £1.5 billion programme (Scottish Government, 2025). In assigning Scotland an Aa3 rating with a stable outlook, Moody's noted that the establishment of the SFC had strengthened the fiscal framework and increased transparency around the fiscal outlook (Moody's Ratings, 2025; Roy, 2025). The point is that transparent forecasting arrangements can matter not only for parliamentary accountability but also for external assessments of fiscal credibility.

References

Alonso-Albarran, V., Arroyo, C., Aydin, O., Davoodi, H. R., Lam, W. R., Balasundharam, V., Hegab, G., Nguyen, A. D. M., Salazar, N., Sher, G., Solovyeva, A., and Tchelishvili, N. (2025). Fiscal rules and fiscal councils: Recent trends and revisions since the pandemic. IMF Working Paper 2025/198, International Monetary Fund.

Australian Government (2021). Intergovernmental agreement on federal financial relations: Schedule d (gst distribution). https://federalfinancialrelations.gov.au/sites/federalfinancialrelations.gov.au/files/2021-05/Schedule_D_April_13.pdf.

Australian Government (2025). Budget paper no. 3: Federal financial relations 2025–26. https://budget.gov.au/content/bp3/download/bp3_2025-26.pdf.

Bates, J. M. and Granger, C. W. J. (1969). The combination of forecasts. *Operational Research Quarterly*, 20(4):451–468.

Beetsma, R. M. W. J. and Debrun, X. (2016). Fiscal councils: Rationale and effectiveness. IMF Working Paper 2016/086, International Monetary Fund.

Boletín Oficial del Estado (2012). Ley orgánica 2/2012, de estabilidad presupuestaria y sostenibilidad financiera. BOE-A-2012-5730, https://www.boe.es/buscar/act.php?id=BOE-A-2012-5730.

Boletín Oficial del Estado (2013). Ley orgánica 6/2013, de creación de la autoridad independiente de responsabilidad fiscal. BOE-A-2013-11935, https://www.boe.es/eli/es/lo/2013/11/14/6.

Commonwealth Grants Commission (2025). Commission's assessment

methodology: Horizontal fiscal equalisation. Technical report, Australian Government.

Commonwealth Grants Commission (n.d.). About gst distribution. https://www.cgc.gov.au/about-gst-distribution.

Cuenca, A. (2016). Las entregas a cuenta en la financiación de las comunidades autónomas: algunas opciones de mejora. *Mediterráneo Económico*, 30:191–210.

European Commission (2025a). Independent fiscal institutions in the revised legal framework. Publications Office of the European Union, https://op.europa.eu/en/publication-detail/-/publication/6bd4c484-6127-11f0-a9d0-01aa75ed71a1/language-en.

European Commission (2025b). Report on the capacity and tasks of independent fiscal institutions in the EU. Institutional Paper 330, European Commission, Directorate-General for Economic and Financial Affairs.

Federal Ministry of Finance (Germany) (n.d.a). Germany's federal debt rule (debt brake). https://www.bundesfinanzministerium.de/Content/EN/Downloads/Public-Finances/germanys-federal-debt-rule.pdf.

Federal Ministry of Finance (Germany) (n.d.b). L"anderfinanzausgleich. https://www.bundesfinanzministerium.de/Content/DE/Standardartikel/Themen/Oeffentliche_Finanzen/Foederale_Finanzbeziehungen/Laenderfinanzausgleich/laenderfinanzausgleich.html.

Federal Ministry of Finance (Germany) (n.d.c). Working party on tax revenue estimates. https://www.bundesfinanzministerium.de/Content/EN/Standardartikel/Topics/Taxation/Articles/working-party-on-tax-revenue-estimates.html.

Federal Republic of Germany (2005). Finanzausgleichsgesetz 2005. https://www.gesetze-im-internet.de/finausglg_2005/BJNR395600001.html.

Gobierno de España (2022). Financiación autonómica: entregas a cuenta. La Moncloa, https://www.lamoncloa.gob.es/serviciosdeprensa/notasprensa/hacienda/Paginas/2022/300822-financiacion-autonomica.aspx.

Ministerio de Hacienda (Spain) (2025). Liquidación de los recursos del sistema de financiación autonómica 2023. https://www.hacienda.gob.es/CDI/Sist%20Financiacion%20y%20Deuda/Informaci%C3%B3nCCAA/Liquidacion-de-los-recursos-SFA-2023.pdf.

Ministerio de Hacienda (Spain) (n.d.). Sistema de financiación de las comunidades autónomas de régimen común. https://www.hacienda.gob.es/es-ES/CDI/Paginas/SistemasFinanciacionDeuda/InformacionCCAAs/Informes%20financiacion%20comunidades%20autonomas2.aspx.

Moody's Ratings (2025). Moody's ratings assigns aa3 rating to scotland; the outlook is stable. Rating action, https://ratings.moodys.com/ratings-news/454328.

Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92(5):1521–1534.

OECD (2014). Recommendation of the council on principles for independent fiscal institutions. OECD/LEGAL/0401, https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0401.

OECD (2019). Oecd review of the scottish fiscal commission 2019. Technical report, OECD Publishing, Paris.

OECD (2020). Oecd review of the uk office

for budget responsibility. Technical report, Organisation for Economic Co-operation and Development.

OECD (2025). Oecd review of the scottish fiscal commission 2025. *OECD Papers on Budgeting*, 2025(4).

Parliament of New South Wales (2010). Parliamentary budget officer act 2010, section 16. https://classic.austlii.edu.au/au/legis/nsw/consol_act/pboa2010269/s16.html.

Parliament of Victoria (2017). Parliamentary budget officer act 2017, section 26. https://classic.austlii.edu.au/au/legis/vic/consol_act/pboa2017269/s26.html.

Roy, G. (2025). Letter from the chair of the Scottish Fiscal Commission to the convener of the Finance and Public Administration Committee. Scottish Parliament, https://www.parliament.scot/-/media/files/committees/finance-and-public-administration-committee/correspondence/2025/chairsfctoconvener_19nov25.pdf. 19 November 2025; published 21 November 2025.

Scharfstein, D. S. and Stein, J. C. (1990). Herd behavior and investment. *American Economic Review*, 80(3):465–479.

Scottish Government (2025). Scottish government bonds launch planned. https://www.gov.scot/news/scottish-government-bonds-launch-planned/. 13 November 2025.

Scottish Government and United Kingdom Government (2023). Fiscal framework: Agreement between the scottish and uk governments. https://www.gov.scot/publications/fiscal-framework-agreement-between-scottish-uk-governments/. ISBN 9781835212103.

Stability Council (Germany) (n.d.). Monitoring compliance with the debt brake. https://www.stabilitaetsrat.de/EN/Aufgaben/Ueberwachung_Einhaltung_Schuldenbremse/Ueberwachung-Einhaltung-Schuldenbremse_node.html.

von Trapp, L., Nicol, S., Fontaine, I., Suyker, W., and Delivorias, A. (2017). Oecd review of the independent authority for fiscal responsibility (AIReF). Technical report, Organisation for Economic Co-operation and Development, Paris.

Wang, X., Hyndman, R. J., Li, F. and Kang, Y. (2023). Forecast combinations: An over 50-year review. *International Journal of Forecasting*, 39(4):1518–1547.