

Department of Econometrics and Business Statistics

<http://business.monash.edu/econometrics-and-business-statistics/research/publications>

Forecasting Swiss Exports using Bayesian Forecast Reconciliation

Florian Eckert, Rob J Hyndman and Anastasios
Panagiotelis

July 2019

Working Paper 14/19

Forecasting Swiss Exports using Bayesian Forecast Reconciliation

Florian Eckert^{*†}, Rob J. Hyndman[‡], and Anastasios Panagiotelis[‡]

[†]KOF Swiss Economic Institute, ETH Zurich

[‡]Department of Econometrics and Business Statistics, Monash University

July 2019

Abstract

This paper conducts an extensive forecasting study on 13,118 time series measuring Swiss goods exports, grouped hierarchically by export destination and product category. We apply existing state of the art methods in forecast reconciliation and introduce a novel Bayesian reconciliation framework. This approach allows for explicit estimation of reconciliation biases, leading to several innovations: Prior judgment can be used to assign weights to specific forecasts and the occurrence of negative reconciled forecasts can be ruled out. Overall we find strong evidence that in addition to producing coherent forecasts, reconciliation also leads to improvements in forecast accuracy.

JEL Classification: C32, C53, E17

Keywords: Hierarchical Forecasting, Bayesian Forecast Reconciliation, Swiss Exports, Optimal Forecast Combination.

*Corresponding Author: Leonhardstrasse 21, 8092 Zürich, eckert@kof.ethz.ch, +41 44 632 29 80

1 Introduction

Export forecasts can support economic policy makers, monetary authorities and exporting firms in making decisions. They may be of interest in their own right but also as inputs into projections of other important macroeconomic quantities such as currency reserves, exchange rates and production growth. As exports of goods are usually measured and published on a highly disaggregate basis, many economic agents are interested in forecasts at a more granular level. For instance, a Swiss manufacturer of precision instruments is concerned about exports of Swiss watches into individual countries in order to manage their inventories. The large collection of 13,118 time series on Swiss goods exports is subject to known linear hierarchical constraints. Total exports from Switzerland can be disaggregated geographically by destination into regions such as Europe, North America or Australia. These regional aggregates can then be divided further by country. Total exports can also be disaggregated into product categories such as precision instruments, textiles or vehicles and then further into subcategories such as road, rail, air and water vehicles. As such, the data has the structure of a so-called grouped hierarchy (see Hyndman and Athanasopoulos, 2018, and references therein). Figure 1 gives a simple example of grouped structure with $k = 3$ levels, $m = 9$ series in total and $q = 4$ series at the most disaggregate or ‘bottom’ level.




Figure 1: Simple Example of a Grouped Hierarchy.

Since it is known that all future realizations of the data will adhere to the constraints implied by the aggregation structure, a desirable property of any forecasts is that they also respect these constraints. Such forecasts are referred to as ‘coherent’. In the case of Swiss exports, incoherent forecasts are problematic because they may lead to contradictory conclusions, non-aligned decision making and are difficult to communicate. Earlier literature reduced the issue of producing coherent forecasts to one of predicting only a specific level of the hierarchy. For example, the ‘bottom-up’ approach (Gross and Sohl, 1990) achieves coherence by producing only forecasts for the bottom level series and then summing these up according to the hierarchical structure. A major shortcoming of this approach is that disaggregate series tend to be noisy and there is a high risk of model misspecification. Features such as seasonality may be difficult to identify in the bottom level data, despite being clearly present in the aggregate series. To address this shortcoming a ‘top-down’ approach was proposed (see Athanasopoulos et al., 2009, and references therein), where the predicted top level series is disaggregated according to historical or forecasted proportions

of lower levels. A compromise is given by the ‘middle-out’ approach, where the forecasts at an intermediate level of the hierarchy are summed up to get the higher levels and disaggregated to obtain lower level predictions. A weakness of these single level methods is information loss because the time series characteristics at other levels are not taken into account. A further shortcoming of the middle-out and top-down methods is that they are not easily applied to grouped hierarchical structures.

In response to these shortcomings, there has been a tendency over the past decade towards producing forecasts for all series in the hierarchy rather than only at a single level. These are referred to as ‘base’ forecasts and they generally do not adhere to aggregation constraints. ‘Forecast reconciliation’, introduced by Hyndman et al. (2011), performs an ex-post adjustment to base forecasts in order to produce a new set of coherent forecasts. This adjustment effectively combines predictions from all levels and in doing so ‘hedges’ against misspecification error across all levels. There is now substantial theoretical and empirical evidence that forecast reconciliation can significantly improve forecast accuracy (see Wickramasuriya et al., 2018, and references therein). The first main contribution of our own paper is therefore to apply existing reconciliation methods to the problem of forecasting Swiss export data.

Despite their success, forecast reconciliation methods suffer a number of shortcomings, some arising in general and others due to the idiosyncrasies of our dataset. First, forecast reconciliation can induce negative forecasts even when the base forecasts are non-negative. This is clearly problematic when forecasting quantities that must be non-negative by construction. Second, an important theoretical assumption to ensure optimality of forecast reconciliation is that forecasts are unbiased prior to reconciliation. Often this assumption fails to hold in practice. Third, existing reconciliation methods combine forecasts in a way that is backward-looking, combining models using weights that depend on the inverse of past variation of forecast errors. There are several instances where this approach may break down for our dataset. One is the case of exports to small trading partners which are zero for most (and in some cases all) of the observations in the training sample. These can be predicted almost perfectly by a naive forecast of zero, leading to no variation in the forecast errors and numerically unstable reconciliation. Another case where the backward-looking nature is inappropriate occurs when forecasters have information that a forecasting model may break down even if it has performed well in the past. An example of this in our dataset is a structural break induced by the reclassification of electricity as a good rather than a service. Finally, the backward-looking nature of existing reconciliation methods also makes it difficult to exploit information from the entire predictive density of a forecast target.

The second major contribution of our paper is to extend existing forecast reconciliation in a way that addresses these shortcomings. The main innovation is to convert the general reconciliation equation into a panel regression. Rather than using a single vector of point

forecasts, we draw samples from the predictive densities of the base forecast models. This panel regression is estimated using Bayesian Markov Chain Monte Carlo algorithms. Our approach has a number of benefits relative to the existing methodology. First, reconciliation biases in the base forecasts are made explicit by introducing fixed effects into the panel regression structure. This overcomes a major weakness of existing approaches which require base forecasts to be unbiased to ensure that reconciliation is optimal. Second, we propose a mechanism for down-weighting the influence of particular series irrespective of past forecasting performance. This is valuable if forecasters have strong judgmental reasons for believing that a particular model will work well in the future while other base forecasts will be less reliable. Third, under our approach, the weights used in reconciliation can depend on the variances of the predictive density rather than in-sample forecast errors. Existing approaches essentially determine weights based on estimates of the unconditional variance. Our innovation is therefore particularly promising for forecasting models that allow for conditional heteroskedasticity. Fourth, the Bayesian estimation procedure conveniently allows the incorporation of prior information to solve issues such as the occurrence of negative reconciled forecasts and singular forecast error covariance matrices.

The remainder of the paper is structured as follows. Section 2 introduces in detail the data on exports of Swiss goods, using modern techniques for exploring and visualizing high-dimensional time series. Section 3 introduces both existing forecast reconciliation as well as our novel Bayesian approach. Section 4 conducts an extensive forecast evaluation that compares our proposed method with existing reconciliation techniques. Section 5 concludes.

2 Data

We use a comprehensive dataset containing exports of Swiss goods. All time series cover a period from 1988 to 2018 in monthly frequency and are denominated in Swiss francs. They are not adjusted for seasonalities or calendar effects. The data can be grouped by export destination and product category. The geographic hierarchy consists of 8 regions, aggregated from 245 countries and dependent territories. The categorical hierarchy follows a national nomenclature covering 12 main economic groups and 48 subgroups¹. This leads to a grouped hierarchy with $m = 13,118$ series containing at least one nonzero entry of which $q = 9,483$ series are at the bottom level. Figure 2 shows the historical development of the regional and categorical hierarchies.

¹Precious metals, precious and semi-precious stones, works of art and antiques are generally omitted in business cycle research due to volatility and structural breaks.

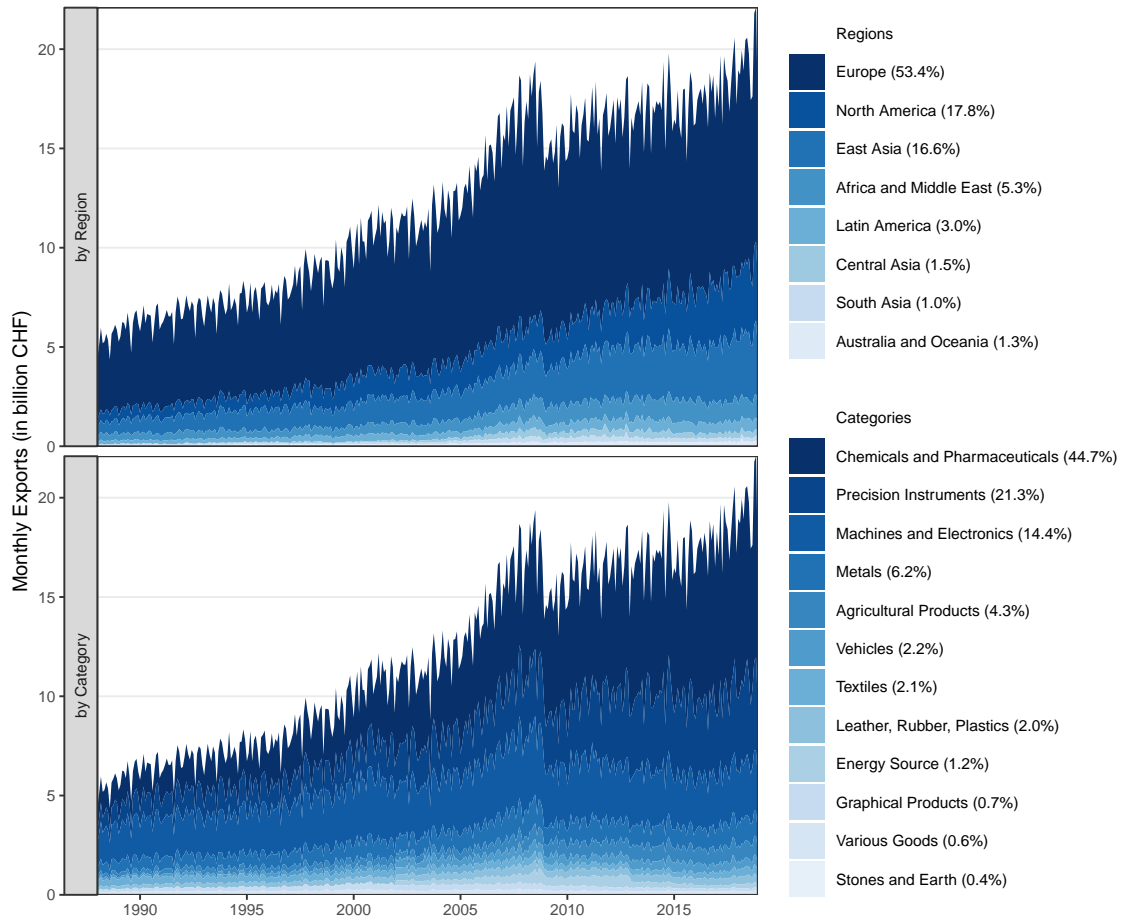


Figure 2: Contribution to Swiss Exports of Goods. *Nominal values, not adjusted for seasonalities or calendar effects. Average export shares of the year 2018 in parentheses.*

As a result of its status as a small open economy in a rapidly globalizing world, Swiss exports have increased significantly since the late 1980s. Accounting for more than half of total exports, Europe is a key market for Swiss goods. Increasingly larger shares of exports also go to North America and East Asia with around 17% each in 2018. Exports to Africa and the Middle East, Latin America, Central and South Asia and Australia account only for about 10% combined.

The hierarchical grouping by categories is more evenly distributed, but has been subject to greater shifts in its composition. The most important categories are ‘Chemicals and Pharmaceuticals’, ‘Precision Instruments’ and ‘Machines and Electronics’. Figure 3 shows the changes in composition between 1988 and 2018.




Figure 3: Regional and Categorical Composition of Swiss Goods Exports.

The two hierarchical groupings are quite different. The geographic hierarchy with 8 groups and 245 subgroups is wide, but with a majority of the export volume going to European countries, it is nevertheless highly concentrated. This has changed slightly in the past 30 years as the relative share of exports to the rest of the world has increased. The categorical hierarchy on the other hand is rather narrow with 12 groups and 48 subgroups. Compared to the regional hierarchy, the export volume is however more evenly distributed, even though an increasing concentration, particularly in chemicals and pharmaceuticals, can be noted.

Due to the aggregation involved, top level series are usually less noisy and exhibit more predictable characteristics such as seasonality or trend. Following Kang et al. (2017), it is possible to construct a measure of predictability for each time series by estimating principal components from a number of time series features that are commonly associated with better predictability. This includes measures such as the strength of seasonality, trend, spectral entropy and serial correlation. On the vertical axis, Figure 4 shows the first principal component, which accounts for a large share of the variation in these predictability features.




Figure 4: Predictability of Different Levels in a Hierarchy. *Predictability is defined as the first principal component of a large number of time series characteristics.*

It is evident that there exists a strong correlation between predictability and export volume. This implies that larger series and consequently those at the top of a hierarchy are easier to forecast.

3 Forecast Reconciliation

3.1 Existing Forecast Reconciliation Methods

In order to encode the aggregation constraints in a hierarchy, we define Y_t to be an m -vector that stacks observations at time t from all series, B_t to be a subvector of Y_t containing only the q bottom level series at time t and S to be an $m \times q$ aggregation matrix. In the simple grouped hierarchy shown in Figure 1, these are given by

$$\begin{array}{c}
 \begin{matrix} Y_0 \\ Y_A \\ Y_B \\ Y_1 \\ Y_2 \\ Y_{A1} \\ Y_{A2} \\ Y_{B1} \\ Y_{B2} \end{matrix} \\
 \begin{matrix} (m \times 1) \end{matrix}
 \end{array}
 =
 \begin{array}{c}
 \begin{matrix} [columns - width = 5mm] 1 & 1 & 1 \\ & 1 & 0 & 0 \\ & 0 & 1 & 1 \\ & 1 & 0 & 1 & 0 \\ & 0 & 1 & 0 & 1 \\ & 1 & 0 & 0 & 0 \\ & 0 & 1 & 0 & 0 \\ & 0 & 0 & 1 & 0 \\ & 0 & 0 & 0 & 1 \end{matrix} \\
 \begin{matrix} (m \times q) \end{matrix}
 \end{array}
 =
 \begin{array}{c}
 \begin{matrix} Y_{A1} \\ Y_{A2} \\ Y_{B1} \\ Y_{B2} \end{matrix} \\
 \begin{matrix} (q \times 1) \end{matrix}
 \end{array}$$

Here and in general, the matrix S is defined so that $Y_t = SB_t$ holds for all realized data. Hyndman et al. (2011) considered a framework whereby forecasts are produced for all m series at every level, referring to these as ‘base forecasts’. To reconcile these base forecasts, the following regression structure was assumed.

$$Y_t(h) = S\beta_h + e_t(h), \quad (1)$$

where $Y_t(h)$ is an m -vector containing the h -periods-ahead base forecasts at time t for each level in the hierarchy, β_h represents the true expected value of the bottom level series, and the error term $e_t(h)$ has mean zero and covariance matrix Σ_h . Reconciled forecasts are given by Sb_h , where b_h is an estimate of β_h that combines information about forecasts at all levels. It can be estimated using the following regression equation,

$$b_h = \left(S'W_h^{-1}S\right)^{-1} S'W_h^{-1}Y_t(h). \quad (2)$$

This choice minimizes the generalized Euclidean distance between $Y_t(h)$ and the reconciled forecasts Sb_h with respect to W_h . Reconciliation is also guaranteed to reduce the distance to the eventual realization targeted by a forecast. There are several potential choices for W_h . Letting $W_h = I$ corresponds to an ordinary least squares estimate. Alternatively, a high degree of heteroskedasticity in the error terms motivates a diagonal W_h or weighted least squares approach (Hyndman et al., 2016). Under so-called ‘variance scaling’, weights are the variances of in-sample h -step ahead forecast variances, and forecasts with less accurate historical performance are down-played in reconciliation. Another alternative is the ‘nseries’ approach due to Athanasopoulos et al. (2017), whereby weights are based on the number of series aggregated at each node. More recently, the ‘MinT’ approach was developed by Wickramasuriya et al. (2018) to allow for a W_h that is not diagonal and exploits the covariances between the h -step-ahead reconciled forecast errors. The nomenclature MinT refers to the fact that this approach minimizes the trace of the covariance matrix of reconciliation errors.

3.2 Bayesian Forecast Reconciliation

In this part, we propose a new methodology for forecast reconciliation. The main insight is that additional information about the uncertainty surrounding base forecasts can be incorporated into the reconciliation procedure. In the spirit of Kapetanios et al. (2015) or Amisano and Geweke (2017), the predictive distributions of the m base forecast models are approximated with a sample of n draws. Possible sources for obtaining these draws are posterior predictive distributions from Bayesian forecasting models, bootstrap aggregating (Bergmeir et al., 2016), model pooling, or sampling from a fitted model. This results in n vectors \hat{y}_i , each of length m , that contain a draw i from each predictive distribution.²

²Since every forecast horizon is reconciled independently, the time subscripts are dropped from now on to simplify notation.

This allows the regression model from equation (1) to be recast as a panel regression. The error term consists of two components, a prediction error e_i and a reconciliation bias α . The latter can be interpreted as a fixed effect that is unique to each forecasted variable. In other words, α is the difference between the unreconciled forecast mean \hat{y} and the reconciled forecast mean \tilde{y} . The interpretation of β depends on the definition of S , but in general it estimates the mean of the bottom level reconciled forecasts. The following equation can then be used to model the forecast reconciliation:

$$\begin{matrix} \hat{y}_i & = & \alpha & + & S & \times & \beta & + & e_i, \\ (m \times 1) & & (m \times 1) & & (m \times q) & & (q \times 1) & & (m \times 1) \end{matrix} \quad (3)$$

where e follows a normal distribution with mean zero and covariance matrix Σ . The reconciliation problem in (3) can be expressed as a system of seemingly unrelated regressions (SUR) to account for cross-equation correlations. Since the explanatory variables are the same for each equation, it is a special case of the SUR model in Zellner (1962). However, the parameters are impossible to estimate directly because of perfect multicollinearity in the regressors I_m and S . This is quite intuitive since there is more than one unique way to reconcile incoherent forecasts. Following Farebrother (1978), the regression is partitioned in order to separate the parameters that cause multicollinearity and they are estimated in separate Gibbs sampling steps. The distribution of the reconciled forecasts can be obtained by sampling from the posterior predictive distribution conditional on α being equal to zero:

$$p(\tilde{y} \mid \alpha = 0, \hat{y}, S) = \iint p(\tilde{y} \mid \alpha = 0, \beta, \Sigma, \hat{y}, S) p(\alpha = 0, \beta, \Sigma \mid \hat{y}, S) d\beta d\Sigma.$$

The joint posterior distribution of α , β and Σ is obtained by combining a prior belief on the parameters with the likelihood according to Bayes' theorem:

$$p(\alpha, \beta, \Sigma \mid \hat{y}, S) \propto p(\hat{y}, S \mid \alpha, \beta, \Sigma) \times p(\alpha, \beta, \Sigma). \quad (4)$$

The likelihood function of the data is given by

$$p(\hat{y}, S \mid \alpha, \beta, \Sigma) \propto \frac{1}{|\Sigma|} \exp \left[-\frac{1}{2} \sum_i (\hat{y}_i - \alpha - S\beta)' \Sigma^{-1} (\hat{y}_i - \alpha - S\beta) \right].$$

The joint posterior distribution is accordingly given by

$$\begin{aligned}
p(\alpha, \beta, \Sigma \mid \hat{y}, S) &\propto \frac{1}{|\Sigma|} \exp \left[-\frac{1}{2} \sum_i (\hat{y}_i - \alpha - S\beta)' \Sigma^{-1} (\hat{y}_i - \alpha - S\beta) \right] \\
&\times \exp \left[-\frac{1}{2} (\alpha - a_0)' A_0^{-1} (\alpha - a_0) \right] \\
&\times \exp \left[-\frac{1}{2} (\beta - b_0)' B_0^{-1} (\beta - b_0) \right] \\
&\times \frac{1}{|\Sigma| (v_0 - m - 1)} \exp \left[-\frac{1}{2} \text{tr}(R_0^{-1} \Sigma^{-1}) \right].
\end{aligned}$$

The Bayesian approach has the advantage that uncertainty surrounding the parameters α , β and Σ is taken into account. Following Percy (1992), we get the marginal distributions by approximating the joint posterior distribution via Gibbs sampling from the conditional distributions. Convergence is achieved quickly, irrespective of the starting values. It can be verified by testing for stability in the recursive means of the Markov chains. A sufficiently large sample of draws from the posterior predictive distribution of \tilde{y} is saved and evaluated to get summary statistics such as the mean and variance of the reconciled forecasts.

Step 1: Draw β conditional on $\alpha, \Sigma, \hat{y}, S$

The parameter β is the mean of the bottom level forecasts, given an appropriate aggregation matrix S . The conditional posterior distribution is then given by

$$\beta \mid \alpha, \Sigma, \hat{y} \sim N(b_1, B_1), \quad (5)$$

where $B_1 = \left(\sum_i S' \Sigma^{-1} S + B_0^{-1} \right)^{-1}$ and $b_1 = B_1 \left(\sum_i S' \Sigma^{-1} (\hat{y}_i - \alpha) + B_0^{-1} b_0 \right)$. Unless there is reason to believe otherwise, the priors b_0 and B_0 should be chosen to be as uninformative as possible. In some cases, this regression approach leads to negative values in the reconciled bottom level forecasts. This might be a concern since many applications such as sales or exports do not allow for negative observations. Using a truncated normal prior, this issue can be resolved in an uncomplicated fashion by simply discarding draws of β that contain negative entries during the sampling process.

Step 2: Draw Σ conditional on $\alpha, \beta, \hat{y}, S$

Σ is the covariance matrix of the prediction errors. While Σ is not singular by definition because the forecasts in \hat{y}_i are not reconciled, it might be near-singular if the base forecasting models are estimated jointly or if the draws are reordered following Jeon et al. (2018). In the latter case Σ can be drawn from an inverse Wishart distribution.

$$\Sigma \mid \alpha, \beta, \hat{y} \sim W^{-1}(v_1, R_1), \quad (6)$$

where $v_1 = v_0 + n$ and $R_1 = \left(R_0^{-1} + \sum_i (\hat{y}_i - \alpha - S\beta)' (\hat{y}_i - \alpha - S\beta) \right)^{-1}$. It is useful to set an almost uninformative prior with v_0 and R_0 close to zero, which introduces a small amount of noise into the reconciled forecasts. This has negligible impact on the posterior

distribution, but ensures that Σ is nonsingular in the case where a base forecast has no variation. A possible simplification is to draw the variances equation-by-equation from an inverse gamma distribution.

Step 3: Draw α conditional on $\beta, \Sigma, \hat{y}, S$

Because the reconciliation regression is an ill-posed problem, it is necessary to impose additional restrictions on the reconciliation biases α in order to achieve identification. The conditional distribution of α can be expressed equivalently by concentrating out β in the following reconciliation identity.

$$\alpha = \frac{1}{n} \sum_i \hat{y}_i - S\beta. \quad (7)$$

In order to eliminate β from equation (7), both sides are multiplied by a projection matrix P . The reconciliation biases depend greatly on the definition of this projection matrix and alternative choices for P will be discussed in Section 3.3. Using $P = S(S'S)^{-1}S'$ implies an orthogonal projection onto the coherent subspace. The resulting terms are then subtracted from both sides of equation (7):

$$(I_m - P)\alpha = (I_m - P)\frac{1}{n} \sum_i \hat{y}_i. \quad (8)$$

It is useful to define the idempotent residual maker $M = I_m - P$. Since M is not invertible due to the presence of multicollinearity, equation (8) cannot be solved for α . Our identifying assumption is that α lies in the span of M in which case $M\alpha = \alpha$. For $P = S(S'S)^{-1}S'$, this implies that the direction of the reconciliation bias is orthogonal to the coherent subspace. This solves the identification problem and leaves the reconciliation biases as a function of the data and the residual maker M . This result is again intuitive since the reconciliation biases are the residuals from a regression of the base forecasts on the aggregation matrix:

$$\alpha = M \left(\frac{1}{n} \sum_i \hat{y}_i \right). \quad (9)$$

Having identified the system in that way, the prior variance A_0 is allowed to be uninformative and the prior mean a_0 is a zero vector. A numerically stable conditional posterior for α is therefore given by

$$\alpha \mid \beta, \Sigma, \hat{y} \sim N(a_1, A_1), \quad (10)$$

where $A_1 = M \left(\frac{\Sigma}{n} \right) M'$ and $a_1 = M \left(\frac{1}{n} \sum_i \hat{y}_i \right)$.

3.3 Bias Weighting

The definition of the projection matrix P is crucial for the estimated parameters. Figure 5 demonstrates the impact of different projections on the estimated reconciliation biases. It features identical unreconciled forecasts of a simple hierarchy with $m = 3$ series, where $Y_A + Y_B = Y_0$. For each series a sample is drawn from the predictive forecasts density, assumed to be $N(4, 2)$ for Y_A (shown in blue), $N(6, 1)$ for Y_B (shown in purple), and $N(16, 3)$ for Y_0 (shown in yellow). The horizontal axis shows draws from the unreconciled base forecasts, which are clearly incoherent. The vertical axis on the other hand shows the means of the reconciled forecasts. The diagonal line shows values where the means of base and reconciled forecasts are equal. For boxes above this diagonal line, reconciliation adjusts forecasts upwards, for boxes below the diagonal line reconciliation adjusts forecasts downwards.

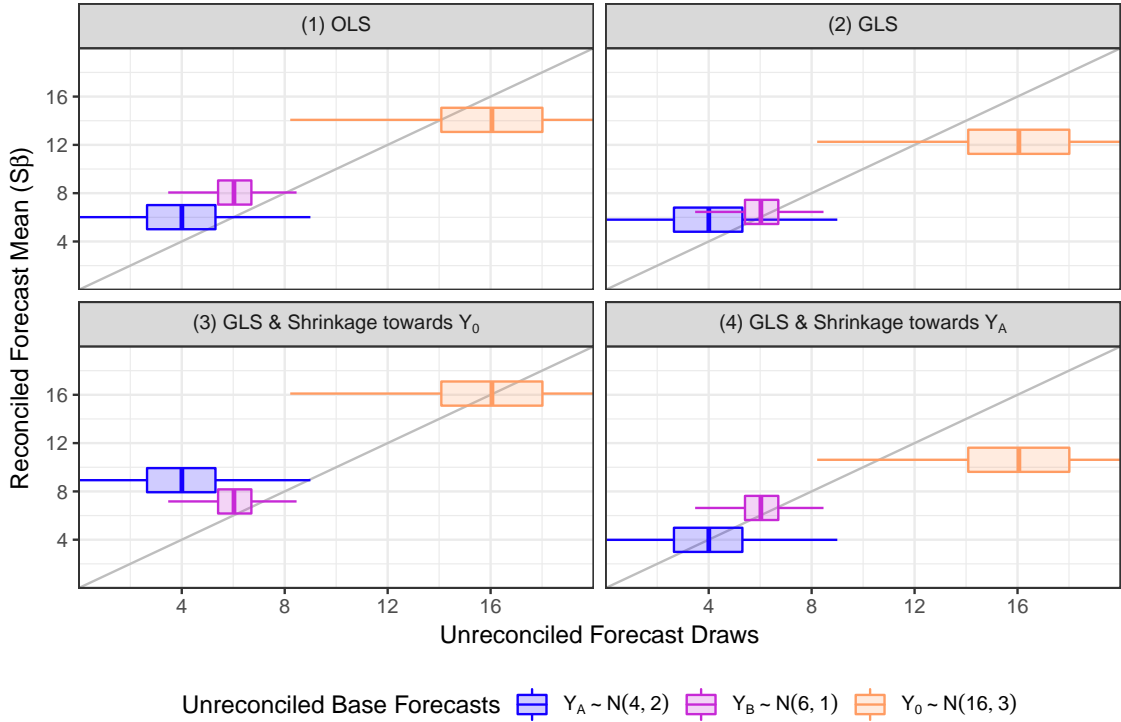


Figure 5: Weighting Schemes. The grey line indicates where the unreconciled base forecasts on the x-axis are equal to reconciled forecast means on the y-axis.

Each panel corresponds to a different choice of P . Using $P = S(S'S)^{-1}S'$ corresponds to the orthogonal projection in an ordinary least squares regression. Subfigure (1) shows that the forecast biases for each margin are treated equally, consequently the means of Y_A and Y_B are adjusted upwards while the mean of Y_0 is adjusted downwards. Using $P = S(S'\Sigma^{-1}S)^{-1}S'\Sigma^{-1}$ implies that the reconciliation biases are weighted with the inverse of their corresponding forecast variances. This leads to a smaller adjustment in Y_B (the reconciled and base means are close) relative the others since it is more accurate.

Even though it is intuitive to weight the reconciliation biases using the predictive accuracy of the corresponding base forecasts, this can be generalized to different weighting schemes. There may exist prior information on the reliability of certain models or the requirement to fix some forecasts at specific values. This could be due to better data availability, higher suitability of a particular model or subjective judgment of the forecaster. Weighting can be achieved using a projection matrix $P = S(S'(\Lambda\Sigma\Lambda')^{-1}S)^{-1}S'(\Lambda\Sigma\Lambda')^{-1}$ that includes a diagonal matrix of weights Λ . It may be of interest to selectively shrink some reconciliation biases in α towards zero by decreasing the corresponding entry in Λ . At the same time, it is necessary to increase the remaining elements such that they are able to capture the higher reconciliation biases at their level of the hierarchy. This is achieved by keeping the total dispersion of the multivariate normal distribution $\Lambda\Sigma\Lambda'$ constant. A common measure is the generalized variance described in Mustonen (1997) and defined as the determinant of a covariance matrix. The weighting matrix Λ is therefore always constructed such that the product of the diagonal elements remains constant at 1. This in turn ensures that the total dispersion of $\Lambda\Sigma\Lambda'$ remains equal to the total dispersion of the unweighted Σ for all Λ . Subfigures (3) and (4) shrink the reconciled forecasts of Y_0 and Y_A towards their base forecasts.

It is important to note that the likelihood is invariant to these choices. Besides the shrinkage of specific reconciliation biases towards zero, there are several other weighting methods conceivable. Possible approaches include the weighting of each series by its level in the hierarchy or by to the number of series at each node in the hierarchy. This allows for the emulation of the ‘nseries’, ‘bottom-up’, ‘middle-out’ and ‘top-down’ results. A convenient feature of this is that the ‘middle-out’ and ‘top-down’ shrinkage work also for grouped time series, which is not the case in the standard approach.

4 Reconciliation of Export Forecasts

4.1 Setup

The large hierarchy of Swiss goods exports is used to test the Bayesian reconciliation framework and various competing methods. Each month from 1995 to 2015, forecasts for all series in the hierarchy are calculated for the next 36 months. For each of the 13,118 series, forecasts are calculated from three models: an autoregressive integrated moving average model (ARIMA), an exponential smoothing state space model (ETS) and a seasonal random walk model (RW). As described in Hyndman and Khandakar (2008), the model for each series is parametrized automatically based on the Akaike Information Criterion. In order to get samples from the predictive densities, $n = 1000$ sample paths are simulated from each fitted model using Gaussian errors. With the exception of the volatile period during the Great Recession, the ARIMA and ETS approaches outperform the Random Walk on average for series at every level and forecasting horizon. All results

in the following subsection will therefore rely on ARIMA forecasts.³

These incoherent forecasts are then reconciled using several basic single level and optimal combination methods. The single level techniques include bottom-up, top-down and middle-out methods. The latter two can only be used for non-grouped time series and are therefore tested on the regional and categorical hierarchies separately. The optimal combination methods used are ordinary and weighted least squares, nseries, MinT and Bayesian reconciliation (BSR). Draws from the predictive distributions are obtained by sampling from the fitted models assuming normality of the errors. If aggregation of the prediction errors is necessary, they are weighted with their respective export share.

The resulting coherent forecasts are evaluated using several accuracy measures such as the root mean squared error (RMSE), mean absolute percentage error (MAPE) and mean absolute scaled error (MASE). The method of Diebold and Mariano (1995) is used to test whether reconciled forecasts are significantly more accurate than unreconciled forecasts. The Diebold-Mariano test checks for significance in the difference between two squared forecast errors at various forecasting horizons, accounting for serial correlation in the squared error loss. In addition, another significance test is used to compare the mean squared errors directly. Since the prediction errors are assumed to be normally distributed, the ratio between mean squared errors of an unreconciled forecast and a specific reconciled forecast has an F -distribution with degrees of freedom corresponding to the number of errors. This allows us to test for equality of the unreconciled and reconciled mean squared prediction errors.

4.2 Comparative Results

This section provides empirical evidence for the benefits of optimal hierarchical combination. It compares the performance of different reconciliation methods and explores which data characteristics profit in particular from hierarchical combination.

Benefits of Hierarchical Combination. Figure 6 shows the accuracy of forecasts, defined as the mean squared error of the base forecasts relative to the mean squared errors of the coherent forecasts from each method. Higher bars indicate therefore better forecast performance. The 95% confidence interval shows the acceptance region of an F -test for equality of the mean squared errors. It is worth noting that reconciliation methods and bottom-up forecasts are the only techniques that allow for coherence across all levels of a grouped hierarchy. Top-down and middle-out reconciliations are not applicable in the case of grouped time series.

It is evident that single level methods do not consistently improve forecasting accuracy. The bottom-up and middle out methods fare reasonably well for the top level series, but fail to outperform the unreconciled forecasts at lower levels and are sometimes even

³A comparison of forecasting methods, horizons and accuracy measures can be found in Appendix A.1.

significantly worse. Optimal combination on the other hand tends to outperform the base forecasts especially for top and intermediate level series. Especially the methods using variance scaling, such as MinT, WLS and BSR, work well at all levels.

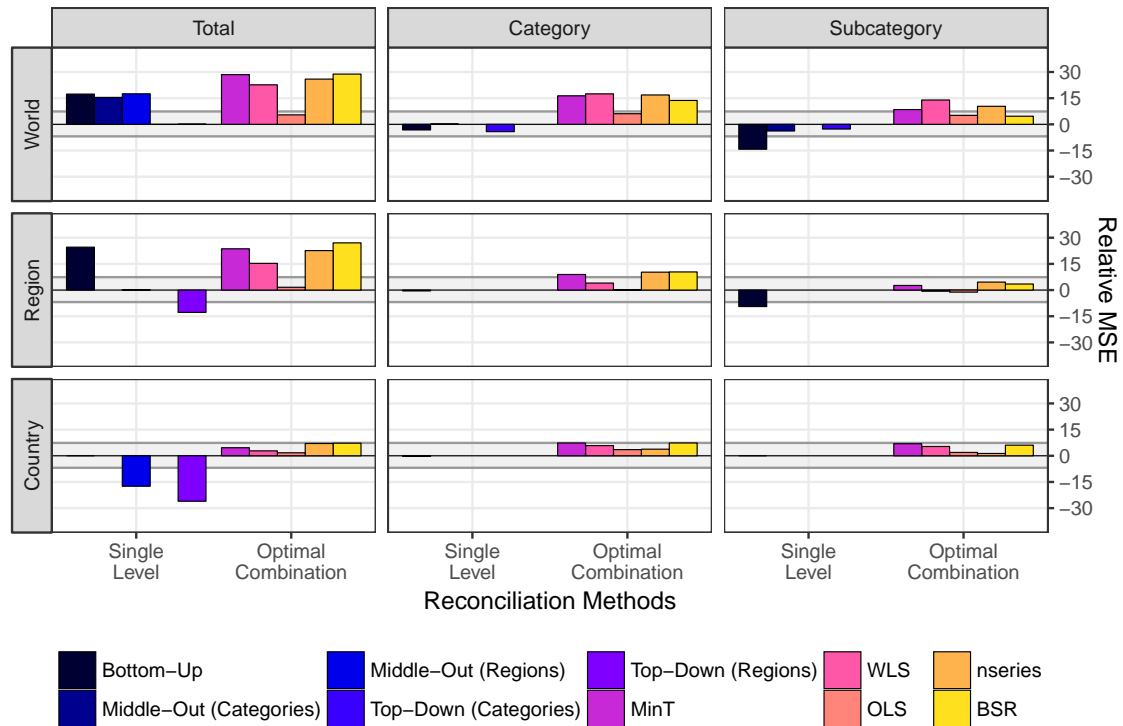


Figure 6: Relative Accuracy of Reconciliation Methods. Higher bars indicate better forecasts relative to the unreconciled case. The zero line shows the accuracy of unreconciled predictions, with a 95% confidence interval for the ratio of mean squared errors to be one. Average of all forecast horizons.

Comparison of Combination Methods. It is also instructive to look at the development of the relative forecasting accuracy over time in Figure 7. Even though the variance scaling methods are more accurate on average, they do not consistently outperform the unreconciled forecasts. It also appears that MinT, WLS and BSR perform fairly similar over time. For the top level series, the benefits of reconciliation accrue mostly during times of global economic distress and corresponding appreciations of the Swiss franc. This is due to the fact that the simpler models at lower levels provide stability at times when the top level model is biased. The biggest gains can be observed during the early 2000s recession following the burst of the dot-com bubble, the global financial crisis and the following sovereign debt crisis in Europe, and the sudden appreciation of the Swiss franc after the Swiss National Bank stopped supporting the currency peg to the Euro in early 2015. Interesting is the forecasting accuracy after January 2002, when electrical energy was reclassified as a good instead of a service. The structural break in the time series leads to misspecified models, but the rigid structure imposed by the hierarchy increases forecast accuracy substantially relative to the unreconciled case.




Figure 7: Relative Accuracy of Reconciliation Methods over Time. Higher lines indicate better forecasts. Zero line shows the accuracy of unreconciled predictions, with a 90% confidence interval for the ratio of mean squared errors to be one. Base forecasts are generated using ARIMA models. Average of all forecast horizons.

Significance of Results. In order to check whether the accuracy improvements are significant, one-sided Diebold-Mariano tests are used for the top-level series. Figure 8 shows the p-values when testing for equality of reconciled and unreconciled forecasts. The alternative hypothesis is that the accuracy of reconciliation methods is greater. With the exception of the middle-out approach, single level methods are not significantly more accurate than the unreconciled forecasts at all horizons. Optimal combinations are associated with lower P-values, in particular OLS. This perhaps reflects the stability of OLS reconciliation which does not require the estimation of a large weighting matrix and therefore leads to forecast errors with a smaller variance.




Figure 8: Significance of Forecast Accuracy Improvements. *Test for top-level total exports series, where points indicate p-values of Diebold-Mariano test for greater accuracy of reconciled versus unreconciled forecasts. Lines indicate locally estimated scatterplot smoothed (LOESS) p-values.*

Implications for Data Characteristics. Another way to dissect the results is to identify which time series see the greatest gains in forecast accuracy from using reconciliation. Figure 9 provides an overview of the relative forecast accuracy by geographic classification, using the Bayesian reconciliation framework. It is again obvious that reconciled forecasts are on average more accurate than in the unreconciled case, but not in every instance. It appears that series with a larger export volume benefit most from reconciliation. Forecasts of exports to countries in Europe, North America and East Asia are almost entirely better off than in the unreconciled case, whereas forecasts of exports to countries with a lower share, such as the islands in Oceania, tend to be worse off. In addition, time series at higher levels in a hierarchy do not necessarily profit more from reconciliation than their corresponding subcategories.




Figure 9: Relative Accuracy of Reconciliation Methods by Regions. *Higher points indicate better forecasts. Zero line shows the accuracy of unreconciled predictions. Reconciliation using unweighted BSR.*

The same results also hold true for the relative forecast accuracy by categories, as shown in Figure 10. Because the export shares in the categorical hierarchy are more evenly distributed, the pattern of smaller export volumes being worse off due to reconciliation is less pronounced. The results for other variance scaling methods such as weighted least squares and MinT are similar.




Figure 10: Relative Accuracy of Reconciliation Methods by Categories. *Higher points indicate better forecasts. Zero line shows the accuracy of unreconciled predictions. Reconciliation using unweighted BSR.*

4.3 Benefits of Weighting

An advantage of the general weighting scheme is that selected series can be shrunk towards their base forecast. This is particularly useful if there exists judgmental information for a specific forecast that would require adjustments for all other base forecasts in a hierarchy. An example is the reclassification of electricity as a good instead of a service. Figure 11 shows total exports of energy sources for unweighted and weighted reconciliation.

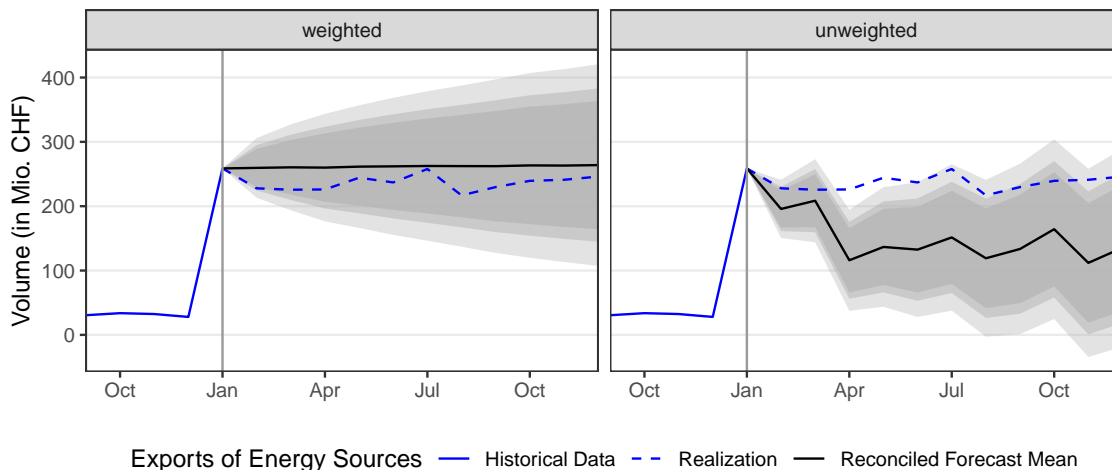


Figure 11: Forecast with Bias Weighting. *Left panel shows weighted, right panel unweighted reconciliation with identical random walk base forecast for energy exports. Grey ribbons indicate 90%, 95% and 99% prediction intervals.*

After observing the first value including electric energy in January 2002, a random walk forecast is used for this particular series. Then the series are reconciled with and without an appropriate weighting scheme. For the unweighted model on the right, the other base forecasts assume the structural break to be an outlier and dominate any information from the random walk forecast. Even though the forecaster has prior knowledge that the random walk forecast is appropriate, it is overruled in the reconciliation procedure. The only possibility is a cumbersome adjustment of the base forecasts for all other series as well.

The model on the left shows the reconciliation of the same forecasts with more weight on the random walk forecast. This is done as described in Section 3.3. The diagonal entry in Λ that corresponds to the random walk forecast is scaled down. At the same time, the remaining entries are scaled up such that the determinant of Λ remains at 1. This forces the reconciled forecast for energy sources to stay close to its random walk base forecast. The remaining series adjust accordingly during the reconciliation procedure. As a result, the mean squared error of the forecast for energy sources in 2002 is more than 90% lower than in the unweighted case. This also leads to significant accuracy gains at other levels; the forecast of total exports for instance is 14% more accurate in 2002.

5 Conclusion

This paper extends the existing literature on hierarchical forecast combination by establishing a Bayesian estimation framework and introducing an explicit definition of the reconciliation bias. This leads to several innovations: It is possible to use subjective judgment of the forecaster to assign weights to the forecast biases and shrink selected reconciled forecasts towards their corresponding base forecasts. The Bayesian sampling procedure allows for the incorporation of prior information on the parameters. This avoids some issues such as the occurrence of negative reconciled forecasts and singular forecast error covariance matrices. The use of predictive densities allows for greater flexibility in the choice of the base forecast models, taking for instance conditional heteroskedasticity into account when weighting the forecasts at different horizons. Confidence intervals can be computed by drawing from the predictive posterior distribution of the reconciled forecasts. However, the approach tends to be slower than established reconciliation techniques because it requires repeated sampling from the joint posterior distribution.

Using a comprehensive dataset of Swiss goods exports, this paper demonstrates that optimal combination methods using variance scaling improve the forecasting accuracy significantly compared to the unreconciled case and simpler reconciliation methods. The results are robust to changes in the forecasting horizon, the underlying base forecast models and the measure used to determine forecasting accuracy. Optimal combination methods are shown to be particularly useful in the case of misspecified models and during periods of high volatility in the time series. Even though the forecasting accuracy is significantly better on average, no reconciliation method consistently outperforms the unreconciled forecasts across the hierarchy or over time. Forecasts at the top of level tend to benefit more from reconciliation than the noisy series at the bottom of a hierarchy. At the same level, forecasts that account for a larger share of the total are on average more accurate after reconciliation. The paper also provides an example where weighting based on subjective judgment can significantly improve forecasting performance.

References

- Amisano, G. and Geweke, J. (2017). Prediction Using Several Macroeconomic Models. *Review of Economics and Statistics*, 99(5):912–925.
- Athanasopoulos, G., Ahmed, R. A., and Hyndman, R. J. (2009). Hierarchical Forecasts for Australian Domestic Tourism. *International Journal of Forecasting*, 25(1):146–166.
- Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., and Petropoulos, F. (2017). Forecasting with Temporal Hierarchies. *European Journal of Operational Research*, 262(1):60–74.
- Bergmeir, C., Hyndman, R. J., and Benítez, J. M. (2016). Bagging exponential smoothing methods using STL decomposition and Box-Cox transformation. *International Journal of Forecasting*, 32(2):303–312.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263.
- Farebrother, R. W. (1978). Partitioned Ridge Regression. *Technometrics*, 20(2):121–122.
- Gross, C. W. and Sohl, J. E. (1990). Disaggregation Methods to Expedite Product Line Forecasting. *Journal of Forecasting*, 9(3):233–254.
- Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., and Shang, H. L. (2011). Optimal Combination Forecasts for Hierarchical Time Series. *Computational Statistics and Data Analysis*, 55(9):2579–2589.
- Hyndman, R. J. and Athanasopoulos, G. (2018). Forecasting Hierarchical or Grouped Time Series. In *Forecasting: Principles and Practice*.
- Hyndman, R. J. and Khandakar, Y. (2008). Automatic Time Series Forecasting: The forecast Package for R. *Journal of Statistical Software*, 27(3):1–22.
- Hyndman, R. J., Lee, A. J., and Wang, E. (2016). Fast computation of reconciled forecasts for hierarchical and grouped time series. *Computational Statistics and Data Analysis*, 97:16–32.
- Jeon, J., Panagiotelis, A., and Petropoulos, F. (2018). Reconciliation of probabilistic forecasts with an application to wind power.
- Kang, Y., Hyndman, R. J., and Smith-Miles, K. (2017). Visualising forecasting algorithm performance using time series instance spaces. *International Journal of Forecasting*, 33(2):345–358.
- Kapetanios, G., Mitchell, J., Price, S., and Fawcett, N. (2015). Generalised density forecast combinations. *Journal of Econometrics*, 188(1):150–165.

- Mustonen, S. (1997). A Measure for Total Variability in Multivariate Normal Distribution. *Computational Statistics and Data Analysis*, 23(3):321–334.
- Percy, D. F. (1992). Prediction for Seemingly Unrelated Regressions. *Journal of the Royal Statistical Society*, 54(1):243–252.
- Wickramasuriya, S. L., Athanasopoulos, G., and Hyndman, R. J. (2018). Forecasting hierarchical and grouped time series through trace minimization. *Journal of the American Statistical Association*, 114(526):804–819.
- Zellner, A. (1962). An Efficient Method of Estimating Seemingly Unrelated Regressions and Tests for Aggregation Bias. *Journal of the American Statistical Association*, 57(298):348–368.

A Appendix

A.1 Robustness




Figure 12: Accuracy of Forecasting Methods at the Top Level.




Figure 13: Accuracy of Forecasting Methods at the Bottom Level.




Figure 14: Forecast Errors of Reconciliation Methods. Lower bars indicate better forecasts. Horizontal lines show the mean absolute scaled error of unreconciled predictions. Average of all forecast horizons.




Figure 15: Forecast Errors of Reconciliation Methods. Lower bars indicate better forecasts. Horizontal lines show the mean absolute percentage error of unreconciled predictions. Average of all forecast horizons.

A.2 Data

The data is compiled by the Swiss Federal Customs Administration⁴.

Table 1: Description of Categorical Hierarchy

No.	Description
01	Forestry and agricultural products, fisheries
01.1	Food, beverages and tobacco
01.2	Feeding stuffs for animals
01.3	Live animals
01.4	Horticultural products
01.5	Forestry products (not firewood)
01.6	Products for commercial/industrial further processing such as oils, fats, starches, plants and vegetable parts, etc.
02	Energy source
02.1	Solid combustibles
02.2	Petroleum and distillates
02.3	Gas
02.4	Electrical energy
03	Textiles, clothing, shoes
03.1	Textiles
03.2	Articles of apparel and clothing
03.3	Shoes, parts and accessories
04	Paper, articles of paper and products of the printing industry
04.1	Basic materials for paper production, such as cellulose and cellulose fibre and paper and carton waste
04.2	Paper and carton in rolls, strips or sheets
04.3	Goods from paper or carton
04.4	Products of the printing industry
05	Leather, rubber, plastics
05.1	Leather
05.2	Rubber
05.3	Plastics
06	Products of the chemical and pharmaceutical industry
06.1	Chemical raw materials, basic materials and unformed plastics
06.2	Chemical end products, vitamins, diagnostic products, including active substances
07	Stones and earth
07.1	Mineral raw materials and basic products
07.2	Goods from stone and cement
07.3	Ceramic wares
07.4	Glass
08	Metals
08.1	Iron and steel
08.2	Non-ferrous metals

continued ...

⁴<https://www.ezv.admin.ch/ezv/en/home/topics/swiss-foreign-trade-statistics.html>

... continued

08.3	Metal goods
09	Machines, appliances, electronics
09.1	Industrial machinery
09.2	Agricultural machines
09.3	Household appliances
09.4	Office machines
09.5	Electrical and electronic industry appliances and devices
09.6	Military equipment
10	Vehicles
10.1	Road vehicles
10.2	Railed vehicles
10.3	Air- and spacecraft
10.4	Watercraft
11	Precision instruments, clocks and watches and jewellery
11.1	Precision instruments and equipment
11.2	Watches
11.3	Jewellery and household goods made from precious metals
12	Various goods such as music instruments, home furnishings, toys, sports equipment, etc.
12.1	Exposed film
12.2	Music instruments
12.3	Home furnishings
12.4	Toys and sports equipment
12.5	Stationery goods
12.6	Various goods such as umbrellas, neon signs, festive articles, brushes, lighters, pipes, etc.
13	Precious metals, precious and semi-precious stones
13.1	Precious and semi-precious stones
13.2	Precious metals (including gold and silver bars from 1.1.2012)
14	Works of art and antiques
14.1	Works of art
14.2	Antiques and collectors' items

Table 2: Description of Geographical Hierarchy

Country	valid from	valid to
Europe		
DE Germany	01/1988	-
FR France	01/1988	-
IT Italy	01/1988	-
NL Netherlands	01/1988	-
BE Belgium-Luxembourg	01/1988	12/1998
BE Belgium	01/1999	-
LU Luxembourg	01/1999	-
AT Austria	01/1988	-

continued ...

... continued

GB	United Kingdom	01/1988	-
DK	Denmark	01/1988	-
NO	Norway	01/1988	-
SE	Sweden	01/1988	-
PT	Portugal	01/1988	-
FI	Finland	01/1988	-
HR	Croatia, Republic of	02/1992	-
SI	Slovenia	02/1992	-
BA	Bosnia and Herzegovina	05/1992	-
MK	North Macedonia	05/1992	-
ME	Montenegro	05/1992	12/1996
ME	Montenegro	01/2007	-
XM	Montenegro	01/2006	12/2006
SQ	Serbia	05/1992	12/1996
RS	Serbia	01/2007	-
XS	Serbia	01/2006	12/2006
YU	Federal Republic of Yugoslavia	01/1997	12/2003
CS	Serbia and Montenegro	01/2004	12/2005
XK	Kosovo	01/2006	-
IS	Iceland	01/1988	-
IE	Ireland	01/1988	-
ES	Spain	01/1988	-
GR	Greece	01/1988	-
TR	Turkey	01/1988	-
DD	GDR	01/1988	10/1990
PL	Poland	01/1988	-
CZ	Czech Republic	01/1993	-
CS	Czechoslovakia	01/1988	02/1992
SK	Slovakia	01/1993	-
HU	Hungary	01/1988	-
AL	Albania	01/1988	-
BG	Bulgaria, Republic of	01/1988	-
RO	Romania	01/1988	-
SU	USSR	01/1988	12/1991
YU	Yugoslavia	01/1988	04/1992
CY	Cyprus	01/1988	-
SJ	Svalbard and Jan Mayen Island	01/1999	-
MT	Malta	01/1988	-
GI	Gibraltar	01/1988	-
FO	Faeroe Islands	01/1988	-
SM	San Marino	01/1999	-
VA	Holy See	01/1999	-
AD	Andorra	01/1988	-
EE	Estonia	01/1992	-
LV	Latvia	01/1992	-
LT	Lithuania	01/1992	-
QX	Countries not specified	01/2002	-

Central Asia

continued ...

... continued

RU	Russian Federation	01/1992	-
AM	Armenia	01/1992	-
AZ	Azerbaijan	01/1992	-
BY	Belarus	01/1992	-
GE	Georgia	01/1992	-
KZ	Kazakhstan	01/1992	-
KG	Kyrgyz, Republic	01/1992	-
MD	Moldova, Republic of	01/1992	-
TJ	Tajikistan	01/1992	-
TM	Turkmenistan	01/1992	-
UA	Ukraine	01/1992	-
UZ	Uzbekistan	01/1992	-

Africa and Middle East

EG	Egypt	01/1988	-
SD	Sudan	01/1988	-
SS	South Sudan, Republic of	09/2011	-
LY	Libya	01/1988	-
TN	Tunisia	01/1988	-
DZ	Algeria	01/1988	-
XA	Canary Islands	01/1988	-
MA	Morocco	01/1988	-
EH	Western Sahara	01/1999	-
XB	Ceuta and Melilla	01/1988	12/2010
GQ	Equatorial Guinea	01/1988	-
XC	Ceuta	01/2001	-
XL	Melilla	01/2001	-
TG	Togo	01/1988	-
SN	Senegal	01/1988	-
ML	Mali	01/1988	-
MR	Mauritania	01/1988	-
CI	Côte d'Ivoire	01/1988	-
BF	Burkina Faso	01/1988	-
BJ	Benin	01/1988	-
NE	Niger	01/1988	-
GN	Guinea	01/1988	-
GM	Gambia	01/1988	-
SL	Sierra Leone	01/1988	-
LR	Liberia	01/1988	-
GH	Ghana	01/1988	-
NG	Nigeria, Federal Republic of	01/1988	-
CM	Cameroon	01/1988	-
GA	Gabon	01/1988	-
CG	Congo, Republic of the	01/1988	-
CF	Central African Republic	01/1988	-
TD	Chad	01/1988	-
CD	Congo, Democratic Republic of the	06/1997	-
ZR	Zaire	01/1988	05/1997
AO	Angola	01/1988	-

continued ...

... continued

GW	Guinea-Bissau	01/1988	-
BW	Botswana	01/1988	-
CV	Cabo Verde, Republic of	01/1988	-
LS	Lesotho	01/1988	-
ST	Sao Tomé and Príncipe	01/1988	-
NA	Namibia	01/1988	-
ZA	South Africa	01/1988	-
SZ	Swaziland	01/1988	-
ZM	Zambia	01/1988	-
ZW	Zimbabwe	01/1988	-
MW	Malawi	01/1988	-
MZ	Mozambique	01/1988	-
MG	Madagascar, Republic of	01/1988	-
RE	RÉunion	01/1988	-
SH	St Helena, Ascen. and Tristan da Cunha	01/1988	-
KM	Comoros, Union of	01/1988	-
AQ	Antarctica	01/1988	-
MU	Mauritius	01/1988	-
IO	British Indian Ocean Territory	01/1988	-
TZ	Tanzania, United Republic of	01/1988	-
SC	Seychelles, Republic of	01/1988	-
RW	Rwanda	01/1988	-
BV	Bouvet Island	01/1999	-
BI	Burundi	01/1988	-
YT	Mayotte	01/1999	-
SO	Somalia, Federal Republic of	01/1988	-
TF	French Southern Territories	01/1999	-
DJ	Djibouti	01/1988	-
ER	Eritrea	01/1994	-
ET	Ethiopia, Fed. Democratic Republic of	01/1988	-
KE	Kenya	01/1988	-
UG	Uganda	01/1988	-
SY	Syrian Arab Republic	01/1988	-
LB	Lebanon	01/1988	-
IL	Israel	01/1988	-
PS	Palestine, the State of	01/1997	-
JO	Jordan	01/1988	-
SA	Saudi Arabia	01/1988	-
YE	Yemen (Nord)	01/1988	12/1990
YE	Yemen	01/1991	-
YD	Yemen (Sud)	01/1988	12/1990
QA	Qatar	01/1988	-
BH	Bahrain	01/1988	-
AE	United Arab Emirates	01/1988	-
OM	Oman	01/1988	-
KW	Kuwait	01/1988	-
IQ	Iraq	01/1988	-
IR	Iran, Islamic Republic of	01/1988	-

continued ...

... continued

South Asia			
AF	Afghanistan	01/1988	-
PK	Pakistan	01/1988	-
BD	Bangladesh	01/1988	-
IN	India	01/1988	-
LK	Sri Lanka	01/1988	-
MV	Maldives	01/1988	-
NP	Nepal, Federal Democratic Rep.	01/1988	-
BT	Bhutan	01/1988	-
<hr/>			
East Asia			
MM	Myanmar, Union of	01/1988	-
TH	Thailand	01/1988	-
MY	Malaysia	01/1988	-
BN	Brunei Darussalam	01/1988	-
SG	Singapore	01/1988	-
KH	Cambodia	01/1988	-
LA	Lao, People's Democratic Republic	01/1988	-
VN	Viet Nam, Socialist Republic of	01/1988	-
MN	Mongolia	01/1988	-
CN	China, People's Republic of	01/1988	-
HK	Hong Kong	01/1988	-
TW	Taiwan	01/1988	-
MO	Macau	01/1988	-
KP	Korea, People's Democratic Republic of	01/1988	-
KR	Korea, Republic of	01/1988	-
JP	Japan	01/1988	-
PH	Philippines	01/1988	-
ID	Indonesia	01/1988	-
TL	East Timor	01/2004	-
TP	East Timor	01/1999	12/2003
<hr/>			
North America			
CA	Canada	01/1988	-
PM	St Pierre and Miquelon	01/1988	-
US	United States	01/1988	-
GL	Greenland	01/1988	-
<hr/>			
Latin America			
MX	Mexico	01/1988	-
BZ	Belize	01/1988	-
GT	Guatemala	01/1988	-
HN	Honduras	01/1988	-
SV	El Salvador	01/1988	-
NI	Nicaragua	01/1988	-
CR	Costa Rica	01/1988	-
PA	Panama	01/1988	-
KY	Cayman Islands	01/1988	-
TC	Turks and Caicos Islands	01/1988	-
BS	Bahamas	01/1988	-

continued ...

... continued

BM	Bermuda	01/1988	-
JM	Jamaica	01/1988	-
CU	Cuba	01/1988	-
HT	Haiti	01/1988	-
DO	Dominican Republic	01/1988	-
VI	American Virgin Islands	01/1988	-
PR	Puerto Rico	01/1988	12/2005
DM	Dominica	01/1988	-
VC	St Vincent and the Grenadines	01/1988	-
LC	St Lucia	01/1988	-
MS	Montserrat	01/1988	-
AG	Antigua and Barbuda	01/1988	-
BB	Barbados	01/1988	-
GD	Grenada	01/1988	-
KN	St Kitts and Nevis	01/1988	-
AI	Anguilla	01/1988	-
GP	Guadeloupe	01/1988	-
VG	British Virgin Islands	01/1999	-
MQ	Martinique	01/1988	-
TT	Trinidad and Tobago	01/1988	-
BL	Saint BarthÉlemy	01/2013	-
AN	Netherlands Antilles	01/1988	12/2012
AW	Aruba	01/1999	-
BQ	Bonaire, Sint Eustatius and Saba	01/2013	-
CW	Curacao	01/2013	-
SX	Sint Maarten (NL)	01/2013	-
CO	Colombia	01/1988	-
VE	Venezuela, the Bolivarian Republic of	01/1988	-
GY	Guyana	01/1988	-
SR	Suriname	01/1988	-
GF	French Guiana	01/1988	-
BR	Brazil	01/1988	-
PY	Paraguay	01/1988	-
UY	Uruguay	01/1988	-
AR	Argentina	01/1988	-
FK	Falkland Islands	01/1988	-
GS	South Georgia and South Sandwich Islands	01/1999	-
CL	Chile	01/1988	-
BO	Bolivia, the Plurinational State of	01/1988	-
PE	Peru	01/1988	-
EC	Ecuador	01/1988	-

Australia and Oceania			
AU	Australia	01/1988	-
PG	Papua New Guinea	01/1988	-
CC	Cocos (Keeling) Islands	01/1999	-
HM	Heard and McDonald Islands	01/1999	-
NF	Norfolk Island	01/1999	-
CX	Christmas Island	01/1999	-

continued ...

... continued

NZ	New Zealand	01/1988	-
CK	Cook Islands	01/1998	-
WS	Samoa	01/1988	-
NU	Niue Island	01/1999	-
KI	Kiribati, the Republic of	01/1988	-
TK	Tokelau Islands	01/1999	-
TV	Tuvalu	01/1988	-
PN	Pitcairn Islands	01/1988	-
SB	Solomon Islands	01/1988	-
PF	French Polynesia	01/1988	-
NC	New Caledonia	01/1999	-
WF	Wallis and Futuna	01/1999	-
PU	American Oceania	01/1988	12/1996
UM	American Oceania	01/2006	-
UM	American Oceania	01/1997	12/2005
MP	Northern Mariana, Islands	01/1997	-
MH	Marshall Islands	01/1997	-
FM	Micronesia, Federated States of	01/1997	-
PW	Palau	01/1997	-
FJ	Fiji, Republic of	01/1988	-
AS	American Samoa	01/2006	-
GU	Guam	01/2006	-
VU	Vanuatu	01/1988	-
NR	Nauru	01/1988	-
TO	Tonga	01/1988	-
