CPB Netherlands Bureau for
Economic Policy Analysis

Forecasting wage growth and price inflation in the Netherlands with a BVAR model

We developed a BVAR model for hourly wage growth and HICP price inflation to help reduce the forecast error of these variables. We find that a BVAR model can be a useful additional tool support the CPB forecast of wage and price developments in the Netherlands.

Predictions of HICP inflation have a smaller error than predictions of hourly wage growth.

CPB Background Document

Jurriaan Paans, Marente Vlekke

August 2020

Contents

1 Introduction

During the past few years, forecasting wage and price developments in the Netherlands has become more challenging. At the CPB Netherlands Bureau for Economic Policy Analysis, the macroeconometric model SAFFIER[1] produces the main forecast of wage and price inflation. These forecasts serve as the baseline forecast. Using additional analyses, the wage and price experts make regular adjustments to the baseline forecasts, and take into account developments that are not necessarily captured by macroeconometric models. Since the great recession multiple macroeconomic relationships appear to have changed: inflation and wage inflation have been subdued and appear to respond with greater delay to the business cycle, the Phillips-curve appears to have weakened, productivity growth has declined, and energy prices have increasingly been decoupled from the oil price. As a result, forecasting wage and price inflation has become more challenging.

Therefore, the wage and price experts at the CPB are developing additional tools to support the forecasting process, such as Bayesian Vector Autoregression (BVAR) models. Using different types of models can help improve the forecast, as each model has its own strengths and weaknesses. One approach complementary to a macroeconometric model such as SAFFIER is to apply Bayesian VAR (BVAR) models to forecast wage and price developments. At the CPB, additional models such as BVARs are increasingly used to inform the forecasts, for example for GDP (De Wind, 2015) and unemployment (Adema et al., 2018). This approach is also common at other institutes involved in forecasting, such as central banks[2]. VAR models capture the linear interdependencies among multiple variables: each variable in the model has an equation which relates the evolution of that variable to its own lagged values, and the lagged values of the other variables. VAR models do not require much prior knowledge about the specific characteristics of the economic relationships between variables, which is convenient if these characteristics are unknown or uncertain. That is however also a drawback if one seeks to explain (rather than describe) changes in the economy. VAR models are therefore complementary to macroeconometric models such as SAFFIER, of which the equations are derived from economic theory. We use Bayesian techniques to estimate our VAR models, hence the acronym 'BVAR'. Bayesian approaches have a number of advantages compared to frequentist approaches, which we will delineate below.

The BVAR presented here is relatively parsimonious, easy to use and quick to estimate. The BVAR we use has an analytical solution (we do not require MCMC sampling methods[3]), so they are quick to estimate and easy to use. These are essential requirements for the models we use for the regular forecasting exercise of the CPB, which occurs four times per year. The final BVAR specification contains 29 variables, a model size that is shown to be tractable while delivering a good forecasting performance (Banbura et. al, 2008).

We find that BVARs can be useful additional tools for forecasting hourly wage growth and HICP inflation. BVARs outperform simpler time series models, such as AR(1) models. The forecasting performance of the BVAR for HICP inflation and the CPB forecast in the 2016 – 2019 period are roughly comparable. The BVAR for hourly wage growth outperforms the CPB forecast and other simple models in the 2016 – 2019 period, revealing the merits of using a time series approach to model dynamics in wage inflation.

[1] For relevant documentation on SAFFIER, see Verbruggen, Kranendonk en Smid (2010).
[2] See e.g. the BEAR toolbox used at the ECB (link).
[3] Markov Chain Monte Carlo (MCMC) sampling methods are algorithms that can be used to numerically approximate probability distributions of parameters if the distributions cannot be derived analytically. These approaches usually are time-consuming compared to Bayesian methods with analytical solutions.

2 Forecasting with BVAR models

2.1 The challenge of forecasting wage and price dynamics in the Netherlands

In the past few years, hourly wage growth was mostly overestimated, while HICP inflation was mostly underestimated; moreover predicting wage inflation was more prone to error than predicting inflation. At the CPB we make regular forecasts of HICP and hourly wage growth. Wages and prices tend to co-move: as prices of goods and services increase, so does the price of labor, i.e. wages. Therefore, modelling these two variables together makes both theoretical and empirical sense. However, while dynamics in these two variables are positively correlated, predicting these two variables poses different challenges. Figure 2.1 shows the forecast error of hourly wage growth per forecasting round for the years 2016-2019. The figure reveals that in the past few years hourly wage growth was consistently overestimated, often by more than 1%-point. By contrast, during this period HICP inflation was mostly underestimated, see figure 2.2. The forecast error of inflation is also substantially smaller than that of wage inflation: hourly wage growth was less easy to predict than HICP inflation.

Figure 2.1 Forecast error[4] for hourly wage growth per forecasting round, 2016-2019

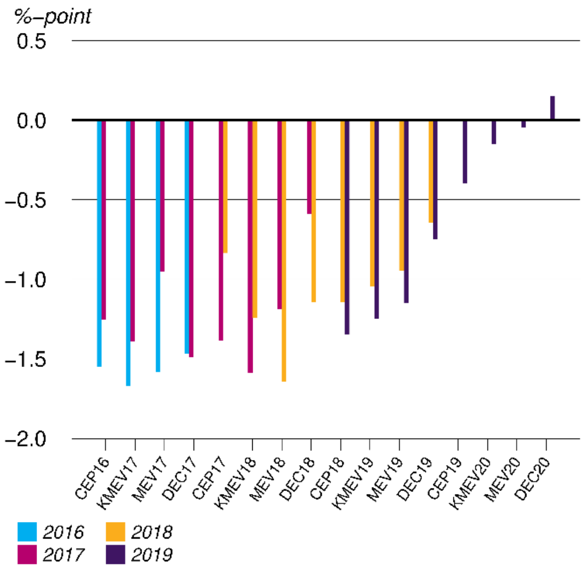

[4] The forecast error is computed as realization – forecast. The CPB publishes its main economic forecast four times per year, in March ('CEP'), June ('KMEV'), in August and September ('MEV') and in December ('DEC').

Figure 2.2 Forecast error[5] for HICP inflation per forecasting round, 2016-2019

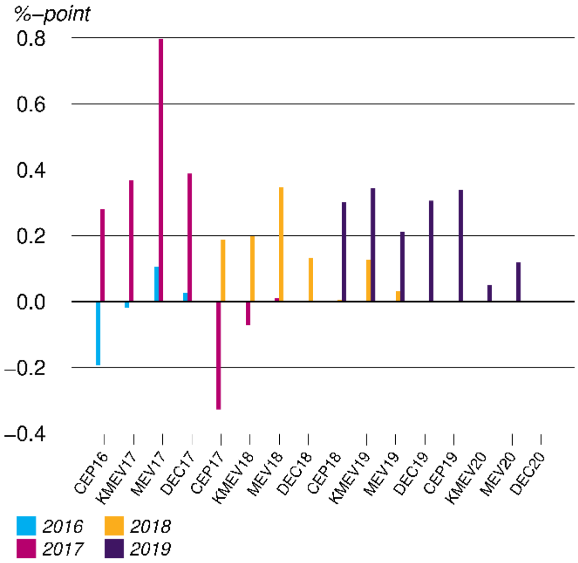

A BVAR model can be used as an additional tool to forecast wage drift inflation. Hourly gross wage growth can be decomposed into negotiated wage growth and a residual called the wage drift. The wage drift includes the growth of additional salary components, such as bonusses in addition to the agreed upon contract wages, and changes in the labor market composition. In recent years we found that wage drift is difficult to predict. This is shown in figure 2.3, which decomposes the root mean squared forecast error (RMSFE) of hourly wage growth into the RMSFE of negotiated wage growth and of wage drift for the 2016-2019 period. Annual average negotiated wage growth is relatively predictable, since negotiated wage increases for year *t* are usually announced in year *t-1* or earlier. Therefore, most of the forecast error is due to the unpredictable dynamics in wage drift. In general the wage drift is expected to be procyclical. However surprisingly, in recent years (2014-2018) the wage drift has decreased despite a tightening labor market. This has contributed to a higher forecast error during this period. Paans and Volkerink (2020) find that this is partly the result of a large inflow of low paid workers in the labor market, which suppresses the average wage and therefore the wage drift. A BVAR may be a useful additional tool in capturing and hence predicting these dynamics.

[5] The forecast error is computed as realization – forecast. The CPB publishes its main economic forecast four times per year, in March (‘CEP’), June (‘KMEV’), in August and September (‘MEV’) and in December (‘DEC’).

Figure 2.3 Decomposition of the RMFSE of hourly wage per forecasting round[6]

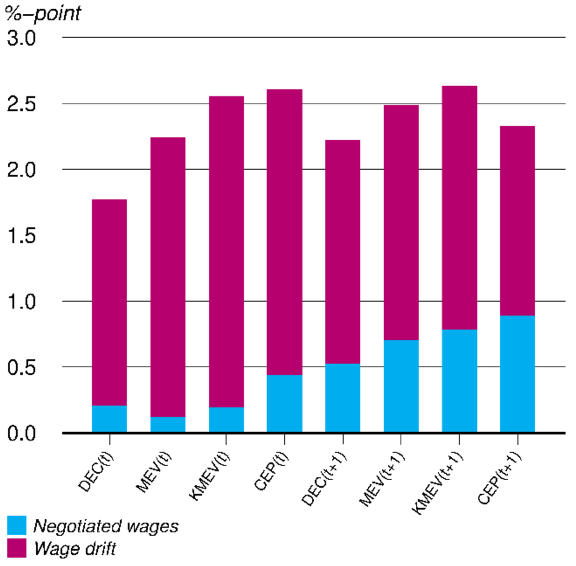

2.2 Using a BVAR model to forecast wage growth and price inflation

Vector Autoregression (VAR) models can contribute to reducing forecast error. VAR models use a simple autoregressive specification to model the (reduced-form) interdependencies between a relatively large number of variables. VARs are standard models of the forecasting toolkit, and generally produce good forecasting results. A VAR model with *n* variables and *p* lags describes the evolution of a variable to its own lagged values, and the lagged values of the other variables in the data set, that is

$$y_t = c + \sum_{i=1}^{p} B_i y_{t-i} + u_t ,$$

where y_t is an $n \times 1$ vector of observed endogenous variables, c is an $n \times 1$ vector of intercepts, B_i is an $n \times n$ matrix of autoregressive parameters, and u_t is an $n \times 1$ vector of disturbance terms or 'noise', which are assumed to be normally distributed, i.e. $u_t \sim N(0,\Sigma)$.

The main advantage of the VAR approach is that these models capture the linear interdependencies among multiple variables without requiring much prior knowledge about the specific characteristics of the economic relationships between variables. This is convenient if these characteristics are unknown or uncertain. That is however also a drawback if one seeks to explain, rather than describe, changing trends in the economy. VAR models are therefore complementary to macroeconometric models such as SAFFIER, of which the equations are derived from economic theory.

A potential drawback of VARs is that they have a large number of parameters that have to be estimated, while generally macro-economic time series have a relatively short length. This can cause 'overfitting', i.e. a situation in which the in-sample forecast of the optimized model is good, but the out-of-sample forecast is

[6] The CPB publishes its main economic forecast four times per year, in March ('CEP'), June ('KMEV'), in August and September ('MEV') and in December ('DEC').

poor. There is an inherent trade-off between overfitting and forecasting bias. A parsimonious model will likely not suffer from overfitting, but may exclude important explanatory variables. A large model on the other hand will fit in such a way that its in-sample-forecast is adequate, but the out-of-sample forecast performance is poor.

Through the use of a *prior*, a Bayesian approach to estimate VAR models can reduce the problem of overfitting. Pioneered by Litterman (1979) and Sims (1980) and Doan, Litterman and Sims (1984), Bayesian estimation techniques make use of a *prior*, which refers to parameters in the model that capture the researcher's knowledge and assumptions of the relevant economic relationships and characteristics that are not necessarily captured by the data. These parameters are pre-specified by the researcher. In Bayesian analysis the *prior* is combined with the *likelihood function*, i.e. the joint probability of the sample data as a function of the model's parameter values. The prior can be used to restrict the possibility that the estimated or *posterior* distributions of the parameter values take on economically implausible values.

We combine three types of priors: the Minnesota prior, the sum-of-coefficients prior and the dummy-initial-observation prior. We follow the literature on improving forecasting performance of BVARs for macroeconomic time series by relying on the so-called *combination prior* (see e.g. Sims and Zha (1998) and Giannone, Lenza and Primiceri (2013)). As the name suggests, the combination prior combines three priors for B_i, the *Minnesota prior*, the *sum-of-coefficients prior* and the *dummy-initial-observation prior*. Here, we briefly describe the goal of these priors, for a detailed description of the application to our BVAR model we refer to De Wind (2015). The Minnesota prior (first introduced by Litterman 1979) is used to 'shrink' the parameters in B_i to a random walk (possibly with drift) in case of level data, and essentially to white noise in case the variables are defined as growth rates, with stronger shrinkage for coefficients on longer lags and across variables. This prior is useful as a benchmark, because univariate random walk models are typically good in forecasting macroeconomic time series data. Naturally, the posterior distributions of the parameters may allow for a more complicated process if there is sufficient information in the data. This prior hence reduces the risks of overparameterization and overfitting. The sum-of-coefficients prior and the dummy-initial-observation prior are both used to reduce the impact of deterministic components (i.e. the impact of the initial conditions of the VAR process) on the forecasting performance (Doan, Litterman, and Sims (1984)).

For each of these priors we specify hyperparameters that determine how much impact the prior has on the posterior distributions. In the Bayesian economic literature this is referred to as the 'tightness' of the prior. The tightness is controlled by so-called hyperparameters, of which the values, specified beforehand by the researcher, determines how informative the prior is relative to the likelihood function for the posterior distributions of the parameters.

Model training and evaluation

To optimize the model parameters and hyperparameters, we apply statistical learning techniques by splitting the sample into a training, validation and test set. The literature on BVARs for macroeconomic forecasting proposes default values for these hyperparameters which have been shown to produce good forecasting results, see Sims and Zha (1998). However, these values may not necessarily be suitable for modelling wage and price dynamics in the Netherlands. To examine this, we optimize the prior hyperparameters based on a forecasting competition. For this we follow a common statistical learning procedure. These are useful in selecting the optimal number of lags, and the values of the hyperparameters in our model while reducing the risk of overfitting, see also James et al. (2013). This approach consists of three steps:

1. *Training*: we first estimate different versions of the model (i.e. different lags and prior hyperparameter values) on a data set[7] starting in 2001Q4 and ending in 2011Q2. This is the 'training' sample.
2. *Cross-validation*:
   i. we use the fitted models to produce out-of-sample forecasts of wage and price inflation in the validation test set, which consists of the period 2011Q3-2015Q4.
   ii. We then incorporate one extra period in the training set and re-estimate the model and use this slightly larger data set to produce another out-of-sample forecast on an expanded validation set.
   iii. We continue this process until the end of the validation period.
   iv. We choose the model which has the lowest root mean squared forecast error (RMSFE) for a one-step ahead forecast of our variable of interest in the validation period. The RMSFE is computed as

   $$\sqrt{\frac{1}{m-n}\sum_{t=n}^{m}(y_t - \hat{y}_t)^2},$$

   where for the period starting at $t = n$ and ending at $t = m$, we take the average of the squared differences between the forecast and the actual value of our variable at interest, after which we take the mean.
3. *Out-of-sample testing*: the optimal model specification results are then used to produce out-of-sample forecasts for the test set, which is the 2016-2019 period. Similar to the cross-validation, the model is evaluated in the test period using a rolling window. In the test period forecast error is calculated for each horizon (e.g. 1 month ahead, 3 months ahead etc.).

We optimize the model's parameters separately for wage and price inflation, resulting in separate models for these two variables.

3 Data

We use both a monthly and quarterly data set of seasonally adjusted variables to forecast HICP inflation and hourly wage growth. Gross wage data is only available at the quarterly level, while HICP inflation is a monthly variable. Therefore, we estimate the BVAR both on a monthly and quarterly basis. The monthly BVAR is only evaluated for HICP inflation, while we evaluate the quarterly BVAR for both HICP inflation and hourly wage growth. For the monthly data set, we impute[8] the quarterly variables, and for the quarterly data set we take three-month-averages of the monthly variables. All variables are seasonally adjusted.

We take growth rates of all variables which are not already rates. There is no consensus in the literature as to whether BVARs should be estimated in growth rates or in levels. While levels are more commonly used, Carriero (2015) among others finds that for PCE inflation, first differences result in slightly smaller errors than level data. Importantly for our goal of developing an easily maintainable tool for forecasting, there is a practical advantage to using growth rates compared to level data. Using growth rates means that we can handle changes in the base year of indices in our data set more easily. Therefore, we take growth rates of all variables except for those that are not already rates (such as interest rate data).

We include 29 variables as well as policy dummy variables in our BVARs. Here we follow the literature and include key macro-economic indicators, monetary variables, consumer, business and inflation expectations, energy prices and labor market characteristics. In addition, the model includes policy shock dummy variables

[7] We start our analysis in 2001 because of the data availability of some variables of interest.
[8] For imputation we use a Piecewise Cubic Hermite Interpolating Polynomial (PCHIP).

for VAT-increases and wage freezes. An overview of all the variables used is shown in table A1 in the Appendix. We tested different model sizes and found that overall the BVAR with these variables delivered the most reliable forecasting results. A complete data set is available from 2001Q4 onward. The data are obtained from Statistics Netherlands and Datastream.

4 Forecasting Results

4.1 HICP Inflation

Table 4.1 compares the RMSFE of HICP inflation for the BVAR and a simple AR(1) specification for the test period of 2016-2019. We show the RMSFEs for both the default and optimized prior and lags. The table also shows the results for a simple autoregressive model of order 1 or 'AR(1)-with drift' model, in which the dependent variable is modelled as a function of its lagged value times a parameter, an intercept and a disturbance term. The table displays the RMSFE for different forecasting horizons, including 1 month, 3 months (or 1 quarter), 6 months (2 quarters), 12 months (4 quarters) and 24 months (8 quarters) ahead. The forecasting results are computed with an expanding window, meaning that for an horizon *h* and period *t*, the forecast *t+h* is based on all the available data up to and including data from period *t*. We show the results for the model based on both monthly and quarterly data.

Table 4.1 RMSFE for HICP inflation for different horizons, test sample (2016-2019)

	1 month	3 months/ 1 quarter	6 months/ 2 quarters	12 months/ 4 quarters	24 months/ 8 quarters
Monthly BVAR*					
optimal lags and priors	0.202	0.187	0.176	0.178	0.186
default lags and priors	0.207	0.184	0.173	0.181	0.186
AR(1) with drift	0.210	0.189	0.179	0.182	0.194
Quarterly BVAR**					
optimal lags and priors	-	0.251	0.196	0.199	0.176
default lags and priors	-	0.230	0.189	0.191	0.200
AR(1) with drift	-	0.240	0.221	0.228	0.245

* Both the default and optimal number of lags is 12 in the monthly BVAR.
** In the quarterly BVAR the default number of lags is 4 and the optimal number of lags is 3.

The difference in forecasting performance between the BVAR and AR(1) - with drift models is small. For the test period, the RMSFE for the monthly and quarterly models hovers around 0.2, indicating that on average the difference between the predicted and realized HICP month-on-month inflation is around 0.2 percentage points. The differences in RMSFE between the optimal and default lags and priors is negligible. Particularly for the quarterly model, the RMSFE of the optimal BVAR is higher than for the BVAR with the default settings. This could be an indication of overfitting in the test period. It hence appears to be the case that optimizing the prior settings and lags has a relatively limited impact on the forecasting performance. Overall, the monthly models outperform the quarterly models. The most likely reason for this is that forecast of the monthly BVAR is based on more granular data. The differences between the RMSFEs of the BVARs and the AR(1) - with drift model is also small, suggesting that the first lag of HICP inflation adds significant value to the forecasting performance of the model.

One counterintuitive result is that for the test period, the forecasting performance of the BVAR-model increases as the forecasting horizon increases. An advantage of our validation method is that it reduces the probability of overfitting. A drawback that arises from the short series is that to sufficiently train and validate the model, we are left with a short testing period. This implies that outliers or volatility in the test period can have a large influence on the test results. We found that this is indeed what is causing odd results in the RMSFEs for different horizons. Overall the BVAR predictions are better for 2018 and 2019 compared to 2017 and 2016 (see the RMSFEs per year in the tables in the Appendix). The RMSFE in 2018 is particularly low. In the test period sample of 2016-2019, the year 2018 dominates the RMSFEs of larger horizons, resulting in a decreasing RMSFE as the horizon increases.

We also computed the RMSFE for the entire post-training sample period and found that these results are more regular. Note that because in this table the RMSFEs are in part computed with observations from the validation period, this table does not allow us to properly evaluate the added benefit of using the optimized priors and lags. For this, we can still use table 4.1.

Table 4.2 RMSFE for HICP inflation for different horizons, post-training sample (2012-2019)

	1 month	3 months/ 1 quarter	6 months/ 2 quarters	12 months/ 4 quarters	24 months/ 8 quarters
Monthly BVAR*					
optimal lags and priors	0.183	0.185	0.191	0.191	0.185
default lags and priors	0.191	0.185	0.199	0.220	0.229
AR(1) with drift	0.193	0.193	0.197	0.202	0.204
Quarterly BVAR**					
optimal lags and priors	-	0.296	0.304	0.387	0.441
default lags and priors	-	0.287	0.189	0.392	0.430
AR(1) with drift	-	0.306	0.343	0.365	0.365

* The default and optimal number of lags is 12 in the monthly BVAR.
** In the quarterly BVAR the default number of lags is 4 and the optimal number of lags is 3.

4.2 Hourly wage growth

Quarterly data publications necessarily lag behind the publication of monthly data. As table 4.3 shows, the wage data we use for the forecast of hourly wage growth even lags two quarters behind. Therefore, forecasting gross wages at the CPB often involves predicting both the future and the present. Predicting the present is commonly called ‘nowcasting’.

Table 4.3 Publication months and data availability of hourly gross wages

	Data available
CEP (March forecast)	Q3 (t-1)
KMEV (June forecast)	Q4 (t-1)
MEV (August/September forecast)	Q1
DEC (December forecast)	Q2

We can make an *unconditional* and *conditional* two-quarters-ahead nowcast to predict hourly wage growth. An *unconditional* forecast is the standard approach to forecasting, and was also used in section 4.1. For an un unconditional prediction we restrict the dataset to the last period for which a complete dataset is available, even though for some variables later periods would be available. In tables 4.4 and 4.5 these are the *unconditional* predictions.

We can use the Kalman filter-smoother to make a conditional two-quarter-ahead nowcast. Ideally, we would like to use all available data to make the best possible forecast or nowcast. Fortunately, we can use the Kalman-filter smoother (KFS) to impute missing values of quarterly variables *conditional* on monthly data that is already available. In tables 4.4 and 4.5 these are the *conditional* predictions. For a general discussion on the KFS we refer to Durbin and Koopman (2012), and for a description of its application to the VAR models used here we refer to De Wind (2015). The KFS can be used for the recursive estimation of unobserved values in a system of equations. To apply the KFS to a BVAR, we rewrite the VAR in the so-called 'state space' form, which essentially consists of two equations: a measurement equation relating the observed variables to unobserved components, and a state equation describing the dynamics of the unobserved component. The resulting imputation in a system with both observed and unobserved values is hence a conditional forecast or 'conditional nowcast'. We evaluate the potential merits of this approach by applying the KFS to two types of incomplete data sets: one in which the quarterly variables lag behind one quarter to the monthly variables, and one in which the lag is two quarters, simulating the actual data availability of gross wages per hour.

The conditional predictions obtained with the Kalman filter-smoother result in a substantially smaller RMSFE compared to the unconditional predictions. The unconditional results show that hourly wage growth is harder to predict than HICP inflation. The unconditional RMSFEs vary between 0.4 and 0.5 in the test sample. They are roughly two times larger than those for HICP inflation forecasts[9]. By applying the KFS, we can reduce the one- and two-quarters-ahead RMFSE with around 35% to 50%. Note that we do not show the conditional results for four or eight-quarter-ahead conditional predictions, because in practice wage data never lags behind more than two quarters. Note also that if we want to forecast a four quarters ahead or more, we also have to rely on unconditional forecasts of hourly wage growth. The tables below show that this results in a much larger RMSFE for hourly wage growth than for HICP inflation. Hence, despite the fact that we can reduce the prediction error in a nowcast exercise, forecasting wage inflation remains more challenging than HICP inflation.

Table 4.4 RMSFE for hourly wage growth for different horizons, test sample (2016-2019)

		1 quarter	2 quarters	4 quarters	8 quarters
BVAR*					
optimal lags and priors	unconditional	0.414	0.459	0.456	0.286
	conditional	0.266	0.289	-	-
default lags and priors	unconditional	0.410	0.456	0.545	0.440
	conditional	0.372	0.377	-	-
AR(1) with drift	unconditional	0.509	0.536	0.498	0.324

* In the quarterly BVAR the default number of lags is 4 and the optimal number of lags is 2.

[9] We found that hourly wage growth is more difficult to forecast, because total hours is more difficult to forecast. Forecasting gross wages separately results in RMSFEs that are comparable to those of HICP inflation. Since hourly wage growth is our variable of interest, we investigated whether we could improve upon the forecast by making a separate forecast for total hours and optimizing the BVAR for this variable, but this approach did not improve upon optimizing the model for hourly wage growth. [10] This comparison does not take into account the effect of different vintages of the data. The BVAR-forecasts are based on the latest vintage, while the CPB forecasts were made with the vintages available at that time. Fortunately, revisions of inflation and wage data are rare. In our sample period, there was only one revision of inflation in 2018. If we take this revision into account, our results hardly change.

Table 4.5 RMSFE for hourly wage growth for different horizons, post-training sample (2012-2019)

		1 quarter	2 quarters	4 quarters	8 quarters
BVAR*					
optimal lags and priors	unconditional	0.671	0.805	0.919	1.015
	conditional	0.417	0.395	-	-
default lags and priors	unconditional	0.819	0.873	0.934	0.976
	conditional	0.524	0.474	-	-
AR(1) with drift	unconditional	0.797	0.866	0.928	0.968

* In the quarterly BVAR the default number of lags is 4 and the optimal number of lags is 2.

Despite the relatively high RMSFEs for hourly wage growth, the BVAR models have a better forecasting performance than a simple AR(1)-with drift model. Tables 4.4 and 4.5 show that the RMSFE of the BVAR models is smaller than those of a simple AR(1)-with drift model for hourly wage growth.

As was the case for HICP inflation, we again find that the RMSFEs decrease as the forecasting horizon increases. Therefore, we also evaluated the RMSFEs for the entire post-training sample, and again we see that for this period the results are more regular.

For hourly wage growth, the RMSFEs for the entire post-training sample are higher than that of the test period. The RMSFEs for hourly wage growth are higher for the post-training than for the test sample. This appears to be a counterintuitive result, since the model's lags and priors settings were optimized in the validation period. However, we found that the model's forecast errors, shown in the Appendix, are larger in the validation period, particularly in 2014. This may explain the relatively large overall RMFSE in the validation period.

4.3 Comparing annual forecasts, 2016 - 2019

In this section we compare the forecasting performance of the BVAR and AR(1) models for *annual average* inflation with predictions from the quarterly forecast publications of the CPB. For the quarterly forecast publication, the CPB publishes the *annual average* growth of inflation of the HICP index and gross wages per hour. The annual average growth rate is computed as the percentage difference between the average index in year t and the average index in year $t-1$. Therefore, it is useful to review the performance of the BVAR and the AR(1) -with drift model on an annual basis as well. Figures 4.2 and 4.3 compare the RMSFE for the current year (year t) and the upcoming year (year $t+1$) made in the current year t, per forecast publication. The CPB publishes its main economic forecast four times per year, in March ('CEP'), June ('KMEV'), in August and September ('MEV') and in December ('DEC'), this is compared with the BVAR and AR(1) forecasts for these fixed moments. The publication dates can be found in table 4.6. Note that small differences in the month-on-month forecast error discussed in the previous two sections can accumulate rapidly when an annual forecast is made.

Table 4.6 Publication months and data availability HICP inflation

	Data available up until and including:	
	Monthly model	Quarterly model
CEP (March forecast)	January	Q4 (t-1)
KMEV (June forecast)	April	Q1
MEV (August/September forecast)	July	Q2
DEC (December forecast)	November	Q3

For HICP inflation, the monthly BVAR has a considerably smaller annual RMSFE than the quarterly BVAR and the AR(1) with drift model. In the previous section we saw that the differences in the RMSFE of the month-on-month forecasts were small. However, once these month-on-month predictions are compounded into an annual prediction, the differences in forecast error are also compounded, see figure 4.1. Hence, based on annual forecast errors, we conclude that the monthly BVAR is the optimal model.

For HICP inflation, the annual forecast performance of the monthly BVAR and the CPB are comparable during 2016-2019 forecasting rounds. We made this comparison based on the assumption that the BVAR has the same data available as the forecaster at that point in time. An overview of data availability is shown in table 4.6[10]. As is shown in figure 4.1, the annual RMSFEs of the BVAR are mostly similar those of the CPB: for some forecasting rounds the RMSFE is slightly higher, for others it is lower or about the same. As was explained in the introduction, the CPB forecast is the result of the SAFFIER prediction to which the wage and price experts make adjustments if this is deemed prudent. Given the comparable RMSFEs, this finding demonstrates the potential usefulness of a BVAR for HICP inflation as an additional tool in this process.

Figure 4.1 RMSFE for year *t* and *t*+1 per forecast round of HICP inflation in year *t*, 2016-2019

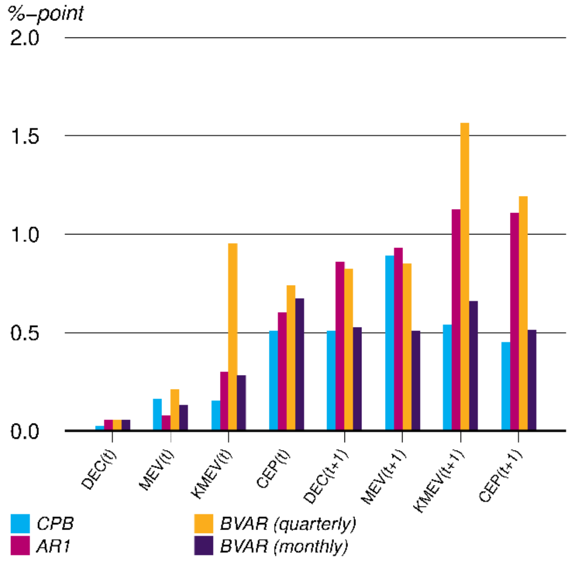

[10] This comparison does not take into account the effect of different vintages of the data. The BVAR-forecasts are based on the latest vintage, while the CPB forecasts were made with the vintages available at that time. Fortunately, revisions of inflation and wage data are rare. In our sample period, there was only one revision of inflation in 2018. If we take this revision into account, our results hardly change.

For hourly wage growth, the annualized BVAR predictions have a consistently lower RMSFE than the CPB forecasts during the 2016-2019 period. The results are shown in figure 4.2. The figure shows that the BVAR forecast outperforms the CPB forecast during the 2016-2019 period. As was mentioned above, the main challenge in predicting gross wages lies in predicting wage drift. These results suggest that a pure time series approach is a useful tool in capturing and predicting the dynamics in wage drift.

Figure 4.2 RMSFE for year *t* and *t*+1 per forecast round of hourly wage growth in year *t*, 2016-2019

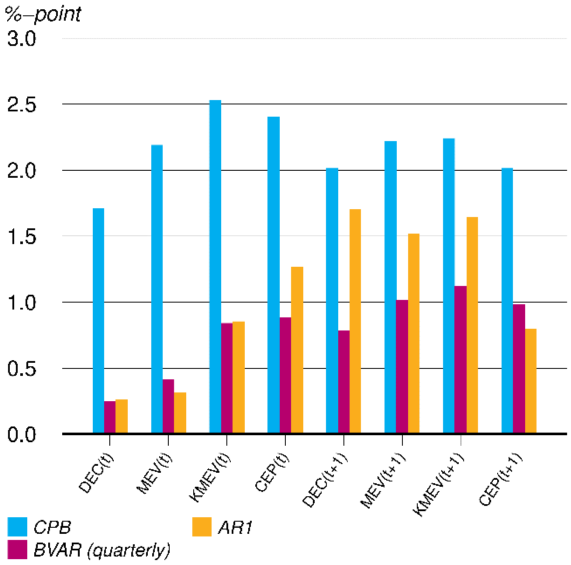

5 Conclusions

The wage and price experts at the CPB developed a BVAR model for wage and price inflation to help reduce the forecast error of these variables. We tested monthly and quarterly specifications with standard and optimized lags and priors. We find that optimizing the lags and prior hyperparameters yields a small reduction in prediction error compared to the default settings.

For HICP and wage inflation, we find that a BVAR can be a useful additional tool for the forecasting suite at the CPB. The forecasting performance of the BVAR for HICP inflation and the CPB forecast in the 2016 – 2019 period are roughly comparable. From this we conclude that the BVAR can be a useful additional tool for forecasting HICP inflation. The BVAR for hourly wage growth outperforms the CPB forecast and other simple models in the 2016 – 2019 period, revealing the merits of using a time series approach to model dynamics in wage inflation. Nevertheless, the performance is still relatively poor compared to the performance of the BVAR for HICP inflation. We found that this is mainly due to the prediction error of hours worked, and we recommend future research to explore this further.

The BVARs for wage and price inflation serve as a benchmark for new additions to the modelling suite. The findings of this study indicate several promising directions for future research. The conditional analysis showed that in forecasting a quarterly variable such as hourly wage growth in the current year is made easier if we can use monthly data that is already available. This finding also points to the potential gains from developing mixed-frequency models. Furthermore, given the volatility of the forecast error for wage inflation, we should investigate the advantages of allowing for stochastic volatility in the disturbance term. In addition, it is worthwhile evaluating whether related forecasting approaches (such as dynamic factor modelling), as well as machine learning approaches provide additional advantages. Here we have to consider the trade-off between parsimony and forecasting accuracy, and between transparency and model complexity. For example, in practical use, the effects of incoming data on changes in the predictions of wage and price inflation ought to be transparent and easy to interpret. Here, classic econometric approaches such as BVARs have proven their merits.

References

Adema, A., K. Folmer, H. van Heuvelen, S. Kuijpers, R. Luginbuhl en B. Scheer, 2018, Voorspellen van de werkloosheid: kan het beter?, CPB Background Document, 8 March 2018.

Banbura, M., D. Giannone, L. Reichlin, 2008, Large Bayesian VARs. Working paper series, 966, European Central Bank.

Carriero, A., T.E. Clark, M. Marcellino, 2015, Bayesian VARs: specification choices and forecast accuracy, *Journal of Applied Econometrics*, 30: 46-73.

Doan, T., R. Litterman, and C. Sims, 1984, Forecasting and conditional projection using realistic prior distributions, *Econometric Reviews*, 3(1): 1-100.

Durbin, J. and S. J. Koopman, 2012, *Time Series Analysis by State Space Methods: Second Edition*. Oxford University Press.

James, G., D. Witten, T. Hastie en R. Tibshirani, 2013, *An Introduction to Statistical Learning*, New York: Springer.

Litterman, R. B., 1979, Techniques of forecasting using vector autoregressions, Working Papers 115, Federal Reserve Bank of Minneapolis.

Paans, J. and M. Volkerink, 2020, De incidentele loongroei onder de loep: samenstellingseffecten en loondynamiek, CPB Background Document, July 2020.

Sims, C. A., 1980, Macroeconomics and reality, *Econometrica*, 48(1):1-48.

Sims, C. A. and Zha, T., 1998, Bayesian methods for dynamic multivariate models, *International Economic Review*, 39(4):949 - 968.

Verbruggen, J., H. Kranendonk en B. Smid, 2010, SAFFIER II: 1 model voor de Nederlandse economie, in 2 hoedanigheden, voor 3 toepassingen, CPB Document, no. 217.

Wind, J., 2015, Technical background document for BVAR models used at the CPB, CPB Background Document, February 2015.

Appendix

Table A.1 Variable overview[11]

Variable	Frequency: monthly (M) or quarterly (Q)
Δ HICP Netherlands	M
Δ HICP EU	M
Inflation expectations consumers NL	M
Δ GDP deflator	Q
Δ Hours worked	Q
Δ Gross Wage per hour	Q
Δ Household consumption	Q
Δ Government consumption	Q
Δ Real Investment	Q
Δ Real Imports	Q
Δ Real Exports	Q
Consumer confidence	M
Producer confidence	M
Business survey - personnel	M
Business survey - sales	M
Business survey - production	M
Δ Productivity	Q
Δ Capacity utilization rate *	Q
Δ Employment level (employees)	Q
Unemployment rate	M
Δ M3	M
USD/EUR exchange rate	M
10-year government NL bond yield	M
Δ Brent oil spot price	M
EONIA	M
Δ Electricity Base futures	M
Δ Number of flexible workers *	Q
Δ Number of workers aged 15-30 *	Q
Δ Workers education level *	Q
Wage freeze and VAT increase dummies[12]	M

[11] The symbol 'Δ' refers to the growth rate. Variables marked with a '*', are defined as a first difference.
[12] In the 2011 – 2014 period wages in the government have not increased as a result of policy, and VAT-increases in 2012 and 2019.

Table A.2 RMSFE for hourly wage growth and HICP inflation for different horizons, 2012

	2012				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.376	0.454	-	-
Baseline lags and default Sims priors	-	0.368	0.274	-	-
HICP inflation - quarterly model					
Optimal lags and priors	-	0.321	0.402	-	-
Baseline lags and default Sims priors	-	0.339	0.399	-	-
HICP inflation - monthly model					
Optimal lags and priors	0.176	0.202	0.212	-	-
Baseline lags and default Sims priors	0.187	0.183	0.181	-	-

Table A.3 RMSFE for hourly wage growth and HICP inflation for different horizons, 2013

	2013				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.662	0.764	0.747	-
Baseline lags and default Sims priors	-	0.787	0.668	0.608	-
HICP inflation - quarterly model					
Optimal lags and priors	-	0.411	0.408	0.456	-
Baseline lags and default Sims priors	-	0.394	0.444	0.469	-
HICP inflation - monthly model					
Optimal lags and priors	0.153	0.146	0.153	0.163	-
Baseline lags and default Sims priors	0.185	0.193	0.219	0.206	-

Table A.4 RMSFE for hourly wage growth and HICP inflation for different horizons, 2014

	2014				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	1.216	1.702	1.964	2.042
Baseline lags and default Sims priors	-	1.803	1.931	1.974	1.892
HICP inflation - quarterly model					
Optimal lags and priors	-	0.207	0.294	0.538	0.656
Baseline lags and default Sims priors	-	0.206	0.309	0.542	0.630
HICP inflation - monthly model					
Optimal lags and priors	0.135	0.137	0.146	0.148	0.163
Baseline lags and default Sims priors	0.152	0.143	0.183	0.259	0.308

Table A.5 RMSFE for hourly wage growth and HICP inflation for different horizons, 2015

	2015				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.691	0.512	0.454	0.666
Baseline lags and default Sims priors	-	0.915	0.755	0.381	0.650
HICP inflation - quarterly model					
Optimal lags and priors	-	0.334	0.471	0.466	0.647
Baseline lags and default Sims priors	-	0.394	0.428	0.485	0.636
HICP inflation - monthly model					
Optimal lags and priors	0.181	0.195	0.182	0.180	0.180
Baseline lags and default Sims priors	0.201	0.207	0.192	0.217	0.262

Table A.6 RMSFE for hourly wage growth and HICP inflation for different horizons, 2016

	2016				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.523	0.801	0.964	1.034
Baseline lags and default Sims priors	-	0.665	0.837	1.136	1.088
HICP inflation - quarterly model					
Optimal lags and priors	-	0.263	0.359	0.390	0.477
Baseline lags and default Sims priors	-	0.269	0.335	0.355	0.440
HICP inflation - monthly model					
Optimal lags and priors	0.244	0.235	0.248	0.248	0.246
Baseline lags and default Sims priors	0.254	0.240	0.251	0.269	0.266

Table A.7 RMSFE for hourly wage growth and HICP inflation for different horizons, 2017

	2017				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.508	0.345	0.321	0.312
Baseline lags and default Sims priors	-	0.410	0.368	0.326	0.509
HICP inflation - quarterly model					
Optimal lags and priors	-	0.183	0.216	0.279	0.332
Baseline lags and default Sims priors	-	0.187	0.195	0.255	0.294
HICP inflation - monthly model					
Optimal lags and priors	0.163	0.166	0.158	0.153	0.160
Baseline lags and default Sims priors	0.154	0.157	0.169	0.153	0.184

Table A.8 RMSFE for hourly wage growth and HICP inflation for different horizons, 2018

	2018				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.290	0.331	0.232	0.258
Baseline lags and default Sims priors	-	0.173	0.169	0.279	0.315
HICP inflation - quarterly model					
Optimal lags and priors	-	0.275	0.099	0.173	0.153
Baseline lags and default Sims priors	-	0.202	0.080	0.147	0.170
HICP inflation - monthly model					
Optimal lags and priors	0.164	0.150	0.150	0.160	0.153
Baseline lags and default Sims priors	0.193	0.153	0.153	0.177	0.166

Table A.9 RMSFE for hourly wage growth and HICP inflation for different horizons, 2019

	2019				
	Horizon (months)				
	1	3	6	12	24
Hourly wage growth - quarterly model					
Optimal lags and priors	-	0.274	0.298	0.298	0.268
Baseline lags and default Sims priors	-	0.306	0.277	0.218	0.259
HICP inflation - quarterly model					
Optimal lags and priors	-	0.355	0.266	0.199	0.231
Baseline lags and default Sims priors	-	0.306	0.277	0.218	0.259
HICP inflation - monthly model					
Optimal lags and priors	0.174	0.184	0.188	0.181	0.180
Baseline lags and default Sims priors	0.170	0.186	0.193	0.172	0.175