

INTERNATIONAL MONETARY FUND

Smooth Forecast Reconciliation

Sakai Ando

WP/24/66

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.**

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2024
MAR**



WORKING PAPER

IMF Working Paper
Research Department

Smooth Forecast Reconciliation
Prepared by Sakai Ando*

Authorized for distribution by Prachi Mishra
March 2024

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: How to make forecasts that (1) satisfy constraints, like accounting identities, and (2) are smooth over time? Solving this common forecasting problem manually is resource-intensive, but the existing literature provides little guidance on how to achieve both objectives. This paper proposes a new method to smooth mixed-frequency multivariate time series subject to constraints by integrating the minimum-trace reconciliation and Hodrick-Prescott filter. With linear constraints, the method has a closed-form solution, convenient for a high-dimensional environment. Three examples show that the proposed method can reproduce the smoothness of professional forecasts subject to various constraints and slightly improve forecast performance.

RECOMMENDED CITATION: Ando, Sakai (2024), "Smooth Forecast Reconciliation," IMF Working Paper.

JEL Classification Numbers:	C53, E17, E37
Keywords:	Smoothness; Forecast Reconciliation; Minimum Trace Reconciliation; Hodrick-Prescott filter; Cross-sectional; Temporal
Author's E-Mail Address:	sando@imf.org

* The author would like to thank Michal Andrle, Tommaso Di Fonzo, Jean-Jacques Forneron, Domenico Giannone, Shinya Kotera, Marco Marini, Petia Topalova, and participants in various seminars.

WORKING PAPERS

Smooth Forecast Reconciliation

Sakai Ando¹

¹ The author would like to thank Michal Andrie, Tommaso Di Fonzo, Jean-Jacques Forneron, Domenico Giannone, Shinya Kotera, Marco Marini, Petia Topalova, and participants in various seminars.

Contents

1.	Introduction.....	3
2.	Theory.....	5
2.1.	Minimum-Trace Reconciliation and Hodrick-Prescott Filter	5
2.2.	Closed-Form Solution with Given Smoothness Parameters	6
2.3.	Choice of Smoothness Parameters	9
3.	Applications	10
3.1.	Data.....	10
3.2.	Forecasting Quarterly GDP	11
3.2.1.	Input	11
3.2.2.	Output.....	13
3.3.	Mixed-Frequency: Forecasting Both Annual and Quarterly GDP	15
3.3.1.	Input	15
3.3.2.	Output.....	16
3.3.3.	Incorporating Ad-Hoc Information	17
3.4.	Cross-Sectional and Temporal Constraints: Adding Subcomponents	17
3.4.1.	Input	18
3.4.2.	Output.....	19
4.	Conclusion	20
	Annex I. Comparison with Modified Denton	22
	Annex II. Alternative First-Step Methods	23
	References.....	24
FIGURES		
	Figure 1. Times Series of US GDP in WEO.....	11
	Figure 2. Forecast of US quarterly GDP	13
	Figure 3. Forecast Performance Comparison	14
	Figure 4. Smoothness Comparison.....	15
	Figure 5. Forecast of US Annual and Quarterly GDP	16
	Figure 6. Forecast of US Annual and Quarterly GDP with Ad-Hoc Constraints.....	17
	Figure 7. Forecast of Mixed-Frequency GDP and Subcomponents	20
	Figure 8. Comparison with Modified Denton	22
	Figure 9. Forecast Performance Comparison When the First Step is AR(4)	23
	Figure 10. Forecast Performance Comparison When the First Step is AutoETS	23
TABLES		
	Table 1. Structure of Data and Forecast Horizon.....	11
	Table 2. Expanding Windows of Time Series Split	12

1. Introduction

Economists often need to impose constraints, such as accounting identities, when forecasting multiple macroeconomic variables to ensure internal consistency. At the same time, they also want each economic series to be smooth over time in the forecast horizon. For example, when economists forecast seasonally adjusted quarterly GDP, quarterly GDP forecasts need to aggregate up to the annual GDP forecast and evolve smoothly unless a shock is expected to hit the economy in a specific quarter. The problem shows up in other macroframework settings, including forecasting the monthly consumer price index (CPI), which needs to be consistent with annual CPI, the balance of payments accounts that consist of dozens of time series linked by accounting identities, etc. These problems are widely faced by macroeconomists in policy institutions, including central banks, ministries of finance, international financial institutions, and the private sector.

Both respecting constraints and ensuring smoothness over the forecast horizon are important for the quality of the forecast. Satisfying accounting identities, within and across time series, is a property of the true data and crucial for the internal consistency of the forecast. The smoothness over the forecast horizon is a property of the optimal forecast in many time series models, including random walk and autoregressions. Smoothness also helps forecasters communicate the story behind the forecast by suppressing the noise generated by algorithms and highlighting intentionally introduced kinks.

Existing methods, however, do not achieve both objectives. In the example of quarterly GDP forecast, economists often apply their favorite technique to the forecast of Q1, Q2, and Q3 GDP, and use Q4 GDP as the residual variable to manually enforce consistency with annual GDP. This approach ensures the forecast satisfies accounting identities but can generate unwanted kinks in Q4 and the following year's Q1. Forecasters often need to re-adjust all quarters manually until the quarterly path looks smooth and satisfies the accounting identity. Such manual adjustments could easily lead to mistakes and become prohibitively costly as the number of time series increases.

The literature on forecast reconciliation offers a more systematic approach to producing forecasts that satisfy linear constraints. Di Fonzo and Girolimetto (2023) and Girolimetto and Di Fonzo (2023) propose a projection method to satisfy linear cross-sectional and temporal constraints by extending the minimum-trace (min-T) reconciliation proposed by Wickramasuriya et al. (2019) and Athanasopoulos et al. (2017). Taieb (2017) proposes an approach based on regularization. Di Fonzo and Marini (2011) extend Denton (1971) in the context of national accounts statistics. The existing methods in the reconciliation literature, however, do not necessarily generate smooth time series. A separate strand of literature proposes various filtering methods to smooth time series, including the Hodrick and Prescott filter (HP filter) (Hodrick and Prescott, 1997), moving averages, and others, but they do not allow for forecasts to be subject to constraints. Applying the filtering and reconciliation methods repeatedly need not converge, leading to a whack-a-mole exercise: solving one problem creates another. To our knowledge, there is no paper that combines the two discrete strands of literature. The main contribution of this paper is to integrate them into a practical method to achieve both objectives.

This paper proposes a new forecast reconciliation method to impose both constraints and smoothness. The method integrates the min-T reconciliation and HP filter. Intuitively, the min-T component optimally determines how much each forecast horizon of each time series should be adjusted to satisfy the constraints imposed by the forecaster, while the HP component controls each time series' smoothness over the forecast horizon. Mathematically, since both min-T reconciliation and HP filter are quadratic programming, they can be naturally

integrated into one constrained minimization problem. The proposed method can accommodate any number of time series with any frequencies, any constraints between or within the time series, and any smoothness parameters of choice. Moreover, the proposed method has a closed-form solution, which makes it amenable to a high-dimensional environment.

We demonstrate the application of the proposed method using three examples. Each example compares the forecast generated by the proposed method, with the forecasts published in the International Monetary Fund (IMF) World Economic Outlook (WEO) database, and the forecasts generated by min-T reconciliation. We consider IMF WEO forecasts as the benchmark, which incorporates economists' extensive country-specific knowledge and quality controls, including the validation of accounting identities and unintended kinks in the forecast horizon. (Genberg et al., 2014)

In the first example, we take the annual US GDP series as fixed and forecast the seasonally adjusted quarterly GDP subject to the constraint that the quarterly GDP aggregates up to the annual number. We show that the proposed method can replicate the smoothness of the expert-generated WEO forecasts with slightly smaller forecast errors than the reconciled-but-not-smoothed alternative generated by min-T reconciliation. In other words, our method allows forecasts to be smoothed subject to constraints without loss in forecast performance and resource-intensive manual adjustments.

The proposed method can accommodate ad-hoc constraints in addition to imposing accounting identities. In the second example, we forecast both the annual and quarterly US GDP. We show that the forecast based purely on historical data can deviate substantially from the WEO forecasts since WEO forecasts incorporate ad-hoc information about the future such as recessions and recoveries in the short run and potential growth rates in the long run. We demonstrate that the proposed method can incorporate such ad-hoc information by imposing short-run and long-run constraints, which bring back the forecast closer to the one in the WEO.

The last example illustrates the flexibility of the proposed method by incorporating both cross-sectional and temporal constraints. We build on the second example by expanding the forecast to include the subcomponents of annual GDP. We show that the proposed method can smooth mixed-frequency multivariate time series subject to a wide range of constraints, including cross-sectional, temporal, and ad-hoc information.

It is important to note that the proposed method is a post-forecasting process: adopting reasonable forecasting methods and imposing reasonable constraints are crucial. The examples described in the paper intentionally rely on naïve forecasts for the first-step forecast to highlight the kinks introduced by min-T reconciliation. In practice, however, an accurate first-step forecast is important for forecast performance. The constraints should also be chosen wisely since imposing constraints improves forecast only when they are correct. In this sense, the proposed method is not a replacement but a complement to the existing forecasting methods. It could, however, replace some of the ex-post ad-hoc and resource-intensive manual adjustments to the first-step forecasts that many macroeconomists use to ensure internal consistency and smoothness of their forecasts.

The rest of the paper proceeds as follows. Section 2 discusses the theoretical justification of the proposed smooth forecast reconciliation method. Section 3 illustrates the application with three examples. Section 4 concludes.

2. Theory

2.1. Minimum-Trace Reconciliation and Hodrick-Prescott Filter

This section reviews the min-T reconciliation method and HP filter. We use unified notation to highlight that both methods take some data as exogenously given and adjust them to satisfy certain properties. Thus, they can be integrated into a single method as a natural extension.

The min-T reconciliation, proposed by Wickramasuriya et al. (2019), is a projection of a forecast on the space spanned by linear constraints. The distance used in the projection is weighted by the covariance matrix of the first-step forecast errors, which is optimal in the sense that it minimizes the sum of the variances of the second-step forecast errors when the first-step forecast is unbiased (Wickramasuriya et al., 2019; Ando and Narita, 2022).

Specifically, let $\hat{y} \in \mathbb{R}^N$ be the first-step forecast, $W = V(\hat{y} - y^*)$ be the variance of its error from the ground truth y^* , and (C, d) be a pair of the $K \times N$ matrix and $K \times 1$ vector in the linear constraints that the ground truth y^* satisfies. The second-step forecast $\tilde{y} \in \mathbb{R}^N$ is defined as

$$\tilde{y} = \arg \min_{y \in \mathbb{R}^N} (y - \hat{y})' W^{-1} (y - \hat{y}) \quad s. t. \quad Cy = d. \quad (1)$$

The solution can be written as

$$\tilde{y} = \hat{y} + WC'(CWC')^{-1}(d - C\hat{y}). \quad (2)$$

Intuitively, the second-step forecast \tilde{y} adjusts the first-step forecast \hat{y} to satisfy the constraints in a way that the i -th variable \hat{y}_i is adjusted more if the first-step forecast \hat{y}_i is less accurate. Geometrically, the second-step forecast \tilde{y} is an oblique projection of the first-step forecast \hat{y} on the space of constraints (Panagiotelis et al., 2023), so that applying the reconciliation twice is the same as applying it once. The true weight $W = V(\hat{y} - y^*)$ is unfeasible since it uses the ground truth y^* and becomes degenerate when the first-step forecast \hat{y} satisfies the constraints $C\hat{y} = d$. A non-singular estimate of the weight matrix $\hat{W} = \hat{V}(\hat{y} - y^*)$, however, can be obtained by various shrinkage methods (Ledoit and Wolf, 2004; Schafer and Strimmer, 2005; Chen et al., 2010; and Ando and Xiao, 2023), as illustrated in section 3.

The HP filter, developed by Hodrick and Prescott (1997), is a filter to smooth times series. Let $\hat{y} \in \mathbb{R}^T$ be a times series and λ be a smoothness parameter, both of which are exogenously given. The smoothed time series \tilde{y} is

$$\{\tilde{y}_t\}_{t=1}^T = \arg \min_{\{y_t\}_{t=1}^T} \sum_{t=1}^T (\hat{y}_t - y_t)^2 + \lambda \sum_{t=2}^{T-1} \{(y_{t+1} - y_t) - (y_t - y_{t-1})\}^2 = \arg \min_{y \in \mathbb{R}^T} (\hat{y} - y)' (\hat{y} - y) + \lambda y' F y, \quad (3)$$

where F is a $T \times T$ degenerate penta-diagonal matrix

$$F = \begin{bmatrix} 1 & -2 & 1 & 0 \\ -2 & 5 & -4 & 1 \\ 1 & -4 & 6 & -4 \\ 0 & 1 & -4 & 6 & \ddots \\ & & & & 6 & -4 & 1 & 0 \\ & & & & -4 & 6 & -4 & 1 \\ & & & & 1 & -4 & 5 & -2 \\ & & & & 0 & 1 & -2 & 1 \end{bmatrix}. \quad (4)$$

With a $T \times T$ identity matrix I_T , the solution can be written as

$$\tilde{y} = (I_T + \lambda F)^{-1} \hat{y}. \quad (5)$$

Intuitively, the smoothed time series \tilde{y} adjusts the original time series \hat{y} in a way that the difference in differences is minimized. The smoothness parameter λ for quarterly time series is typically chosen to be $\lambda = 1600$. For annual and monthly data, Hodrick and Prescott (1997) suggest 100 and 14400, while Ravn and Uhlig (2002) propose 6.25 and 129600.

Although the literature has debated whether the HP filter can estimate the trend accurately, this paper abstracts from this debate and interprets the HP filter as a smoothing device. For example, Hamilton (2018) proposes Maximum Likelihood estimates and points out drawbacks of the HP filter, including the endpoint problem. Hodrick (2020) and Dritsaki and Dritsaki (2022) offer counterarguments. Instead of assuming a data-generating process and discussing whether the HP filter estimates the trend accurately, this paper interprets the HP filter as a smoothing device and uses actual data to check the forecast performance.

The min-T reconciliation and HP filter methods have both similarities and differences. Both adjustment methods (1) take some data as exogenously given and (2) are quadratic programming. These similarities allow the two methods to be integrated as natural extensions. The second similarity also enables a closed-form solution, which is useful for high-dimensional data. The main differences between the two methods are that (1) min-T reconciliation can accommodate multiple time series by stacking them into a single vector, while the HP filter can handle only a single time series, and (2) min-T reconciliation weights variables, while the HP filter does not. Sections 2.2 and 2.3 combine the two methods by exploiting their similarities and bridge the differences by allowing multiple smooth parameters with appropriate units.

2.2. Closed-Form Solution with Given Smoothness Parameters

This section defines the smooth forecast reconciliation for a given set of smoothness parameters and derives its closed-form solution. The solution boils down to min-T reconciliation and HP filter in special cases.

Suppose there are M time series $\{y_m\}_{m=1}^M$ with potentially different lengths $\{T_m\}_{m=1}^M$. Different time series can have different frequencies or units. Let y denote the stacked vector and N denote its length.

$$y = \begin{bmatrix} y_1 \\ \vdots \\ y_M \end{bmatrix} = \begin{bmatrix} y_{11} \\ \vdots \\ y_{1T_1} \\ \vdots \\ y_{M1} \\ \vdots \\ y_{MT_M} \end{bmatrix}, \quad N = \sum_{m=1}^M T_m. \quad (6)$$

The problem takes three objects as exogenously given. The first object is the set of first-step forecast $\hat{y} \in \mathbb{R}^N$ and its associated covariance matrix of forecast errors $W = V(\hat{y} - y^*) \in \mathbb{R}^{N \times N}$ where y^* is the ground truth. The covariance matrix of forecast errors W is assumed to be invertible.

The first-step forecast \hat{y} can be obtained by any forecasting methods, including judgmental forecast. One caveat is that the optimality of min-T reconciliation assumes unbiasedness of the first-step forecast $E[\hat{y} - y^*] = 0$. The covariance matrix of forecast errors W is not feasible since it depends on the unknown ground truth $y^* \in \mathbb{R}^N$, but it can be replaced by a feasible estimate \hat{W} . One caveat is that the covariance matrix of forecast errors W is singular when both the first-step forecast \hat{y} and the ground truth y^* satisfy the same linear constraints. Section 3 shows that shrinkage methods can be used to estimate a non-singular covariance matrix of forecast errors \hat{W} .

The second object is the constraints (C, d) that the ground truth y^* satisfies.

$$Cy^* = d, \quad (7)$$

where C is a $K \times N$ matrix and d is a $K \times 1$ vector. The constraint matrix C is assumed to be full rank.

Two observations are noteworthy. First, as discussed in Ando and Kim (2023), nonlinear constraints often show up in macroeconomic statistics. Although a closed-form solution may not be available and numerical solution may be challenging to obtain in general, it is straightforward to extend the constraints to nonlinear $C(y^*) = 0$ or inequality $C(y^*) \geq 0$ constraints. Second, since the vector y stacks multiple time series, the formulation allows both intra-temporal/cross-sectional and inter-temporal constraints, as illustrated in section 3.

The third object is the symmetric smoothness matrix Φ , which is block-diagonal and consists of M smoothness parameters $\{\lambda_m\}_{m=1}^M$ and M penta-diagonal matrices $\{F_m\}_{m=1}^M$, where the size of the matrix F_m is $T_m \times T_m$, and the elements of each matrix F_m are analogous to (4).

$$\Phi = \begin{bmatrix} \lambda_1 F_1 & & \\ & \ddots & \\ & & \lambda_M F_M \end{bmatrix}. \quad (8)$$

Given the first-step forecast and its associated covariance matrix of forecast errors (\hat{y}, W) , the constraints (C, d) , and the smoothness matrix Φ , the second-step forecast \tilde{y} is

$$\tilde{y} = \arg \min_{y \in \mathbb{R}^N} (y - \hat{y})' W^{-1} (y - \hat{y}) + y' \Phi y \quad s. t. \quad Cy = d. \quad (9)$$

One can see that the formulation reduces to min-T reconciliation (1) when the smoothness parameters are nullified $\lambda_1 = \dots = \lambda_M = 0$. The formulation also reduces to an independent sum of M HP filters (3) when the weight is an identity matrix $W = I_N$ and the constraints are zeros $(C, d) = (0,0)$.

$$\min_{y \in \mathbb{R}^N} \sum_{m=1}^M \{(y_m - \hat{y}_m)'(y_m - \hat{y}_m) + \lambda_m y_m' F_m y_m\}. \quad (10)$$

The following theorem states that problem (9) has a closed-form solution, which is a convenient feature in a high-dimensional environment with a large N .

Theorem: The second-step forecast \tilde{y} that solves (9) can be written as

$$\tilde{y} = [I_N - (W^{-1} + \Phi)^{-1}C'\{C(W^{-1} + \Phi)^{-1}C'\}^{-1}C](W^{-1} + \Phi)^{-1}W^{-1}\hat{y} + (W^{-1} + \Phi)^{-1}C'\{C(W^{-1} + \Phi)^{-1}C'\}^{-1}d. \quad (11)$$

Proof: The first-order condition is sufficient since the problem is quadratic. The Lagrangian is

$$\mathcal{L} = (\tilde{y} - \hat{y})'W^{-1}(\tilde{y} - \hat{y}) + \tilde{y}'\Phi\tilde{y} + 2\lambda'(d - C\tilde{y}).$$

The first-order condition with respect to \tilde{y} is

$$2(W^{-1} + \Phi)\tilde{y} - 2W^{-1}\hat{y} - 2C'\lambda = 0 \Rightarrow \tilde{y} = (W^{-1} + \Phi)^{-1}(W^{-1}\hat{y} + C'\lambda).$$

From the constraint,

$$C\tilde{y} = d \Rightarrow d = C(W^{-1} + \Phi)^{-1}W^{-1}\hat{y} + C(W^{-1} + \Phi)^{-1}C'\lambda \Rightarrow \lambda = \{C(W^{-1} + \Phi)^{-1}C'\}^{-1}\{d - C(W^{-1} + \Phi)^{-1}W^{-1}\hat{y}\}.$$

By substitution,

$$\begin{aligned} \tilde{y} &= (W^{-1} + \Phi)^{-1}W^{-1}\hat{y} + (W^{-1} + \Phi)^{-1}C'\{C(W^{-1} + \Phi)^{-1}C'\}^{-1}\{d - C(W^{-1} + \Phi)^{-1}W^{-1}\hat{y}\} \\ &= \{I_N - (W^{-1} + \Phi)^{-1}C'\{C(W^{-1} + \Phi)^{-1}C'\}^{-1}C\}(W^{-1} + \Phi)^{-1}W^{-1}\hat{y} + (W^{-1} + \Phi)^{-1}C'\{C(W^{-1} + \Phi)^{-1}C'\}^{-1}d. \end{aligned}$$

Q.E.D.

Equation (11) is a natural extension of both min-T reconciliation (2) and HP filter formula (5). One can reconfirm that equation (11) reduces to min-T reconciliation (2) when the smoothness matrix is nullified $\lambda_1 = \dots = \lambda_M = 0$. Equation (11) also reduces to the stacked version of HP filter (5) when the weight is an identity matrix $W = I_N$ and the constraints are zeros $(C, d) = (0,0)$.

$$\tilde{y} = (I_N + \Phi)^{-1}\hat{y} = \begin{bmatrix} (I_{T_1} + \lambda_1 F_1)^{-1} \hat{y}_1 \\ \vdots \\ (I_{T_M} + \lambda_M F_M)^{-1} \hat{y}_M \end{bmatrix}. \quad (12)$$

Note that equation (11) is no longer an oblique projection, so that applying it twice may result in a different time series than applying it once. Intuitively, the HP filter (5) is not a projection, and thus, equation (11) which inherits the property of the HP filter is also not a projection.

2.3. Choice of Smoothness Parameters

The choice of smoothness parameters needs to reflect the units of the covariance matrix of forecast errors W . Section 2.2 demonstrates that problem (9) reduces to the HP filter when the weight is an identity matrix $W = I_N$. The covariance matrix of forecast errors W , however, is rarely an identity, as fan charts often expand over the forecast horizon. Thus, the typical smoothness parameter values cannot be applied directly.

One way to align the units of the smoothness parameters is to use the minimum variance of the forecast errors

$$\lambda_m = \frac{\lambda_m^*}{\sigma_m^2}, \quad \sigma_m^2 = \min_{t=1, \dots, T_m} V(\hat{y}_{mt} - y_{mt}^*) = \min_{t=1, \dots, T_m} (W_m)_{tt}, \quad m = 1, \dots, M, \quad (13)$$

where λ_m^* is the typical smoothness parameter for the frequency of the m -th time series, like $\lambda_m^* = 1600$ if the m -th time series is quarterly, and W_m is a submatrix of W that corresponds to the m -th time series. This is a natural choice since problem (9) boils down to the HP filter with the typical smoothness parameter in the special case where W is diagonal, W_m has constant diagonals, and the constraints are zeros $(C, d) = (0, 0)$.

$$(y - \hat{y})' W^{-1} (y - \hat{y}) + y' \Phi y = \sum_{m=1}^M \frac{1}{(W_m)_{11}} \left\{ \underbrace{(y_m - \hat{y}_m)' (y_m - \hat{y}_m) + \lambda_m^* y_m' F_m y_m}_{HP \text{ filter}} \right\}. \quad (14)$$

Intuitively, the special case corresponds to the situation where all forecast errors are uncorrelated, and the variances of forecast errors are constant, so the forecast errors come from white noise. This special case is more general than an identity weight $W = I_N$ since each submatrix W_m can take different values. Each submatrix W_m having constant diagonals, however, is restrictive since the variance of forecast errors typically increases as the forecast horizon extends, like a fan chart.

There can be many alternative choices for the smoothness parameters. For example, replacing the minimum variance of the forecast errors (13) with the median, average, or maximum variance of the forecast errors leads to the same expression (14) in the special case. One advantage of using the minimum variance of the forecast errors (13) is that, when the fan chart expands over time, it puts stronger smoothing force on the farther forecast horizon, so the time series is smoother in the more uncertain horizon. To see this, suppose the weight matrix is diagonal, and the diagonals increase monotonically as the index increases, $(W_m)_{11} < \dots < (W_m)_{T_m T_m}$. The part of the objective function that corresponds to the m -th time series becomes

$$\sum_{t=1}^{T_m} \frac{1}{(W_m)_{tt}} (y_{mt} - \hat{y}_{mt})^2 + \frac{\lambda_m^*}{(W_m)_{11}} \sum_{t=2}^{T_m-1} \{(y_{mt+1} - y_{mt}) - (y_{mt} - y_{mt-1})\}^2. \quad (15)$$

The weight on the first term becomes smaller as the time index increases, but the weight on the second term remains the same over time. Thus, the force to avoid deviation from \hat{y} becomes relatively weaker, and the resulting time series is smoother on the horizons where forecasters have relatively weaker confidence.

One can also empirically estimate the smoothness parameters $\{\lambda_m\}_{m=1}^M$, although optimizing M parameters can become computationally challenging. One caveat is that, in estimating the parameters, targeting the smoothness of historical data may not be optimal, as can be seen in Section 3.2.2. For example, if the data-generating process is a random walk, the best forecast is a constant over time, in which case the volatility of the historical data should not be reflected in the forecast. Another caveat is that, even if the estimated smoothness parameters minimize the variance of the forecast errors, the forecast may not be visually smooth enough. This might be unsatisfactory since forecasters often want to explain the economic story to the public. For communication with the public, it is not sufficient to say that the kink is a result of the algorithm, so they want the path to be visually smooth unless the kink is intentionally introduced with associated story to be told.

In this paper, we use the minimum variance of the forecast errors (13) as the baseline smoothness parameter. In practice, forecasters can choose their favorite values depending on their objectives.

3. Applications


This section illustrates the proposed method using as example US GDP forecasts. It shows that the proposed method can replicate the smoothness of the expert-generated forecasts, achieve a slightly smaller forecast error than the reconciled-but-not-smoothed alternative generated by min-T reconciliation, and incorporate various constraints.

3.1. Data

The data used in the following three examples are annual and quarterly US real GDP and the subcomponents of the annual GDP, comprising of annual foreign balance, annual total domestic demand, and annual statistical discrepancy, taken from the 2023 October vintage of the World Economic Outlook (WEO). The WEO database contains both historical data and forecasts, ranging from 1950 to 2028 for annual data and 1979Q1 to 2024Q4 for quarterly data (Figure 1). As of October 2023, the historical data for US GDP have been released up to 2023Q2. In this paper, we assume that data up to 2023Q2 is historical and from 2023Q3 is forecast, abstracting from uncertainties regarding whether some quarters contain flash estimates or will be revised later. The quarterly data are seasonally adjusted and annualized so that the average quarterly GDP over the four quarters equals the annual GDP in each year.

The annual GDP and seasonally adjusted quarterly GDP forecast in WEO are usually smooth unless the country team has prior information to believe there should be a kink or lacks resources to finetune the forecasts. Different economists may use different forecasting techniques, including their judgment, but they typically strive to smooth the forecast as they are published and discussed in various public documents (e.g., country reports, multilateral surveillance products and the like). In addition, forecasts made by individual country teams are checked by the IMF's research department, which validates that accounting identities hold, and no unintended kinks remain, among other quality controls before the forecasts are included in the WEO database. For further details of the IMF's WEO production process, see Genberg et al. (2014).

Figure 1. Times Series of US GDP in WEO



Source: IMF October 2023 World Economic Outlook database.

Note: The left chart shows the level of US GDP in real USD, and the right chart shows the growth of US real GDP used in the applications. Quarterly GDP is annualized and seasonally adjusted.

3.2. Forecasting Quarterly GDP

As a simple example, this section illustrates how the proposed method helps forecast quarterly GDP in a smooth manner when quarterly GDP needs to aggregate up to the exogenously given annual GDP. This example has been studied in the context of national accounts as the benchmarking problem, and the comparison is presented in Annex I. We show that the proposed method can mimic the smoothness of the WEO forecast.

3.2.1. Input

The forecast horizon of interest is from 2023Q3 to 2024Q4. To mitigate the end point problem of the HP filter, the forecast horizon is extended to 2023Q2 and to 2025Q4. Including four quarters after 2024Q4 is excessive but can utilize the annual GDP forecast for 2025. Thus, the objective is to forecast quarterly GDP from 2023Q2 to 2025Q4. Table 1 summarizes the data structure.

Table 1. Structure of Data and Forecast Horizon

	History			Forecast					
Quarterly GDP	...	y_{2023Q1}^{WEO}	\tilde{y}_{2023Q2}	\tilde{y}_{2023Q3}	\tilde{y}_{2023Q4}	...	\tilde{y}_{2024Q4}	...	\tilde{y}_{2025Q4}
Annual GDP (given)	...	y_{2023}^{WEO}			y_{2024}^{WEO}			y_{2025}^{WEO}	
Horizon of interest									
Forecasted				Expand forecast horizon to mitigate HP filter's end point problem					

Note: The table shows the structure of the data and shades the forecast horizon used in the application.

The first-step forecast \hat{y} uses naïve forecast based on the growth rate.

$$\hat{y}_{2023Q2} = y_{2023Q2}^{WEO}, \quad \hat{y}_t - \frac{y_{2023Q2}^{WEO}}{y_{2023Q1}^{WEO}} \hat{y}_{t-1} = 0, \quad t = 2023Q3, \dots, 2025Q4. \quad (16)$$

There can be many alternative choices for the first-step forecast \hat{y} . The naïve forecast has the pedagogical benefit of generating a constant forecast of the GDP growth so that the kinks introduced by reconciliation can be seen clearly. For robustness check, Annex II uses alternative forecasting methods, such as autoregression with four lags and exponential smoothing.

To estimate the covariance matrix of forecast errors W , the historical data is split into $n = 10$ expanding windows. (Table 2) The forecast errors in the ten test sets are used to estimate the covariance matrix W .

Table 2. Expanding Windows of Time Series Split

	Time dimension of historical data		
Fold 1	Training	Test 2011Q2-2013Q4	
Fold 2	Training	Test 2012Q2-2014Q4	
Fold 3	Training	Test 2013Q2-2015Q4	
⋮		⋮	
Fold 10	Training		Test 2020Q2-2022Q4

Note: The table shows the training sets and test sets used to generate the forecast errors.

We use oracle shrinkage approximating estimator with diagonal target (OASD) by Ando and Xiao (2023) instead of the sample covariance matrix of forecast errors $\hat{\Sigma}$.

$$\hat{W} = \hat{\rho} \hat{\Sigma} + (1 - \hat{\rho}) \text{diag}(\hat{\Sigma}), \quad \hat{\rho} = \min \left\{ \frac{1}{n \hat{\phi}}, 1 \right\}, \quad \hat{\phi} = \frac{\text{tr}(\hat{\Sigma}^2) - \text{tr}(\text{diag}(\hat{\Sigma})^2)}{\text{tr}(\hat{\Sigma}^2) + \text{tr}(\hat{\Sigma})^2 - 2 \text{tr}(\text{diag}(\hat{\Sigma})^2)}. \quad (17)$$

The OASD shrinks the sample covariance matrix $\hat{\Sigma}$ toward the diagonals of the sample covariance matrix so that the shrunk matrix \hat{W} is invertible even when the sample covariance matrix is degenerate, which can happen when the number of folds is smaller than the length of the first-step forecast \hat{y} , as is the case here, or when both the first-step forecast \hat{y} and historical data y^{WEO} from WEO satisfy the same linear constraints. One can alternatively use other shrinkage methods discussed in Ando and Xiao (2023), such as Ledoit and Wolf (2004), Schafer and Strimmer (2005), and Chen et al. (2010).

The constraints that the second-step forecast \tilde{y} satisfies are

$$\begin{aligned} \tilde{y}_{2023Q2} &= y_{2023Q2}^{WEO}, \\ \tilde{y}_{2023Q3} + \tilde{y}_{2023Q4} &= 4y_{2023}^{WEO} - y_{2023Q1}^{WEO} - y_{2023Q2}^{WEO}, \\ \tilde{y}_{tQ1} + \tilde{y}_{tQ2} + \tilde{y}_{tQ3} + \tilde{y}_{tQ4} &= 4y_t^{WEO}, \quad t = 2024, 2025, \end{aligned} \quad (18)$$

where the first constraint fixes the forecast of 2023Q2 to be the historical data, the second constraint requires the forecast of 2023Q3 and Q4 to be consistent with the annual forecast of 2023 and the historical data of the first half of 2023, and the third constraint ensures that the average quarterly GDP equals the annual GDP. In matrix form, the constraints can be written, using a 4×11 matrix C and 4×1 vector d , as

$$C\tilde{y} = d, \quad C = \begin{bmatrix} 1 & & & & & & & & & & & & \\ & \mathbf{1}_{1 \times 2} & & & & & & & & & & & \\ & & \mathbf{1}_{1 \times 4} & & & & & & & & & & \\ & & & \mathbf{1}_{1 \times 4} & & & & & & & & & \\ & & & & & & & & & & & & \end{bmatrix}, \quad \tilde{y} = \begin{bmatrix} \tilde{y}_{2023Q2} \\ \tilde{y}_{2023Q3} \\ \tilde{y}_{2023Q4} \\ \vdots \\ \tilde{y}_{2025Q4} \end{bmatrix}, \quad d = \begin{bmatrix} y_{2023Q2}^{WEO} \\ 4y_{2023}^{WEO} - y_{2023Q1}^{WEO} - y_{2023Q2}^{WEO} \\ 4y_{2024}^{WEO} \\ 4y_{2025}^{WEO} \end{bmatrix}. \quad (19)$$

Finally, the smoothness matrix Φ is a penta-diagonal matrix analogous to (4), multiplied by the typical smoothness parameter for quarterly data 1600 and normalized by the minimum variance of the forecast errors.

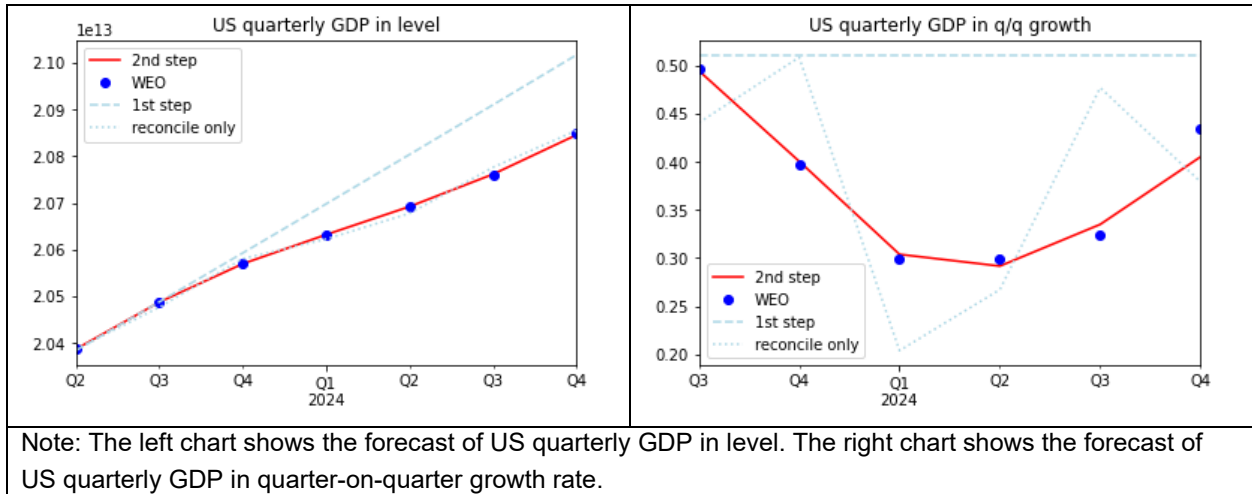
$$\Phi = \frac{1600}{\hat{\sigma}^2} F, \quad F: 11 \times 11, \quad \hat{\sigma}^2 = \min_i \hat{W}_{ii}. \quad (20)$$

Given the first-step forecast and its associated covariance matrix of forecast errors (\hat{y}, \hat{W}) , constraints (C, d) , and the smoothness matrix Φ , the second-step forecast \tilde{y} can be derived by (11).

3.2.2. Output

The second-step forecast \tilde{y} traces well the smoothness of the WEO forecast y^{WEO} . Figure 2 shows that the second-step forecast \tilde{y} traces the WEO GDP level y^{WEO} and the quarter-on-quarter growth with high precision, much better than the first-step forecast \hat{y} and the reconciled-but-not-smoothed alternative where the smoothness parameter is set to $\lambda = 0$. The first-step forecast \hat{y} is smooth but does not aggregate up to the annual GDP. When only min-T reconciliation is applied without smoothing, $\lambda = 0$, the forecast includes unintended boom and bust in 2024Q3 and Q4. The time series remains volatile even if the weight is replaced by an identity I_N . This is an inconvenient property of forecast reconciliation without smoothing since forecasters need to smooth the path manually to highlight the intended kinks and suppress the unintended ones.


Figure 2. Forecast of US quarterly GDP



Back testing shows that smooth reconciliation can slightly improve forecast performance, although the difference is not statistically significant. Figure 3 compares the root mean squared percentage error (RMSPE) of the 6-quarter forecast horizon, averaged over the past 25 years from 1996 to 2020. Expectedly, the RMSPE of the first-step forecast is the largest since it does not reflect the actual annual GDP. An interesting result is that the second-step forecast \tilde{y} has a lower RMSPE than the reconciled-but-not-smoothed forecast. Since the formula of the min-T reconciliation is derived by optimally minimizing the variance of forecast errors, one would expect that imposing smoothness could lead to a suboptimal forecast. It turns out that the Diebold-Mariano test does not reject the null hypothesis that they have the same accuracy, and thus, the reconciled forecast can be smoothed without sacrificing the forecast performance.

Figure 3 also shows the RMSPE for the second-step forecast with random noise, which is drawn from the empirical distribution of the HP filter's cyclical component of the historical data. This confirms that simply adding noise to the second-step forecast \tilde{y} to mimic the volatility of the historical data does not improve forecast accuracy. Annex II shows that the qualitative result remains similar even if the naïve forecaster in the first-step forecast is replaced with other widely used forecasting methods, such as autoregression and exponential smoothing.

Figure 3. Forecast Performance Comparison

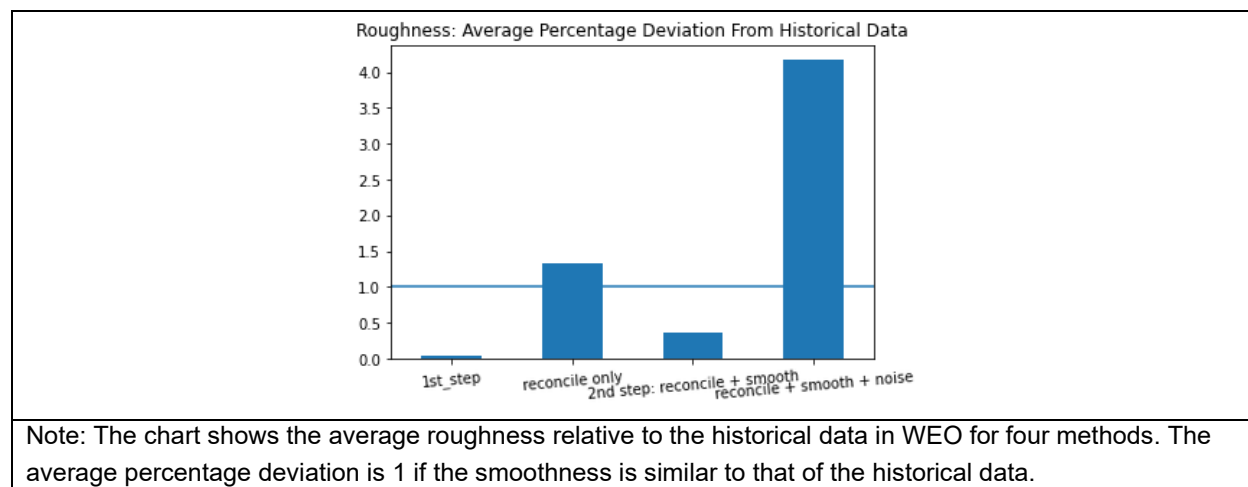


Note that the smoothness of the proposed method may not be similar to that of the historical data. Figure 4 applies the same back testing to the average roughness relative to the historical data over the past 25 years, where the roughness of a times series $y = \{y_t\}_{t=1}^T$ is defined by the penalty term of the HP filter

$$\text{roughness}(y) = \sum_{t=2}^{T-1} \{(y_{t+1} - y_t) - (y_t - y_{t-1})\}^2, \quad (21)$$

and the average is taken for the ratio of a forecasting method's roughness and historical data's roughness. As expected, the proposed method is smoother than the historical data. In many forecasting models, such as random walk and autoregressions, the optimal forecast is smoother than realized times series. Thus, whether a forecasting method has a good forecast performance should be understood separately from whether it generates a smoothness similar to the true data.

Figure 4. Roughness Comparison



3.3. Mixed-Frequency: Forecasting Both Annual and Quarterly GDP

Section 3.2 illustrates the forecast of quarterly GDP. This section extends the example to the forecast of both the annual and quarterly GDP, illustrating how the proposed method handles mixed-frequency data. Although the extension to multiple time series is straightforward, the second-step forecast no longer tracks the WEO forecast. It is shown, however, that the proposed method can incorporate ad-hoc information into the constraints, and the information on the short-run and long-run growth can bring back the second-step forecast closer to the WEO forecast.

3.3.1. Input

The forecast horizon of interest is from 2023 to 2028 for the annual GDP and 2023Q3 to 2028Q4 for the quarterly GDP. The forecast horizon for the quarterly GDP is longer than that in section 3.2 since the annual GDP forecast is available until 2028 in WEO. To mitigate the end point problem of the HP filter, the forecast horizon is extended to 2022-2029 and 2023Q2 to 2029Q4.

The first-step forecast \hat{y} remains to be the naïve forecast based on the growth rate for both annual and quarterly GDP, which takes the growth rate of the last observation. The first-step forecast \hat{y} is a 35×1 vector, stacking annual and quarterly forecasts

$$\hat{y} = \begin{bmatrix} \hat{y}_{2022} \\ \vdots \\ \hat{y}_{2029} \\ \hat{y}_{2023Q2} \\ \vdots \\ \hat{y}_{2029Q4} \end{bmatrix}. \quad (22)$$

As in section 3.2, the estimation method of the 35×35 covariance matrix of forecast errors \hat{W} is OASD (17).

The constraints that the second-step forecast \hat{y} satisfies are

$$\begin{aligned}
\tilde{y}_t &= y_t^{WEO}, & t &= 2022, 2023Q2 \\
4\tilde{y}_{2023} - \tilde{y}_{2023Q3} - \tilde{y}_{2023Q4} &= y_{2023Q1}^{WEO} + y_{2023Q2}^{WEO}, \\
4\tilde{y}_t - \tilde{y}_{tQ1} - \tilde{y}_{tQ2} - \tilde{y}_{tQ3} - \tilde{y}_{tQ4} &= 0, & t &= 2024, \dots, 2029.
\end{aligned} \tag{23}$$

Note that the constraint for 2023 no longer uses the annual GDP for 2023 in the WEO forecast, y_{2023}^{WEO} . In matrix form, the constraints can be written using a 9×35 matrix C and 9×1 vector d .

Finally, the smoothness matrix Φ is a block-diagonal matrix.

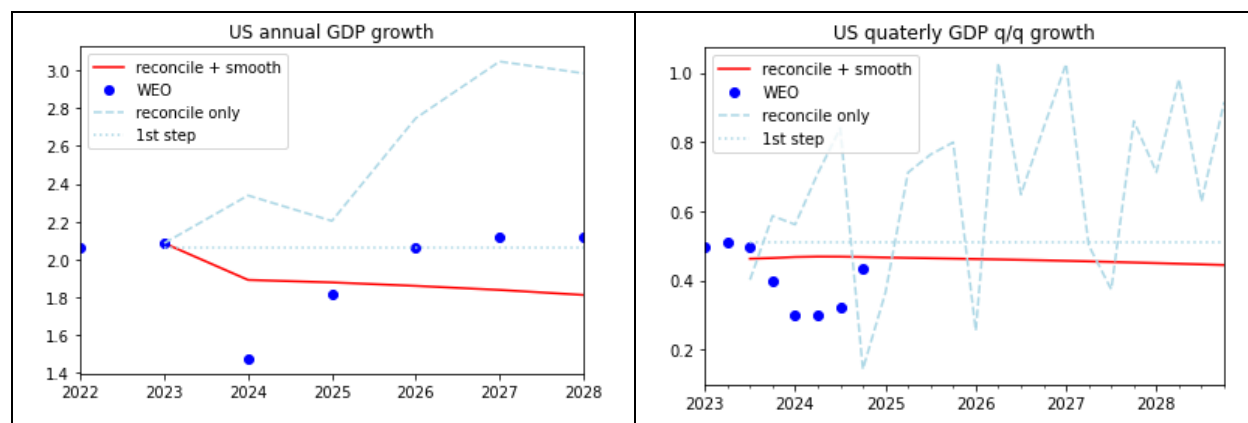
$$\Phi = \begin{bmatrix} \frac{100}{\hat{\sigma}_A^2} F_A & \\ & \frac{1600}{\hat{\sigma}_Q^2} F_Q \end{bmatrix}, \quad \hat{\sigma}_A^2 = \min_{i=1, \dots, 8} \hat{W}_{ii}, \quad F_A: 8 \times 8, \quad \hat{\sigma}_Q^2 = \min_{i=9, \dots, 35} \hat{W}_{ii}, \quad F_Q: 27 \times 27. \tag{24}$$

The smoothness parameter for the annual GDP is the typical value 100 normalized by the minimum variance of the forecast errors of annual GDP's first-step forecast. The smoothness parameter for the quarterly GDP remains to be the typical value of 1600 normalized by the corresponding minimum variance of the forecast errors. The penta-diagonal matrices for the annual GDP F_A and quarterly GDP F_Q are analogous to (4) with different sizes. Given the updated first-step forecast and its associated covariance matrix of forecast errors (\hat{y}, \hat{W}) , constraints (C, d) , and the smoothness matrix Φ , the second-step forecast \tilde{y} can be derived by (11).

3.3.2. Output

Figure 5 shows that the second-step forecast \tilde{y} by the proposed method is smooth and satisfies accounting identities, but the time series do not trace the WEO forecasts. The reconciled-but-not-smoothed time series are volatile and difficult to explain with an intuitive economic story. The second-step forecast \tilde{y} is smoother, but since it only uses historical data for forecasting, it can deviate substantially from the WEO forecast, which may include future information, such as the recession and recovery in 2024 and 2025.

Figure 5. Forecast of US Annual and Quarterly GDP



Note: The left chart shows the US annual GDP growth. The right chart shows the US quarterly GDP quarter-on-quarter growth rate. No constraints other than accounting identities are imposed.

3.3.3. Incorporating Ad-Hoc Information

Incorporating ad-hoc future information is often useful. For example, if the forecasters know some budgets cut in near future or has some targeted potential growth rate in the long run, they want to incorporate such information in the forecast and connect other periods smoothly so that they can explain the economic story.

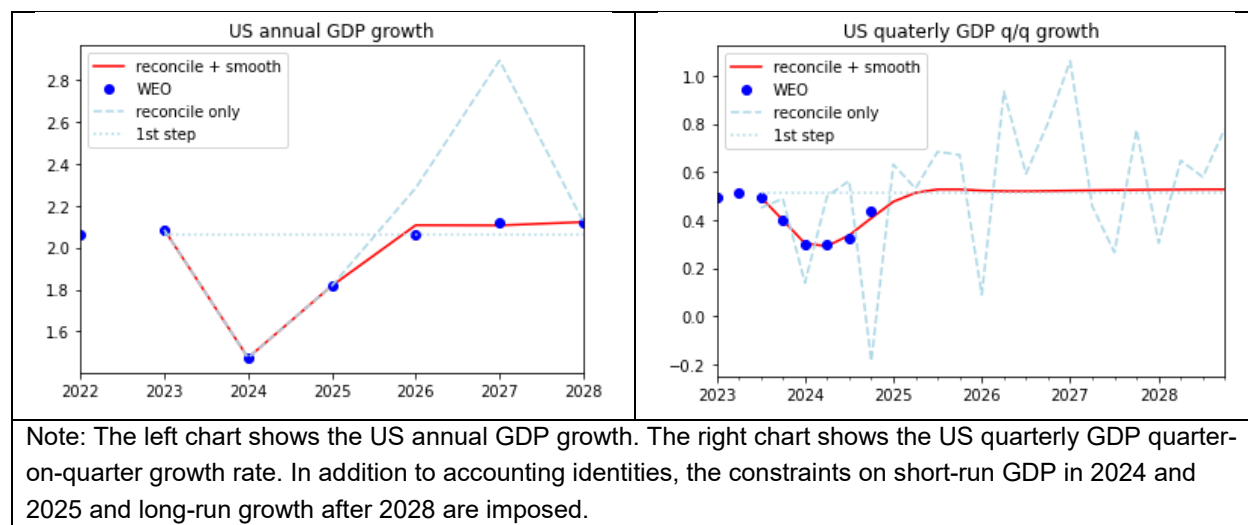
Such ad-hoc future information can be incorporated in the constraints. In the WEO forecast of Figure 5, a rapid slowdown is predicted from 2023 to 2024 and a sharp recovery is predicted from 2024 to 2025. In addition to such a short-run development, the long-run growth rate converges to slightly above two percent. The information can be incorporated by imposing

$$\begin{aligned} \tilde{y}_t &= y_t^{WEO}, & t = 2023, 2024, 2025, \\ \tilde{y}_t - \frac{y_{2028}^{WEO}}{y_{2027}^{WEO}} \tilde{y}_{t-1} &= 0, & t = 2028, 2029. \end{aligned} \quad (25)$$

in addition to the constraints (23) in section 3.3.1.

Figure 6 shows that the information (25) is sufficient to bring back the second-step forecast close to the WEO forecast. The reconciled-but-not-smoothed forecast coincides with the second-step forecast \tilde{y} and WEO forecast y^{WEO} up to 2025 due to the constraints, but it is more volatile than alternatives in other years. The quarterly GDP path also tracks the WEO forecast well, suggesting that there is not much quarter-specific economic story in the quarterly GDP forecast once the annual GDP is conditioned.

Figure 6. Forecast of US Annual and Quarterly GDP with Ad-Hoc Constraints



3.4. Cross-Sectional and Temporal Constraints: Adding Subcomponents

This section accommodates both cross-sectional (intra-temporal) and temporal (inter-temporal) constraints by adding subcomponents of annual GDP. The example with multivariate mixed-frequency time series demonstrates that the proposed method can produce sensible forecasts even in a higher dimensional environment with various types of constraints.

3.4.1. Input

In addition to the annual and quarterly GDP, this example includes annual foreign balance FB_t , annual domestic demand DD_t , and annual statistical discrepancy SD_t . In the expenditure approach of the GDP, the foreign balance is often decomposed into export minus import, and the domestic demand into consumption and investment. Further decomposition into lower subcomponents, such as private vs public, is available. The decompositions can also be made at quarterly frequency. Although exploring the most detailed decomposition at all possible frequencies is straightforward, this paper keeps the decomposition at high level and at annual frequency for expositional purpose.

The forecast horizon of interest is from 2023 to 2028 for the annual variables and 2023Q3 to 2028Q4 for the quarterly GDP. To mitigate the end point problem of the HP filter, the horizon is extended to 2022-2029 and 2023Q2 to 2029Q4 in conducting forecast.

The first-step forecast \hat{y} remains to be the naïve forecast based on growth rate for the quarterly GDP, but for annual variables, the growth rate is replaced by the difference to accommodate foreign balance and statistical discrepancy that can take both positive and negative values. The first-step forecast \hat{y} is a 59×1 vector, stacking annual and quarterly forecasts.

$$\hat{y} = \begin{bmatrix} \widehat{GDP}_{2022} \\ \vdots \\ \widehat{GDP}_{2029} \\ \widehat{FB}_{2022} \\ \vdots \\ \widehat{FB}_{2029} \\ \widehat{DD}_{2022} \\ \vdots \\ \widehat{DD}_{2029} \\ \widehat{SD}_{2022} \\ \vdots \\ \widehat{SD}_{2029} \\ \widehat{GDP}_{2023Q2} \\ \vdots \\ \widehat{GDP}_{2029Q4} \end{bmatrix}. \quad (26)$$

As in section 3.2, the estimation method of the 59×59 covariance matrix of forecast errors \widehat{W} is OASD (17).

We impose both accounting identities and ad-hoc information on the second-step forecast \hat{y} . The accounting identities are analogous to (23) and include the GDP expenditure approach identity.

$$\begin{aligned} FB_{2022} &= FB_{2022}^{WEO}, \\ DD_{2022} &= DD_{2022}^{WEO}, \\ GDP_t &= GDP_t^{WEO}, \quad t = 2022, 2023Q2, \\ 4GDP_{2023} - GDP_{2023Q3} - GDP_{2023Q4} &= GDP_{2023Q1}^{WEO} + GDP_{2023Q2}^{WEO}, \\ 4GDP_t - GDP_{tQ1} - GDP_{tQ2} - GDP_{tQ3} - GDP_{tQ4} &= 0, \quad t = 2024, \dots, 2029, \\ GDP_t - FB_t - DD_t - SD_t &= 0, \quad t = 2022, \dots, 2029. \end{aligned} \quad (27)$$

Ad-hoc information is analogous to (25) and includes

$$\begin{aligned}
 GDP_t &= GDP_t^{WEO}, & t = 2023, 2024, 2025, \\
 GDP_t - \frac{GDP_{2028}^{WEO}}{GDP_{2027}^{WEO}} GDP_{t-1} &= 0, & t = 2028, 2029, \\
 FB_t - \frac{FB_{2028}^{WEO}}{GDP_{2028}^{WEO}} GDP_t &= 0, & t = 2028, 2029, \\
 DD_t - \frac{DD_{2028}^{WEO}}{GDP_{2028}^{WEO}} GDP_t &= 0, & t = 2028, 2029.
 \end{aligned} \tag{28}$$

As in section 3.3.3, the constraints on the medium-term GDP reflect the ad-hoc information on recession and recovery in the WEO forecast. The constraints for 2028 and 2029 reflect the long-run information that forecasters may have from various long-run analyses, such as the modellings of potential growth and external balance. Note that we do not require the foreign balance and domestic demand to coincide with the WEO forecast from 2023 to 2025, unlike the annual GDP. This is to see how subcomponents track the WEO forecast when the annual GDP exhibits a recession. In matrix form, the constraints can be written using a 28×59 matrix C and 28×1 vector d .

Finally, the smoothness matrix Φ is a block-diagonal matrix.

$$\Phi = \begin{bmatrix} \frac{100}{\hat{\sigma}_A^2} F_A & & & & \\ & \frac{100}{\hat{\sigma}_{FB}^2} F_A & & & \\ & & \frac{100}{\hat{\sigma}_{DD}^2} F_A & & \\ & & & \frac{100}{\hat{\sigma}_{SD}^2} F_A & \\ & & & & \frac{1600}{\hat{\sigma}_Q^2} F_Q \end{bmatrix}, \tag{29}$$

where (F_A, F_Q) are the same penta-diagonal matrices as (24), and the variances are estimated by

$$\hat{\sigma}_A^2 = \min_{i=1, \dots, 8} \widehat{W}_{ii}, \quad \hat{\sigma}_{FB}^2 = \min_{i=9, \dots, 16} \widehat{W}_{ii}, \quad \hat{\sigma}_{DD}^2 = \min_{i=17, \dots, 24} \widehat{W}_{ii}, \quad \hat{\sigma}_{SD}^2 = \min_{i=25, \dots, 32} \widehat{W}_{ii}, \quad \hat{\sigma}_Q^2 = \min_{i=33, \dots, 59} \widehat{W}_{ii}. \tag{30}$$

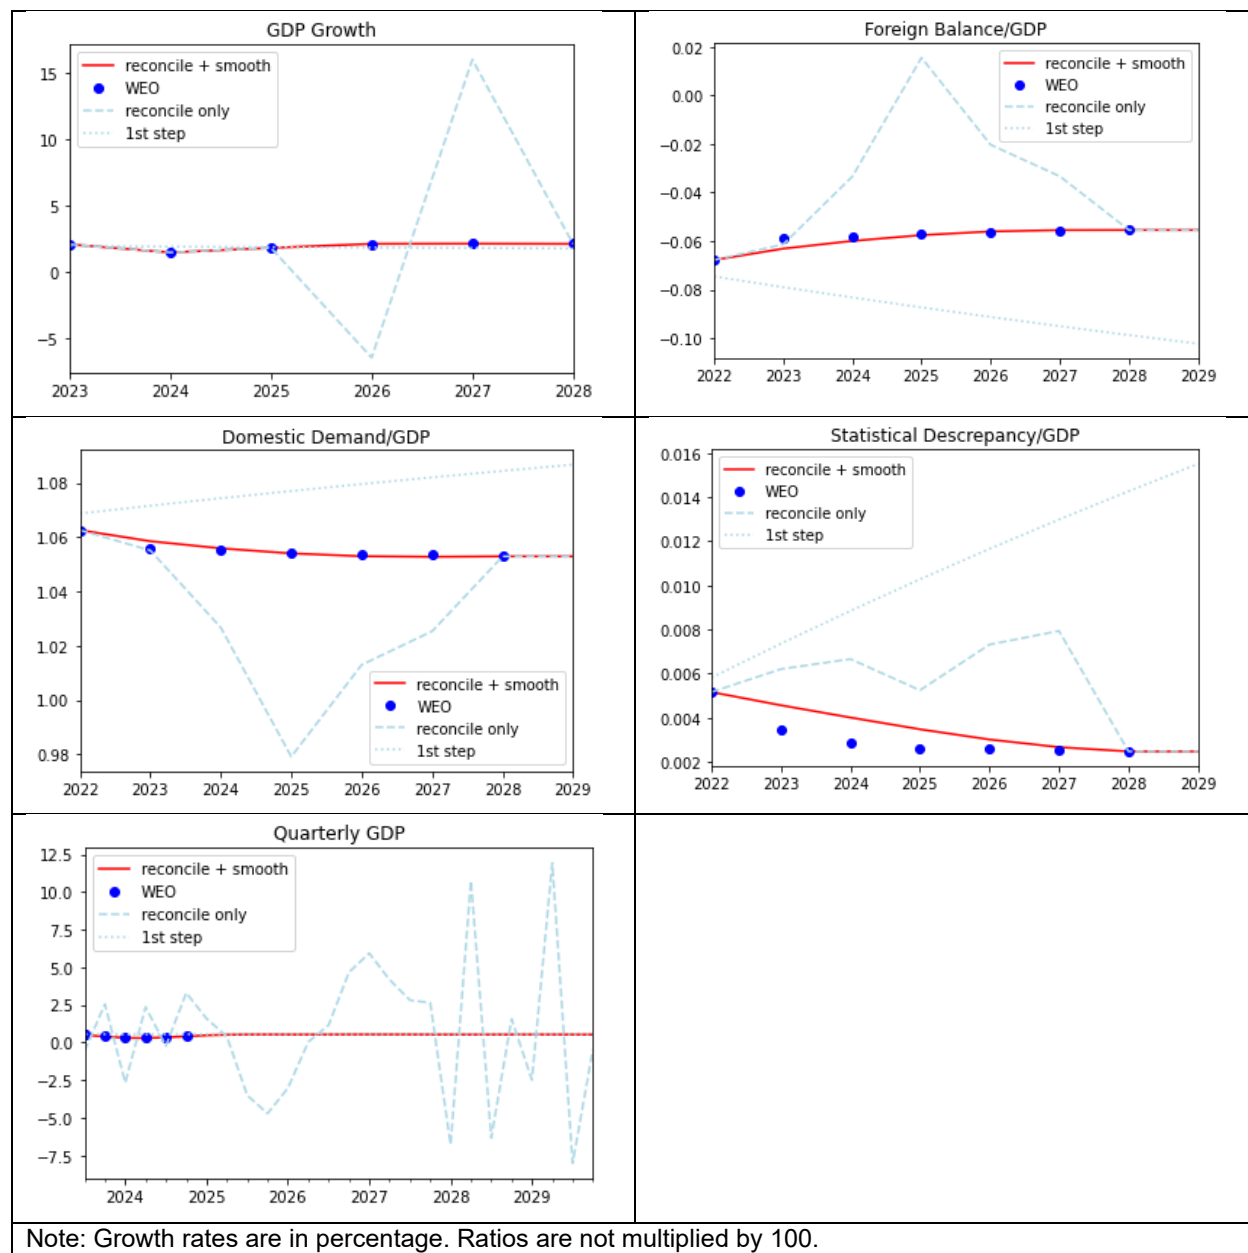
Given the first-step forecast and its associated covariance matrix of forecast errors (\hat{y}, \widehat{W}) , constraints (C, d) , and the smoothness matrix Φ , the second-step forecast \tilde{y} can be derived by (11).

3.4.2. Output

Figure 7 shows that the second-step forecast \tilde{y} by the proposed method is smooth and traces the WEO forecast better than the first-step and reconciled-but-not-smoothed forecasts. The reconciled-but-not-smoothed forecast is more volatile than in section 3.3 and exhibits kinks that are not easy to explain. The kinks are so large that the difference between the WEO forecast and the proposed method is almost invisible except for the

statistical discrepancy. The foreign balance in GDP and domestic demand in GDP trace the WEO forecast satisfactorily with a smoother path, even though no ad-hoc information except for the initial and terminal values is imposed.

Figure 7. Forecast of Mixed-Frequency GDP and Subcomponents



4. Conclusion

This paper proposes a smooth reconciliation method that allows the forecast to satisfy constraints and the time series to be smooth. The method can be used for multivariate forecasts of mixed-frequency time series. Its

closed-form solution is amenable to a high-dimensional environment. The method is flexible enough to incorporate accounting identities, within or across time, and other ad-hoc information that the forecasters may want to incorporate. Imposing smoothness does not incur loss in forecast performance and improves practical usefulness by highlighting the kinks that forecasters intentionally introduce and suppressing unintended noises.

Annex I. Comparison with Modified Denton

The problem of adjusting the quarterly path given the annual path, called benchmarking, has been studied in the context of national accounts statistics. A modified version of Denton (1971) by Cholette (1984) can be applied to the problem as follows.

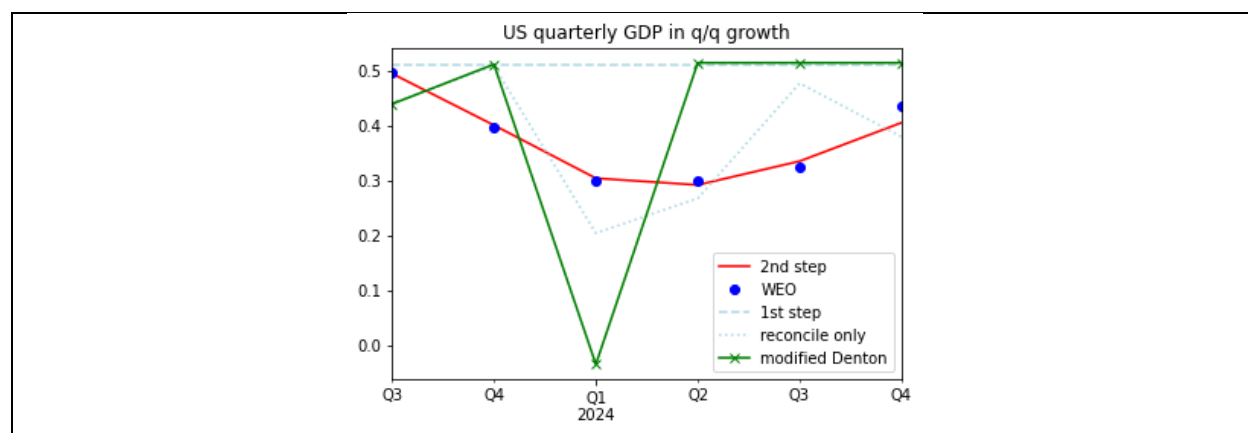
$$\begin{aligned} \{\tilde{y}_t\}_{t=2023Q3}^{2024Q4} &= \underset{\{y_t\}_{t=2023Q3}^{2024Q4}}{\operatorname{argmin}} \sum_{t=2023Q4}^{2024Q4} \left(\frac{y_t - \hat{y}_t}{\hat{y}_t} - \frac{y_{t-1} - \hat{y}_{t-1}}{\hat{y}_{t-1}} \right)^2 \\ \text{s. t. } &\begin{cases} y_{2023Q3} + y_{2023Q4} = 4y_{2023}^{WEO} - y_{2023Q1}^{WEO} - y_{2023Q2}^{WEO} \\ y_{2024Q1} + y_{2024Q2} + y_{2024Q3} + y_{2024Q4} = 4y_{2024}^{WEO} \end{cases} \end{aligned} \quad (31)$$

where \hat{y}_t and y_t^{WEO} denote the same first-step forecast and WEO data as in section 3.2. The problem can be written as a quadratic form in $\{y_t\}_{t=2023Q3}^{2024Q4}$ and interpreted as another projection problem, similar to (1), although it doesn't have a mechanism to ensure smoothness. Since the HP filter is not involved, the time series are not extended beyond the periods of interest and denoted by s representing "short."

$$\min_s (s - \hat{s})' \Omega (s - \hat{s}) \quad \text{s. t. } \begin{cases} As = b, & A = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}, & b = \begin{bmatrix} 4y_{2023}^{WEO} - y_{2023Q1}^{WEO} - y_{2023Q2}^{WEO} \\ 4y_{2024}^{WEO} \end{bmatrix} \\ \Omega = LR, & R = L' = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} 1 \\ \hat{y}_{2023Q3} \\ \vdots \\ 1 \\ \hat{y}_{2024Q4} \end{bmatrix} \end{cases} \quad (32)$$

Although the weight matrix Ω in the modified Denton method is degenerate, the problem can be solved numerically. Figure 8 overlays the modified Denton over Figure 2. The time series is less smooth, like reconciliation-only, as expected since it does not have a mechanism to ensure smoothness.

Figure 8. Comparison with Modified Denton



Annex II. Alternative First-Step Methods

Figure 9 shows the back testing by replacing the naïve forecaster by autoregression with four lags. Although the quantitative magnitudes differ, the qualitative results remain the same.

Figure 9. Forecast Performance Comparison When the First Step is AR(4)





Figure 10 shows the back testing by replacing the naïve forecaster by exponential smoothing with automatic parameter selection, which is available in many statistical software such as the sktime package by Loning et al. (2019). Although the quantitative magnitudes differ, the qualitative results remain the same.

Figure 10. Forecast Performance Comparison When the First Step is AutoETS



References

- Ando, Sakai and Taehoon Kim (2023), "Systematizing Macroframework Forecasting: High-Dimensional Conditional Forecasting with Accounting Identities," IMF Economic Review.
- Ando, Sakai and Futoshi Narita (2022), "An Alternative Proof of Minimum Trace Reconciliation," IMF Working Paper.
- Ando, Sakai and Mingmei Xiao (2023), "High-Dimensional Covariance Matrix Estimation: Shrinkage Toward a Diagonal Target," IMF Working Paper.
- Athanasopoulos, George, Rob Hyndman, Nikolaos Kourentzes, and Fotios Petropoulos (2017), "Forecasting with Temporal Hierarchies," European Journal of Operational Research.
- Chen, Yilun, Ami Wiesel, Yonina Eldar, Alfred Hero (2010), "Shrinkage Algorithms for MMSE Covariance Estimation," IEEE Transactions on Signal Processing.
- Cholette, Pierre (1984), "Adjusting Sub-Annual Series to Yearly Benchmarks," Survey Methodology.
- Denton, Frank (1971), "Adjustment of Monthly or Quarterly Series to Annual Totals: An Approach Based on Quadratic Minimization," Journal of the American Statistical Association.
- Di Fonzo, Tommaso and Daniele Girolimetto (2023), "Cross-Temporal Forecast Reconciliation: Optimal Combination Method and Heuristic Alternatives," International Journal of Forecasting.
- Di Fonzo, Tommaso and Marco Marini (2011), "Simultaneous and Two-Step Reconciliation of Systems of Time Series: Methodological and Practical Issues," Journal of the Royal Statistical Society.
- Dritsaki, Melina and Chaido Dritsaki (2022), "Comparison of HP Filter and the Hamilton's Regression," Mathematics.
- Genberg, Hans, Andrew Martinez, and Michael Salemi (2014), "The IMF/WEO Forecast Process," Independent Evaluation Office Background paper.
- Girolimetto, Daniele and Tommaso Di Fonzo (2023), "Point and Probabilistic Forecast Reconciliation for General Linearly Constrained Multiple Time Series," Statistical Methods & Applications.
- Hamilton, James (2018), "Why You Should Never Use the Hodrick-Prescott Filter," The Review of Economics and Statistics.
- Hodrick, Robert (2020), "An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data," NBER Working Papers.
- Hodrick, Robert and Edward Prescott (1997), "Postwar U.S. Business Cycles: An Empirical Investigation," Journal of Money, Credit and Banking.
- Ledoit, Olivier and Michael Wolf (2004), "A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices," Journal of Multivariate Analysis.
- Loning, Markus, Anthony Bagnall, Sajaysurya Ganesh, Viktor Kazakov, Jason Lines, Franz Kiraly (2019), "Sktime: A Unified Interface for Machine Learning with Time Series," arXiv.
- Panagiotelis, Anastasios, George Athanasopoulos, Puwasala Gamakumara, Rob Hyndman (2023), "Forecast Reconciliation: A Geometric View with New Insights on Bias Correction," International Journal of Forecasting.

Ravn, Morten and Harald Uhlig (2002), "Notes on Adjusting the Hodrick-Prescott Filter for the Frequency of Observations," *The Review of Economics and Statistics*.

Schafer, Juliane and Korbinian Strimmer (2005), "A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics," *Statistical Applications in Genetics and Molecular Biology*.

Taieb, Souhaib Ben (2017), "Sparse and Smooth Adjustments for Coherent Forecasts in Temporal Aggregation of Time Series," *Proceedings of Machine Learning Research*.

Wickramasuriya, Shanika, George Athanasopoulos, and Rob Hyndman (2019), "Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization," *Journal of the American Statistical Association*.



PUBLICATIONS

Smooth Forecast Reconciliation
Working Paper No. WP/2024/066