What were they thinking? Estimating the quarterly forecasts underlying annual growth projections

Jason Blunier, Christian Hepenstrick

SNB Working Papers
5/2022

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK

Legal Issues

ISSN 1660-7716 (printed version)
ISSN 1660-7724 (online version)

What were they thinking? Estimating the quarterly forecasts underlying annual growth projections *

Jason Blunier
jason.blunier@snb.ch

Christian Hepenstrick
christian.hepenstrick@snb.ch

May 9, 2022

Abstract

Many prominent forecasters publish their projections at an annual frequency. However, for applied work, an estimate of the underlying quarterly forecasts is often indispensable. We demonstrate that a simple state-space model can be used to obtain good estimates of the quarterly forecasts underlying annual projections. We validate the methodology by aggregating professional forecasts for quarterly GDP growth in the United States to the annual frequency and then applying our imputation methodology. The imputed forecasts perform as well as the original quarterly forecasts. Applying the imputation methodology to Consensus forecasts for other advanced economies provides further evidence of the good performance of our proposed methodology.

Keywords: *Forecasting, frequency disaggregation, survey expectations*
JEL Classification: *C53, E37*

*We thank Rina Rosenblatt-Wisch, Rolf Scheufele, participants at the internal SNB brownbag seminar and an anonymous referee for very helpful and constructive comments. The views, opinions, findings, and conclusions or recommendations expressed in this paper are strictly those of the author(s). They do not necessarily reflect the views of the Swiss National Bank (SNB). The SNB takes no responsibility for any errors or omissions in, or for the correctness of, the information contained in this paper. Corresponding author: Christian Hepenstrick.

1 Introduction

Many prominent forecasters publish their projections at an annual frequency. Examples include the IMF, the OECD, the monthly Consensus Economics survey and many central banks and government agencies. As a consumer of these projections, one would often like to know more about the underlying quarterly forecasts. An analyst may, e.g., need such a quarterly time series as conditioning variables in a quarterly macro model. Another example would be comparing annual forecasts to Q4/Q4 forecasts (e.g., annual Consensus projections vs. Q4/Q4 forecasts published in the Federal Reserve's Summary of Economic Projections). Yet another example is using professional forecasters' historical projection errors to calculate a measure for forecast uncertainty for quarterly growth h quarters ahead. For this, the analyst would require fixed-horizon forecasts (e.g., GDP growth h quarters ahead) as opposed to the typical fixed-event forecasts (GDP growth in a given calendar year). We demonstrate that a simple state-space model provides good estimates of the underlying quarterly projections.

The methodology assumes that the provider of an annual projection generates her forecasts at a quarterly frequency, i.e., extends the official quarterly data into the future. She then aggregates this underlying quarterly time series to the annual frequency to calculate the implied annual projection, which is then released to the public. Hence, the problem is to find a plausible extension of the quarterly data that results in aggregates that are fully consistent with the published annual projections.

We tackle this problem with a simple state-space framework that we illustrate and validate for GDP, but in principle, can be applied to any variable given appropriate adjustments (which are given in Appendix A). The well-known approximation of annual growth as a weighted average of quarterly growth rates is the key measurement equation, which is combined with a simple law of motion for quarterly GDP as the transition equation. Estimating the underlying quarterly growth rates with the Kalman filter is straightforward, and the data requirement is minimal, requiring only the annual projections and vintage of the GDP series of the day on which the projection was submitted. However, the question that immediately follows is how close the estimates actually get to the truth.

Therefore, we validate our methodology using the Federal Reserve Bank of Philadelphias's Survey of Professional Forecasts (Croushore, 1993). In this survey, participants submit quarterly forecasts for GDP growth in the United States. For validation, we aggregate the quarterly forecasts to the annual frequency and then run our methodology on these annual projections. We find that the implied quarterly projections are on average close to the original underlying quarterly forecasts. Moreover, in regard to forecasting performance, the forecast errors of the implied quarterly forecasts are indistinguishable from the errors of the original quarterly forecast. We then apply the methodology to Consensus forecasts for eight major advanced economies, documenting further desirable properties of the implied quarterly forecasts. We conclude that the methodology works very well and can be used in applied work when analysts need an estimate of the quarterly forecasts that underlie annual growth projections.

Our paper is related to several strands of literature. First, it is related to the literature on temporal disaggregation. Chow and Lin (1971, 1976) and their seminal papers were the first to present a unified framework to interpolate low-frequency time series using a higher-frequency indicator series. Fernandez (1981) and Litterman (1983) propose closely related approaches. National statistical offices use these frameworks to construct quarterly national accounts from annual national accounts (International Monetary Fund, 2017). While the original approaches were not formulated in the state-space framework, they can be easily cast into such a framework (e.g., Bernanke, Gertler, Watson, Sims, and Friedman (1997) and Cuche and Hess (2000) use the Kalman filter to interpolate quarterly GDP to the monthly frequency). In these temporal disaggregation applications, the higher frequency series (e.g., quarterly consumption) is unobserved, but a related higher frequency indicator (e.g., retail sales) is used to estimate the series under the restriction that it aggregates to the observed lower frequency series (e.g., annual consumption). This contrasts with our case, where part of the quarterly series is observed and the problem is to extrapolate the series such that it aggregates to the annual series. Since the unobserved part of the quarterly series lies in the future, there is no indicator series available. Therefore, the quarterly pattern is disciplined by assuming a law of motion as opposed to using an indicator series.[1]

Second, our paper is related to the literature backing out fixed-horizon forecasts from fixed-event ones. In this literature, fixed-horizon forecasts are generally obtained by taking weighted moving averages of fixed-event forecasts. Many studies, such as Dovern, Fritsche, and Slacalek (2012), use ad hoc weights to obtain a particular fixed-horizon forecast. More recently, Knüppel and Vladu (2016) show that the ad hoc weights are often far from the optimal weights derived by minimizing the mean squared approximation error. Both the ad hoc and the optimal weight approaches rely purely on forecasts. The simple state-space approach that we propose, in contrast, combines the quarterly data that the forecaster sees at the time of forming her forecast with the resulting annual projections to directly estimate the entire underlying path of quarterly forecasts. Based on this estimated quarterly time series, any fixed-horizon forecast can be easily calculated. In contrast, with the moving-average approach, one needs to calculate the weights for every forecast horizon individually.

Third, our paper is related to the literature documenting the properties of professional forecasts. While this literature directly discusses the properties of the submitted forecasts, we first disaggregate the projection to the higher frequency and then document the properties of the resulting forecasts. Generally, the properties documented for the original forecasts carry over to our disaggregated forecasts (e.g., differing forecast errors across economies, see Loungani (2001); declining errors with shorter forecast horizons, see Stark (2010); increasing revisions with shorter projection horizons, see Hadzi-Vaskov, Ricci, Werner, and Zamarripa (2021)).

The remainder of the paper proceeds as follows. Section 2 describes the methodology.

[1] Lenz (1999) shows how minimizing the sum of the second difference of a log-level series can be used to obtain smooth indicator-free disaggregations.

Section 3 validates the methodology using quarterly growth forecasts for the United States. Section 4 applies the methodology to a large set of advanced economies. Based on the insights of these sections, Section 5 concludes by making a recommendation to the applied economist for a simple and robust state-space setup to disaggregate annual projections on a regular basis.

2 Methodology

The methodology seeks to back out the quarterly profile underlying annual growth forecasts. It relies on the well-known approximation of annual growth being a weighted average of quarterly growth rates

$$y_t = \frac{1}{4}\left(y_{t,4} + 2y_{t,3} + 3y_{t,2} + 4y_{t,1} + 3y_{t-1,4} + 2y_{t-1,3} + y_{t-1,2}\right) + e_t, \tag{1}$$

where variables with only one subscript are annual variables and those with two subscripts are quarterly variables with the second subscript referring to the quarter within the year. y_t refers to the annual growth of year t, $y_{t,i}$ is the quarterly growth (not annualized) of quarter i in year t and e_t is the approximation error.[2]

For a current year t, one observes the consensus and a subset of the quarterly growth rates required in (1). Using this formula, one can directly calculate the approximate value to which the weighed sum of the unobserved quarters must aggregate. However, as soon as more than one quarter is unobserved, there is an infinite set of combinations of quarterly growth rates that fulfil the aggregation restriction. Hence, to obtain an estimate of the quarterly GDP profile that is behind the annual consensus forecast, some additional structure is needed.

We obtain this required additional structure by assuming a law of motion for quarterly GDP growth.[3] For illustrative purposes, we assume a random walk, but we will experiment later with different assumptions

$$y_{t,i} = Ly_{t,i} + v_{t,i}, \tag{2}$$

where L is the quarterly lag operator.[4] Since growth projections — those resulting from models and those based on expert judgment — are generally quite smooth, the random walk assumption is a good starting point.

Together, (1) and (2) constitute a state-space system with (1) being the measurement equation and (2) being the state equation.[5] Hence, the unobserved quarterly growth rates

[2] See, e.g., Tödter (2010) for a derivation of the approximation formula.

[3] One could argue that actual GDP does not necessarily need to follow the same law of motion as GDP forecasts and, hence, it is necessary to specify two processes — one for GDP and one for the forecast. In the case of the random walk assumption, this would only affect the variance of the error term. Later, when we experiment with an AR(2) law of motion, we will briefly return to this point.

[4] We use the lag operator to avoid notational difficulties associated with the value of i: if $i = 1$, $Ly_{t,i}$ is denoted as $y_{t-1,4}$, whereas for $i > 1$, it is $y_{t,i-1}$.

[5] A second measurement equation trivially links observed quarterly growth rates, $y_{t,i}^{obs}$ to the state variable;

underlying the annual growth forecast can be readily estimated using the Kalman filter. The mean and variance of the approximation error, e_t, can be directly estimated from the GDP data and, thus, are calibrated in the state-space system.[6] Therefore, the variance of the shock in the state equation is the only parameter to be estimated within the state-space system. We perform the estimation using maximum likelihood.

The state equation tends to generate a smooth projection as it wants to minimize the change in quarterly growth from one quarter to the next. However, note that in the measurement equation (1) the weights associated with the current year quarters are declining with the quarter. Hence, the filter will tend to use quarters later in the year to match the annual forecast. Appending an additional annual forecast, i.e., for year $t+1$ generates additional discipline, as, e.g., the last quarter of year t has a low weight for annual growth in t but a high weight for growth in $t+1$. We will document empirically how the forecast performance changes when additional annual growth projections are appended.

3 Validation of methodology using quarterly forecasts for the U.S.

In this section, we use vintages of the SPF to assess the empirical validity of our methodology. The SPF features quarterly growth forecasts that we aggregate to the annual frequency. The resulting annual forecasts are then disaggregated again to the quarterly frequency using the methodology outlined above. We then compare the resulting imputed quarterly forecasts with SPF's original quarterly forecasts. For our methodology, it is important that one uses the quarterly GDP data that the forecaster has at hand when forming her forecast. In applied practice, if one is disaggregating in real-time incoming forecasts, this is typically just the latest GDP vintage available and, hence, straightforward. For the evaluation, however, it is important to use the real-time vintage that was available in the period when the forecast was made.

3.1 Data

The SPF is a quarterly forecast survey conducted by the Federal Reserve Bank of Philadelphia[7]. The questionnaires are sent out after the Bureau of Economic Analysis has published its advanced GDP estimate at the end of the first month of each quarter. The deadline for submitting the answers is typically in the second or third week of the second month of the quarter. The survey collects forecasts for a wide range of variables, but we will only use the median of the forecasters quarterly growth projections and the median of the annual forecasts for years beyond the quarterly forecast horizon. We combine these forecasts

$y_{t,i}^{obs} = y_{t,i}$.

[6]Since GDP is growing on average, the mean approximation error is nonzero. To estimate the mean and variance of the approximation error, we calculate the annual growth rates and compare them to the growth rates obtained via the approximation formula. We choose to calibrate the variance of the measurement equation (as opposed to estimate it together with the other model parameters) to avoid a potential pile-up problem (see Stock and Watson (1998))

[7]Prior to the second quarter of 1990, the survey was conducted by the American Statistical Association and the National Bureau of Economic Research.

with the real-time GDP data that are available to the forecasters at the time of making the forecast as obtained from the Federal Reserve Bank of Philadelphia's real-time dataset (see Croushore and Stark (2001)).

Since the consensus survey always provides a forecast for the current and next year, we replicate this forecast horizon with the SPF. We combine real-time quarterly GDP data and quarterly forecasts to calculate the implied annual growth forecasts. The survey always asks for the first five unobserved quarters. Hence, we can always calculate the implied annual growth forecasts for the current year and, in the case of forecasts made during the fourth quarter, for the next year (in Q4, the survey participants provide forecasts for Q4 of the current year and for all four quarters of the following year). For quarters other than the fourth quarter, we use the annual forecast for the next year directly from the survey.

SPF vintages with quarterly forecasts and real-time GDP data are available from Q4 1968 onward. Starting in Q3 1981, the survey also asks about a growth forecast for the next calendar year; hence, we start the validation in Q3 1981. The annual growth forecasts two and three years into the future were first surveyed in Q2 2009, which will, thus, be the start of the evaluation of appending additional calendar years.

3.2 Baseline results

The baseline validation starts in Q3 1981 and considers the imputed quarterly forecasts that are derived from combining two annual forecasts with the real-time quarterly GDP data.

Figure 1 illustrates the procedure with the example of Q3 2016 — a quarter that is particularly illustrative for all the points that will be raised throughout the validation. The black line is the quarterly GDP growth rates as they were available to the SPF participants at the time (in the first half of October 2016, when the survey was conducted). The survey participants provided forecasts for GDP growth in Q3 2016 and the ensuing four quarters. The red line depicts the median of these original quarterly SPF forecasts. The forecasters expected a rapid rebound of growth after the subdued preceding quarters. After the rebound, growth is projected to moderate somewhat and reach 2.1% at the forecast horizon (Q3 2017).

For the validation of our methodology, we calculate annual growth in 2016 (first horizontal red line) and append the annual growth forecast for 2017 (second horizontal red line). The resulting annual growth forecasts and real-time GDP data are then passed to the Kalman filter to estimate the implied quarterly growth profile (blue line). Similar to the SPF, the imputed quarterly forecasts expect a rebound in Q3 2016. However, quarterly growth is projected to accelerate slightly further into Q4 and to moderate only then. At the forecast horizon, the imputed projection is marginally lower than the SPF original. Considering the typical volatility of GDP, the imputed projection is very close to the original quarterly SPF.

The finding that the imputed projection approximates the original SPF quite closely is not specific to this quarter but a general feature of the methodology. Table 1 reports

Figure 1: Illustration of the imputation procedure

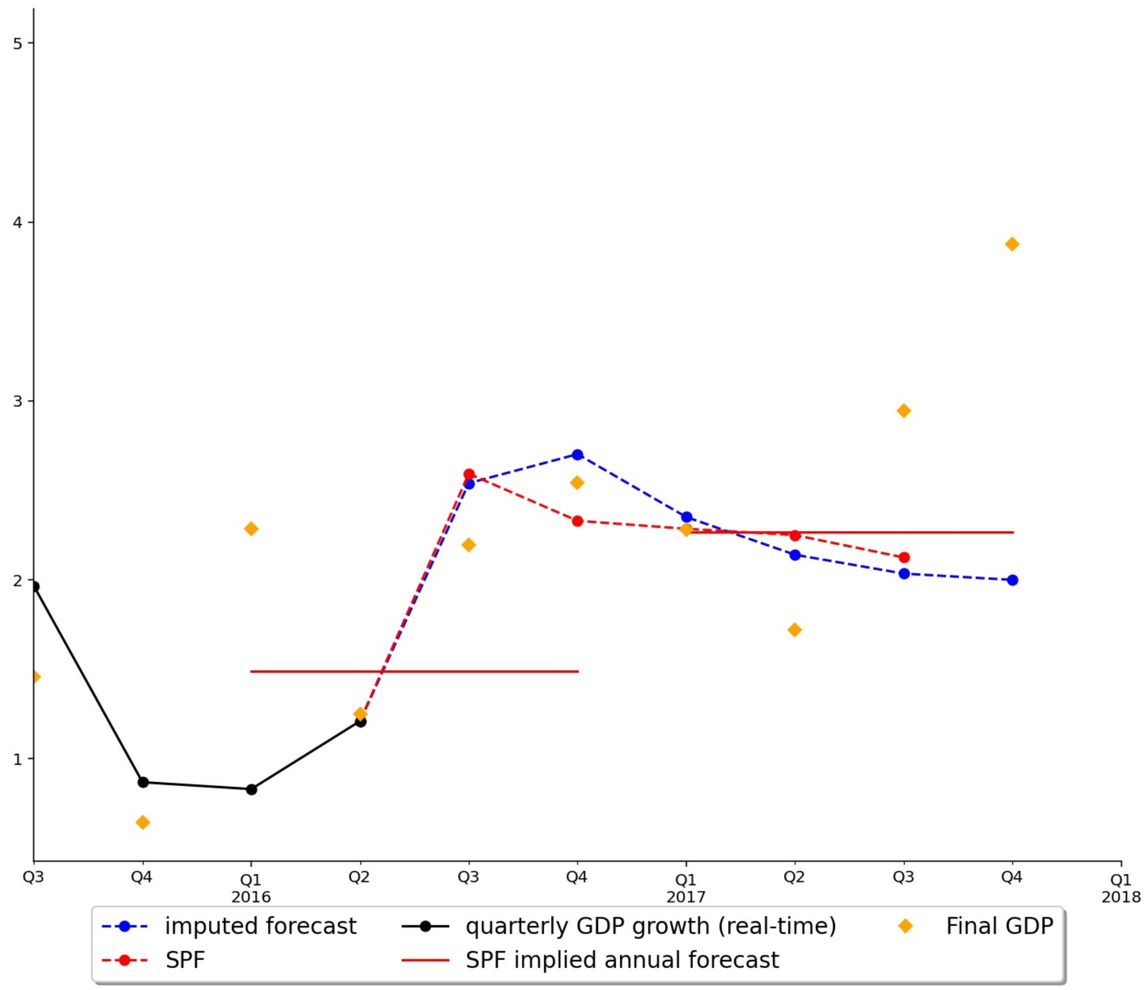

Notes: Illustration of the imputation procedure using real-time data for an imputation performed during Q3 2016. Inputs are the real-time quarterly GDP data and the annual growth projections, and the outputs are the imputed forecasts. These are compared to the original SPF projection and to the final GDP.

descriptive statistics on the difference between the original SPF forecast and the imputed projection separately for each forecasting horizon. For the nowcast (Q3 2016 in the example above), the imputed forecast was on average 0.2 pp higher than SPF. For 90% of quarters, the imputed nowcast was in a band of 1pp to the original SPF (between 0.7 pp above and 0.3 pp below). The standard deviation of the difference between the two forecasts is 0.4 pp, this compares with a standard deviation of GDP growth of 2.6 pp during the evaluation period. For the other forecast horizons, the imputed forecasts are also, on average, close to the original SPF, while the standard deviation of the difference between SPF and imputed quarterly forecasts rises somewhat (to a maximum of 0.6pp for the four-quarter ahead forecast).

Table 1: Descriptive statistics of the difference between original SPF and imputed forecasts

Forecast Horizon (in Quarters):	Current	One ahead	Two ahead	Three ahead	Four ahead
count	150.00	150.00	150.00	150.00	150.00
mean	-0.19	-0.02	0.10	0.20	0.22
std	0.39	0.39	0.43	0.50	0.60
min	-1.82	-1.58	-1.05	-0.94	-1.78
5%	-0.67	-0.68	-0.51	-0.39	-0.38
25%	-0.35	-0.16	-0.10	-0.07	-0.05
50%	-0.21	-0.01	0.05	0.11	0.14
75%	-0.03	0.18	0.27	0.39	0.39
95%	0.31	0.62	0.77	0.92	1.18
max	1.57	1.16	1.72	2.44	3.96

Notes: Difference calculated by subtracting imputed forecasts from the SPF forecasts. The evaluation sample starts in Q3 1981 and runs until Q4 2018. In Q4 2018, the SPF provided forecasts until Q4 2019.

Next, we turn to the question of whether the imputation leads to a deterioration of the forecasting performance compared to the original quarterly SPF. In Figure 1, this corresponds to comparing the SPF and the imputed forecasts (red and blue dashed lines) with realized GDP growth (yellow diamonds). In the example of Q3 2017, the imputed forecast performs better for the nowcast and the one- and three-quarters ahead forecasts, whereas the original SPF better anticipated growth two- and four-quarters ahead. Hence, for this illustrative quarter, it is not the case that the original SPF clearly performs better than the imputation.

This impression is confirmed when considering all quarters since 1981. Figure 2 plots the error with respect to final GDP (blue refers to the original quarterly SPF forecasts, red to the imputed quarterly forecasts) for each of the five forecasting horizons of the SPF. Clearly, the forecasting performance of the imputation is very close to the original SPF performance. This is also confirmed by the green line, which represents the difference of the eight-quarter rolling RMSEs of the two forecasts — sometimes the imputation performs slightly better, sometimes the original SPF, but the difference in the rolling RMSEs is always about one order of magnitude smaller than the forecasting errors per se and, as such, is negligible from an economic point of view.

Table 2 finally considers the average performance since 1981. Panel A compares the root mean squared error (RMSE) of both the original and the imputed forecasts with respect to final GDP, for which we use the vintage of January 28, 2021. The errors are expressed in terms of annual growth rates. We see that the forecast performance is essentially the same, with the difference always less than 10 basis points.

Table 2: RMSE of original SPF and imputed forecasts

Panel A: w.r.t final GDP

	Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
SPF	1.96	2.32	2.51	2.47	2.50
Imputed	1.97	2.36	2.49	2.44	2.52
DM (SPF - imputed)	-0.05	-0.19	0.11	0.19	-0.08

Panel B: w.r.t first release GDP

	Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
SPF	1.43	1.90	2.09	2.12	2.10
Imputed	1.49	1.91	2.05	2.06	2.11
DM (SPF - imputed)	-0.17	-0.02	0.18	0.23	-0.05

Notes: The table presents the root mean squared forecasting errors that one obtains by comparing the GDP realizations (Panel A: final GDP, Panel B: first release) with the forecasts. DM refers to the Diebold-Mariano test (Diebold and Mariano, 1995). Shown values represent the result of regressing the difference of squared forecast errors on a constant. The difference is calculated by subtracting the squared imputed forecast errors from the squared errors of the original quarterly SPF. Thus, positive values indicate that the imputed forecasts outperform the SPF and vice versa for negative values. An asterisk (*) denotes a rejection of the null hypothesis of the constant being equal to 0, i.e., equal forecasting performance. The evaluation sample starts in Q3 1981 and runs until Q4 2018. In Q4 2018, the SPF provided forecasts until Q4 2019.

Panel B presents the RMSE with respect to the first GDP release, i.e., the first official GDP number that is published for a particular quarter. Again, the differences are economically negligible, but interestingly, the original quarterly forecasts dominate statistically according to the Diebold-Mariano test. This makes intuitive sense, since forecasters observe in real-time the monthly indicators that the Bureau of Economic Analysis uses produce the first GDP release. Based on these monthly indicators, the forecasters seem to be able to predict the near-term dynamics slightly better than what follows when the annual projections are disaggregated. From an economic point of view, the differences remain very small, as seen in Figure 10 in the Appendix, which reproduces Figure 2 but replaces final GDP with the first release.

To summarize, our methodology can quite precisely back out the quarterly projections that are behind annual forecasts quite precisely. While the original quarterly projection cannot be estimated exactly, the resulting imputed quarterly forecasts have the same out-of-sample forecasting performance as the original quarterly forecasts.

From an applied perspective, it is common that not only forecast performance matters but that also a certain stability of the forecast is desirable. Consider, for example, a setup in

Figure 2: Evolution of forecast errors over time

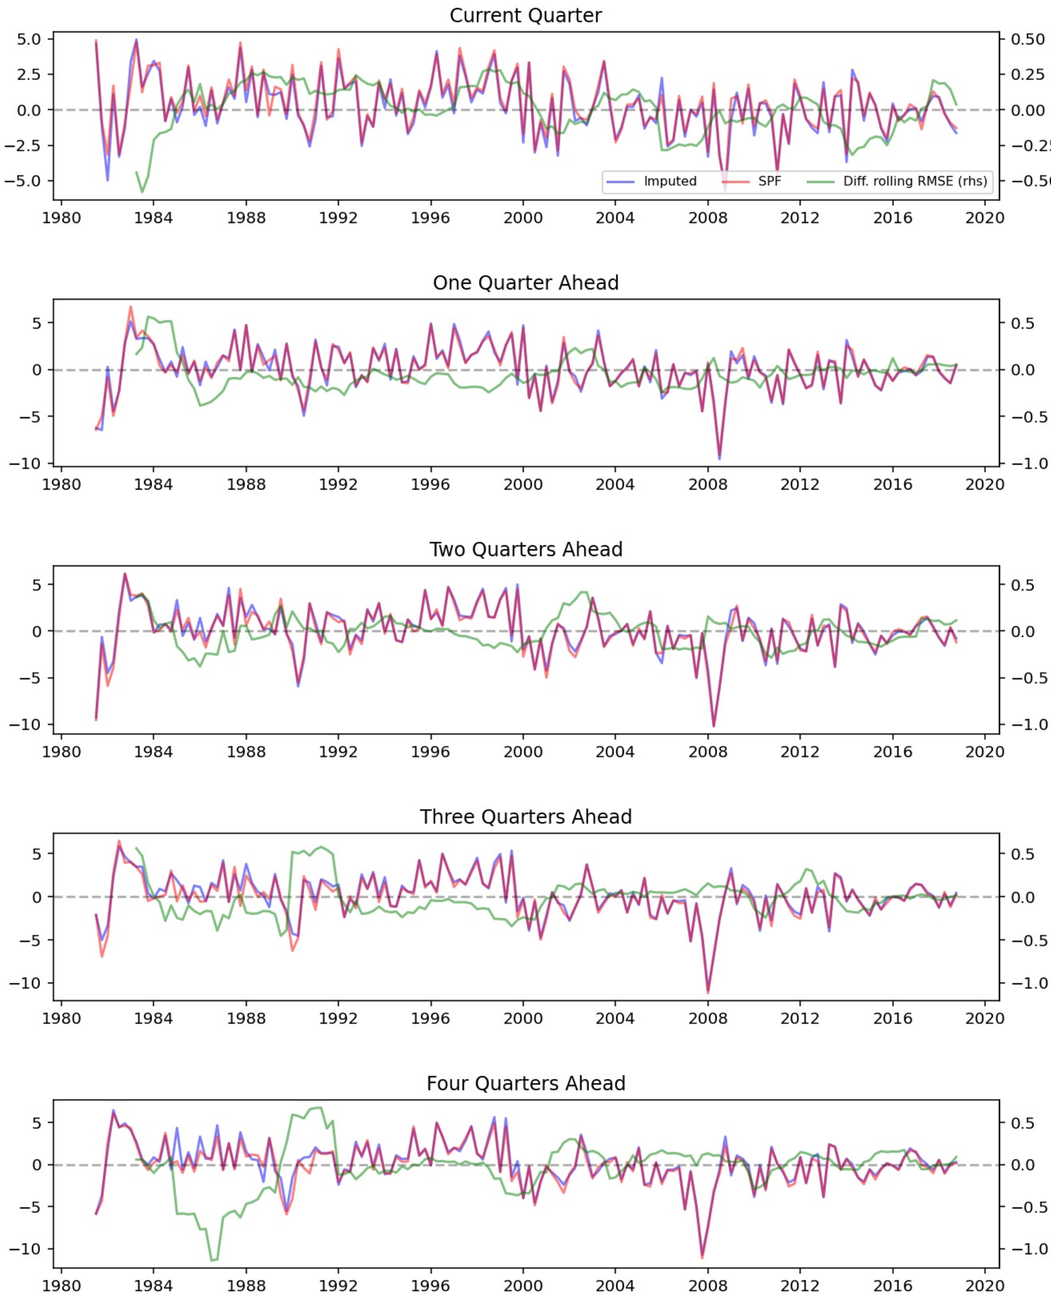

Notes: The five panels present the forecast errors of the original SPF (blue lines) and the imputed forecasts (red lines) together with the difference in the eight quarter rolling root mean squared forecasting errors (green lines). The panels differ by forecasting horizon, with the first panel depicting the nowcasts and the last panel depicting the four-quarters ahead forecast. The evaluation sample starts in Q3 1981 and runs until Q4 2018. In Q4 2018, the SPF provided forecasts until Q4 2019.

which quarterized Consensus forecasts are used as an exogenous conditioning variable for the domestic forecasting models used in a small open economy central bank's monetary policy process. Clearly, forecasts that are as accurate as possible are crucial to correctly condition the domestic forecasting model. Above, we have demonstrated that quarterized annual forecasts provide a competitive input. However, it could still be that for a given quarter that is forecasted, the methodology leads to large revisions from one consensus survey to the next. These revisions would then spill over into the domestic forecasting model so that this typically very important basis for the monetary policy decision becomes volatile and difficult to interpret. Therefore, we close this subsection by comparing the revision pattern of the SPF to those of the imputation methodology.

Figure 3 presents the results. The y-axis presents moments of the distribution of the revision pattern (thick lines denote the mean, thin lines the 25th and 75th percentiles), and the x-axis denotes when the revision is occurring, i.e., how many quarters it still takes to the first release (e.g., 3 compares how the forecast for a given quarter changes when moving from the three-quarters ahead forecast horizon to the two-quarters ahead horizon; 0 compares the nowcast, i.e., the forecast made during the quarter of interest to the first release).

Starting with the original quarterly SPF denoted in red, we see the well-known pattern that forecast revisions are typically larger for shorter-term forecasts when the information set is informative. When moving from the four-quarters ahead horizon to the three-quarters horizon (denoted with 4 on the x-axis), the mean absolute revision in the SPF is slightly below 0.4 percentage points. Only for 25% of the quarters in the evaluation is the revision larger than 0.5 percentage points (75th percentile). Interestingly, the distribution of the revision even compresses slightly when moving one quarter closer to the present (point 3 on the x-axis comparing the three-quarters ahead to the two-quarters forecast), but the compression is only small in economic terms (e.g., compared to the mean and the typical volatility of U.S. GDP growth). Moving even closer to the present, the mean absolute revision becomes larger and the distribution becomes wider as one would expect with an information set that becomes increasingly informative. However, even the revision from the one-quarter ahead forecast to the nowcast is small compared to the mean absolute forecast error (denoted with 0 on the x-axis; 0.6 vs. 1.2 percentage points).

Turning to the imputation denoted in blue, we see that the mean absolute revision is very close to that of the original SPF. The same holds true for the 25th percentile of the distribution of the revisions, but the 75th percentile is slightly larger for the revisions from the three- to the two- and the two- to one-quarter-ahead forecasts (denoted with 3 and 2 on the x-axis). Nevertheless, economically speaking and from a policy-maker perspective, the 75th percentiles are close [8]. These findings are confirmed when comparing the 95th and

[8]In small open economy models, the elasticity of domestic activity with respect to foreign activity is typically on the order of 0.5 to 0.9 (see, e.g., Bäurle and Steiner (2015) for the case of Switzerland). Hence, with an elasticity of 0.7, the 75th percentile of the distribution of the revision from the three- to the two quarter-ahead forecast in the SPF of approximately 0.5 percentage points would lead to a revision of 0.35 percentage points. The corresponding percentile of the imputed forecasts is 0.7, and hence, the forecast for annualized domestic activity growth would be revised by approximately 0.5pp. The difference in the revision of 0.15 percentage

Figure 3: Revision patterns

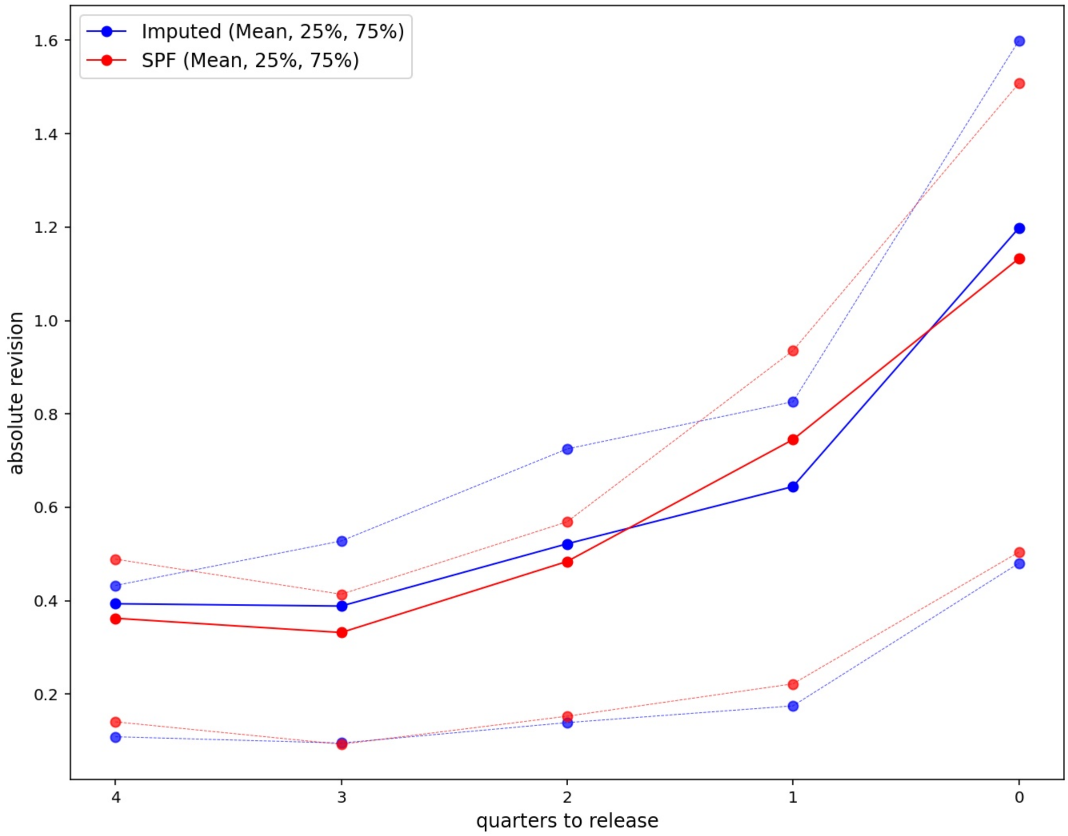

Notes: The figure presents moments of the distribution of the absolute forecast revision from one quarter to the next, i.e., the target quarter is held constant, and the forecast changes as the target quarter is approached are tracked. The x-axis denotes quarters to the release, and 0 denotes the difference of the forecast made during the target quarter and the first official GDP release. The thick lines show the median absolute revision, and the thin lines show the 25th and 75th percentiles. The evaluation sample starts in Q3 1981 and runs until Q4 2018. In Q4 2018, the SPF provided forecasts until Q4 2019.

the maximum of the forecast revisions (Figure 11 in the Appendix). Hence, we conclude that the imputation methodology delivers forecasts that are comparable to the underlying original quarterly forecast not only in terms of forecast performance but also in terms of revision pattern.

3.3 The effects of appending additional annual projections

As argued in section 2, the proposed algorithm tends to use quarters toward the end of the forecast period to match the annual growth rate. Hence, this subsection investigates the effects of appending additional annual growth projections. This can be done because the SPF started in 2009, asking about growth expectations two and three years ahead.

Figure 4 shows the effect of appending additional years, again for the example of forecasts made in Q3 2016. The forecasters expected annual growth to decelerate slowly, as seen from the purple (light blue) horizontal line depicting the annual forecast two (three) years ahead. The fact that the deceleration is only marginal relative to what is expected for 2017 forces the imputed quarterly forecasts up toward the end of 2017. When the three-year ahead forecast is also passed to the filter (light blue lines), the effect on the growth rates during 2017 is only minor, but again, the growth rates toward the end of the projection period, i.e., in H2 2018 are revised considerably.

Table 4 investigates if and how the forecast performance changes when additional forecast years are passed to the filter. The evaluation can only start in 2009, since before then, there are no two- and three-year ahead forecasts available in the SPF. First, the finding reported previously that the forecasting performance of the imputation is essentially identical to that of the original SPF also holds up in this smaller sample (first two rows of table). Appending additional annual forecasts changes the projections, as we have just seen with the example of Q3 2016, although not in a systematic way, as witnessed by the fact that the RMSEs are practically identical across the three versions of the imputation (rows two to four). For the analyst applying our methodology, this implies that she does not gain any forecasting accuracy by appending as many years as are available, but neither does she lose any accuracy. Hence, we would suggest that for convenience, one should append as many years as available.

points is negligible compared to the typical mean and variance of GDP growth and would not lead to different policy implications.

Figure 4: Illustration of how the imputation changes if additional annual projections are appended

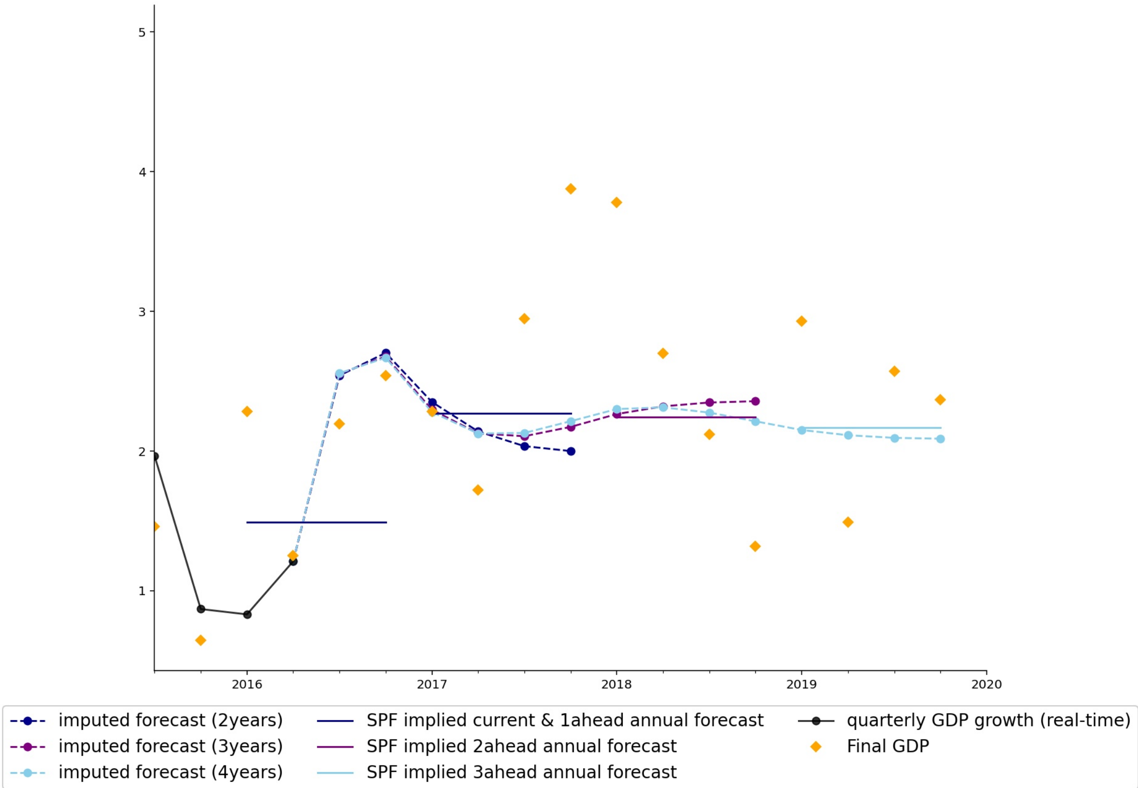

Notes: Illustration of the imputation procedure using real-time data for an imputation performed during Q3 2016. Inputs are the real-time quarterly GDP data and the annual growth projections, and the outputs are the imputed forecasts. The figure shows how the imputed forecasts change if additional annual projections are appended to the inputs.

Table 4: RMSEs when additional annual forecasts are appended

Panel A: w.r.t final GDP

		Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
	SPF	1.45	1.57	1.57	1.51	1.55
Imputed	2years	1.53	1.57	1.61	1.52	1.52
Imputed	3years	1.52	1.57	1.59	1.53	1.52
Imputed	4years	1.52	1.57	1.58	1.52	1.52
DM test	SPF - 2 years	-0.21	0.01	-0.16	-0.02	0.09
DM test	2 - 3 years	0.01	-0.01	0.07*	-0.03	-0.01
DM test	2 - 4 years	0.02	-0.01	0.10*	-0.02	-0.02

Panel B: w.r.t first release GDP

		Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
	SPF	1.13	1.36	1.33	1.22	1.25
Imputed	2years	1.17	1.30	1.33	1.22	1.23
Imputed	3years	1.17	1.30	1.32	1.21	1.22
Imputed	4years	1.17	1.30	1.31	1.20	1.21
DM test	SPF - 2 years	-0.21	0.01	-0.16	-0.02	0.09
DM test	2 - 3 years	0.01	-0.01	0.07*	-0.03	-0.01
DM test	2 - 4 years	0.02	-0.01	0.10*	-0.02	-0.02

Notes: The table presents the root mean squared forecasting errors that one obtains by comparing the GDP realizations (Panel A: final GDP, Panel B: first release) with the forecasts. The imputed forecasts differ by how many annual growth projections are used (2 to 4 years). DM refers to the Diebold-Mariano test (Diebold and Mariano, 1995). Shown values represent the result of regressing the difference of the squared forecast errors on a constant. An asterisk (*) denotes a rejection of the null hypothesis of the constant being equal to 0, i.e., equal forecasting performance. The evaluation sample starts in Q2 2009 and runs until Q4 2018. In Q4 2018, the SPF provided forecasts until Q4 2019.

3.4 The effects of different laws of motion in the transition equation

Next, we turn to the question of the role of the law of motion 2. Instead of a random walk, we use in this section an AR(2) process

$$y_{t,i} = \phi_0 + \phi_1 L y_{t,i} + \phi_2 L^2 y_{t,i} + v_{t,i}. \tag{3}$$

In contrast to the random walk, we now have three additional parameters that need to be estimated (ϕ_0, ϕ_1, ϕ_2).

Figure 5 reproduces Figure 1 but adds, in turquoise, the imputation that results if (3) is used as a law of motion. The resulting dynamics are somewhat different from the baseline, but again, the imputed projection is close to the original SPF, and there is no clear ranking in terms of forecasting performance. We have chosen the quarter presented in Figure 5 so that there is indeed a visible difference between the two imputations. For most other

Figure 5: Illustration of how the imputation changes with an AR(2) transition equation

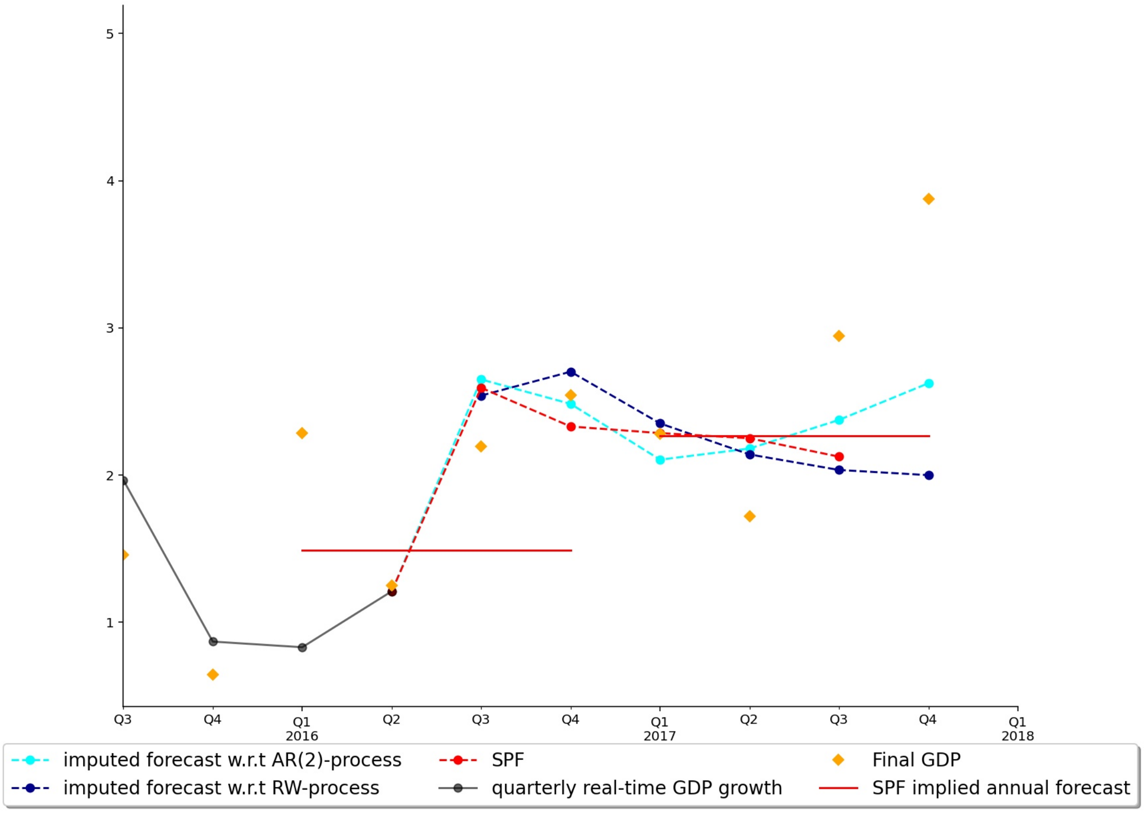

Notes: Illustration of the imputation procedure using real-time data for an imputation performed during Q3 2016. Inputs are the real-time quarterly GDP data and the annual growth projections, and the outputs are the imputed forecasts. The figure shows how the imputed forecasts change if an AR(2) is assumed to be the law of motion of quarterly GDP growth as opposed to the random walk assumption used in the baseline specification.

quarters, the difference is much smaller so that we are convinced that in applied work (e.g., a quarterly model being conditioned on the imputation), the implications of which imputation is chosen are negligible.

Table 6 confirms this impression. There differences in the forecasting performance are infinitesimal both for the long sample with two annual forecasts passed to the filter (Panel A) and for the short sample with four annual forecasts (Panel B).[9]

One may argue that the data generating process for the forecasts is different from the process for actual GDP. Hence, one would need to estimate the parameters of (3) based on historical forecasts and not on actual GDP data. Indeed, if we estimate the parameters with a panel AR regression on the SPF vintages, we obtain coefficients with a higher persistence ($\phi_1 = 0.66$, $\phi_2 = -0.06$, $\phi_0 = 1.11$ as opposed to 0.32, 0.11, and 1.75). However, in terms of forecasting performance, the difference is essentially invisible, as illustrated in Figure 6. Therefore, for applied economists, we recommend parametrizing the state-space system with actual GDP data, since this is straightforward to implement.

[9]Conventional information criteria (AIC and BIC) suggest that the AR model fits the data somewhat better than the RW model.

Figure 6: RMSE of imputed forecasts with different laws of motion

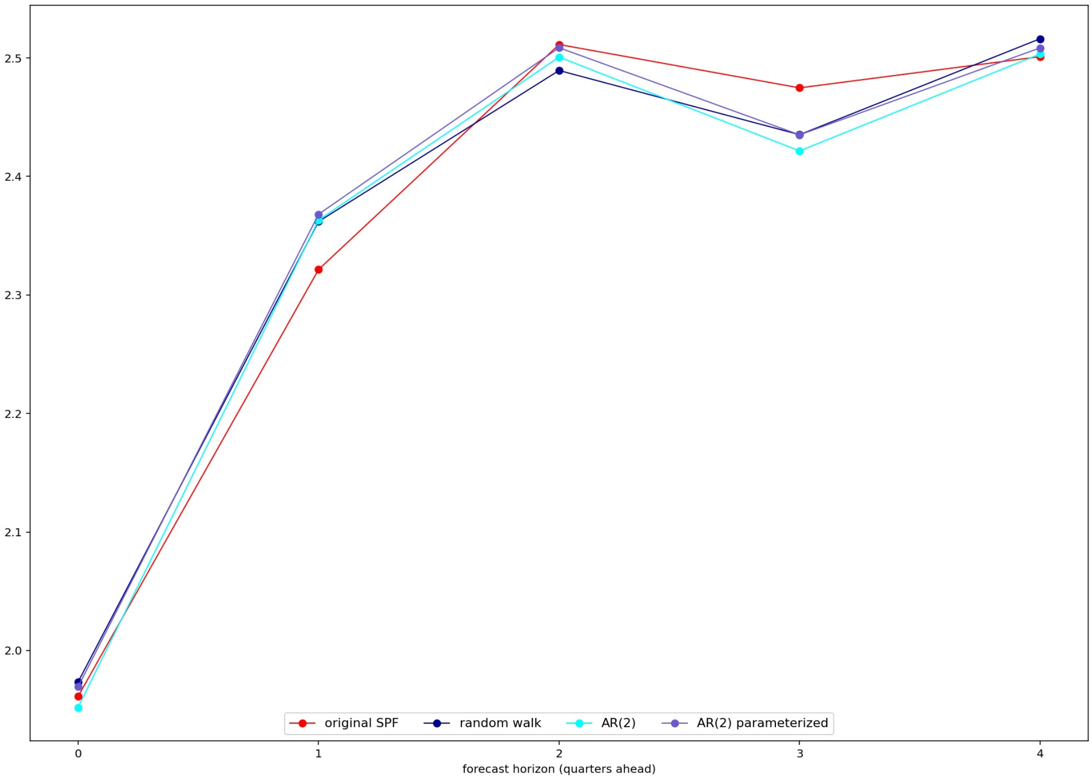

Notes: The figure presents the root mean squared forecasting errors across different forecasting horizons for the original SPF and three different imputations (random walk law of motion, AR law of motion, AR parametrized with parameters estimated from a panel AR on the SPF vintages). The evaluation sample is 1981Q3 up until 2018Q4 with a forecast horizon until 2019Q4.

Table 6: RMSE of Imputed Forecasts with AR(2) in transition equation

Panel A: since 1981, conditional on one and two ahead forecasts, w.r.t final GDP

		Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
	SPF	1.96	2.32	2.51	2.47	2.50
Imputed	AR(2)	1.95	2.36	2.50	2.42	2.50
Imputed	RW	1.97	2.36	2.49	2.44	2.52
DM test	SPF - AR(2)	0.04	-0.20	0.05	0.26	-0.01
DM test	AR(2) - RW	-0.08	0.01	0.06	-0.07	-0.06

Panel B: since 2009, conditional on one & up until four ahead forecasts, w.r.t final GDP

		Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
	SPF	1.45	1.57	1.57	1.51	1.55
Imputed	AR(2)	1.48	1.55	1.58	1.53	1.53
Implied	RW	1.52	1.57	1.58	1.52	1.52
DM test	SPF - AR(2)	-0.09	0.05	-0.03	-0.05	0.04
DM test	AR(2) - RW	-0.11	-0.05	-0.02	0.01	0.03

Notes: The table presents the root mean squared forecasting errors that one obtains by comparing the GDP realizations (Panel A: two annual projections, Panel B: four annual projections) with the forecasts. The imputed forecasts differ by the law of motion that is assumed in the transition equation. DM refers to the Diebold-Mariano test (Diebold and Mariano, 1995). Shown values represent the result of regressing the difference of the squared forecast errors on a constant. An asterisk (*) denotes a rejection of the null hypothesis of the constant being equal to 0, i.e., equal forecasting performance. The evaluation sample in Panel A ranges from Q3 1981to Q4 2018 . The evaluation in Panel B starts in Q2 2009.

In summary, we first conclude that our imputation approach yields quarterly forecasts that are close to the original quarterly SPF forecasts. Second, the forecasting performance of the imputation is very similar to the original quarterly SPF, with the possible exception of the nowcast and the one quarter ahead forecast (when assessed relative to the first GDP release). Here, the SPF has a slight edge relative to the imputation. Third, the forecasting performance of the imputation does not depend on how many years are appended or on whether a random walk or an AR(2) is assumed as the law of motion. Hence, it is possible to back out from annual GDP projections quarterly growth forecasts that are competitive with the quarterly projections underlying the annual forecasts. The next section applies this finding to eight advanced economies and discusses the features of the resulting quarterly forecasts.

4 Results for other economies

This section applies our methodology to eight advanced economies. We use the baseline specification with a random walk law of motion and two annual growth forecasts.

4.1 Data

We use the real GDP forecasts collected by Consensus Economics (CE), one of the leading economic survey organizations. We consider forecasts for the euro area, France, Germany, Italy, Japan, Switzerland, the United Kingdom, and the United States. The sample starts in 1999 and ends in 2018 for all economies except for the euro area, where the sample only starts in 2002. The monthly CE survey asks participants about their forecasts for GDP growth in the current year and in the next year. We combine these forecasts with the quarterly real GDP data as observed at the time of the survey period according to the OECD real-time database (McKenzie, 2006). Later, we compare the forecasting performance of our imputations to the quarterly projections that are surveyed by CE in the last month of the quarter.[10] These forecasts cover the next seven or eight unobserved quarters.

4.2 Results

Figure 7 presents the root mean squared forecasting error as a function of the forecasting horizon (now monthly because the CE survey is performed monthly). Time 0 represents the last month of the target quarter, e.g., March 2015, when forecasting GDP growth in Q1 2015. The RMSEs are generally upward sloping with the forecasting horizon, reflecting the fact that the information set at the time of the forecast contains more information about near-term developments as opposed to developments further in the future. The level of the RMSEs differs across economies, with France, the euro area, the U.S. and the UK featuring the lowest RMSEs.

These differences are driven by different volatilities of the target variables. To illustrate this, Figure 8 plots the RMSEs relative to the standard deviation of an economy's GDP growth. With this standardization, the RMSEs are much closer across economies. Note that the standard deviation of GDP growth is also equal to the RMSE if the in-sample mean was used as a forecast. For most economies, the relative RMSE of the imputation is below 1, suggesting that the imputation performs better than the simple benchmark of a constant forecast.

In the last month of the quarter, CE panel participants submit quarterly growth forecasts in addition to their annual projections. For most economies, these forecasts are y/y. Hence, to compare the performance of our imputation to the quarterly projection, we calculate the y/y growth rate that is implied by our imputation of the annual forecasts that are submitted in the same month. Table 8 presents the results. For most economies, the RMSE of the quarterly CE projections is marginally lower compared to the imputation's RMSE. For approximately one-quarter of the economy-horizon combinations, the forecasting performance of the quarterly CE projection is statistically better according to the Diebold-Mariano test. However, in an economic sense, the performance is very similar, as the difference is typically on the order of a few basis points, which is negligible compared to the variability of the target that is between one and two percentage points (on a

[10] We use the GDP vintages as downloaded in October 2020 as final GDP.

Figure 7: RMSE across economies

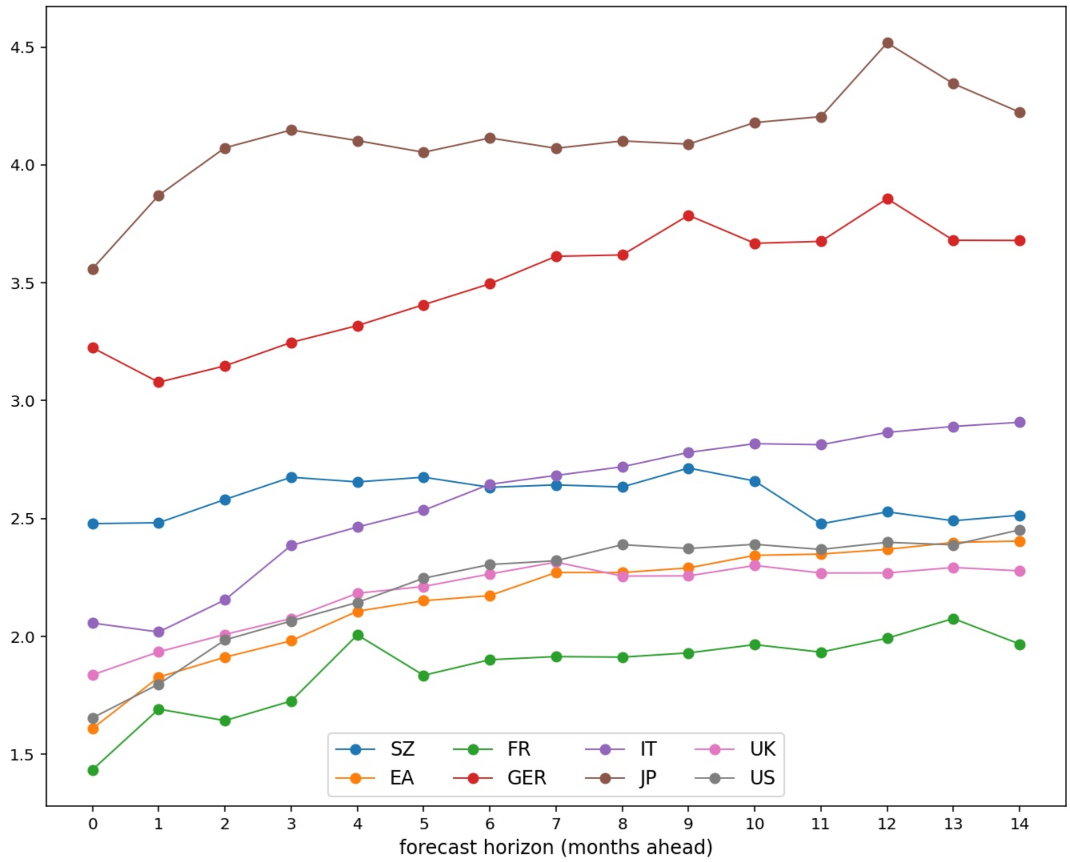

Notes: Each line represents one economy's root mean squared forecasting error for different forecasting horizons. Time period 0 represents a forecast that is made in the last month of the target quarter (e.g., forecasting Q1 2015 in March 2015). Time period 14 represents a forecast made four quarters ahead in the first month of a quarter (e.g., forecasting Q1 2016 in January 2015).

y/y basis). Moreover, not all forecasters submitting annual projections submit quarterly forecasts. Hence, part of the difference may also come from the fact that these two groups have differing forecasting performances (even when considering annual projections only).

Table 8: RMSE of CE forecasts and imputed forecasts (baseline)

Year over year forecasts

		Current Quarter	One Quarter Ahead	Two Quarters Ahead	Three Quarters Ahead	Four Quarters Ahead
SZ	CE	2.55	2.61	2.68	2.70	2.49
	imputed	2.58	2.70	2.75	2.71	2.54
	DM-test	-0.17	-0.48**	-0.37**	-0.08	-0.25
EA	CE	0.66	0.91	1.20	1.46	1.69
	imputed	0.69	0.94	1.24	1.50	1.69
	DM-test	-0.04***	-0.05	-0.12	-0.10	-0.01
FR	CE	0.61	0.76	1.00	1.24	1.41
	imputed	0.62	0.78	1.02	1.28	1.45
	DM-test	-0.01	-0.03	-0.05	-0.10***	-0.13*
GER	CE	0.95	1.27	1.63	2.02	2.32
	imputed	1.12	1.45	1.77	2.09	2.39
	DM-test	-0.36***	-0.51***	-0.48**	-0.30*	-0.35
IT	CE	0.82	1.15	1.50	1.85	2.10
	imputed	0.85	1.19	1.55	1.88	2.10
	DM-test	0.04	-0.10*	-0.15	-0.13	-0.03
JP	CE	1.12	1.48	1.78	2.09	2.33
	imputed	1.25	1.56	1.81	2.10	2.48
	DM-test	-0.31***	-0.24*	-0.09	-0.06	-0.75
UK	CE	1.13	1.22	1.37	1.51	1.64
	imputed	1.14	1.23	1.38	1.53	1.65
	DM-test	-0.03	-0.02	-0.03	-0.06	-0.06
U.S.	CE	0.65	0.88	1.10	1.34	1.56
	imputed	0.66	0.85	1.08	1.30	1.53
	DM-test	-0.01	0.04	0.06**	0.10**	0.09

Notes: DM refers to the Diebold-Mariano test (Diebold and Mariano, 1995). Shown values represent the result of regressing the difference of squared forecast errors on a constant. The difference is calculated by subtracting the squared imputed forecast errors from the squared CE forecast errors. Thus, positive values indicate that the imputed forecasts outperform the quarterly Consensus forecasts. An asterisk (*) denotes a rejection of the null hypothesis of the constant being equal to 0, i.e., equal forecasting performance. Forecast vintages starting in March 1999 for all economies, except the EA that starts in December 2002.

Finally, Figure 9 presents the forecast revisions from one month to the next. These revisions feature clear seasonal patterns that are driven by the release of a new GDP data point that occurs once a quarter. In the months when these data are released, the new data point is the main driver of the forecast revision. In the other two months, the forecast revision is only driven by changes in the annual Consensus projection (or back-data revisions to the official quarterly GDP series). This within-quarter seasonality dominates the tendency for the revisions to increase with a shorter forecasting horizon. Figure 12 in the Appendix presents the revision pattern over the whole quarters, and there the tendency for revisions to increase with a shorter forecasting horizon becomes apparent for all economies.

Figure 9: Revision pattern across economies

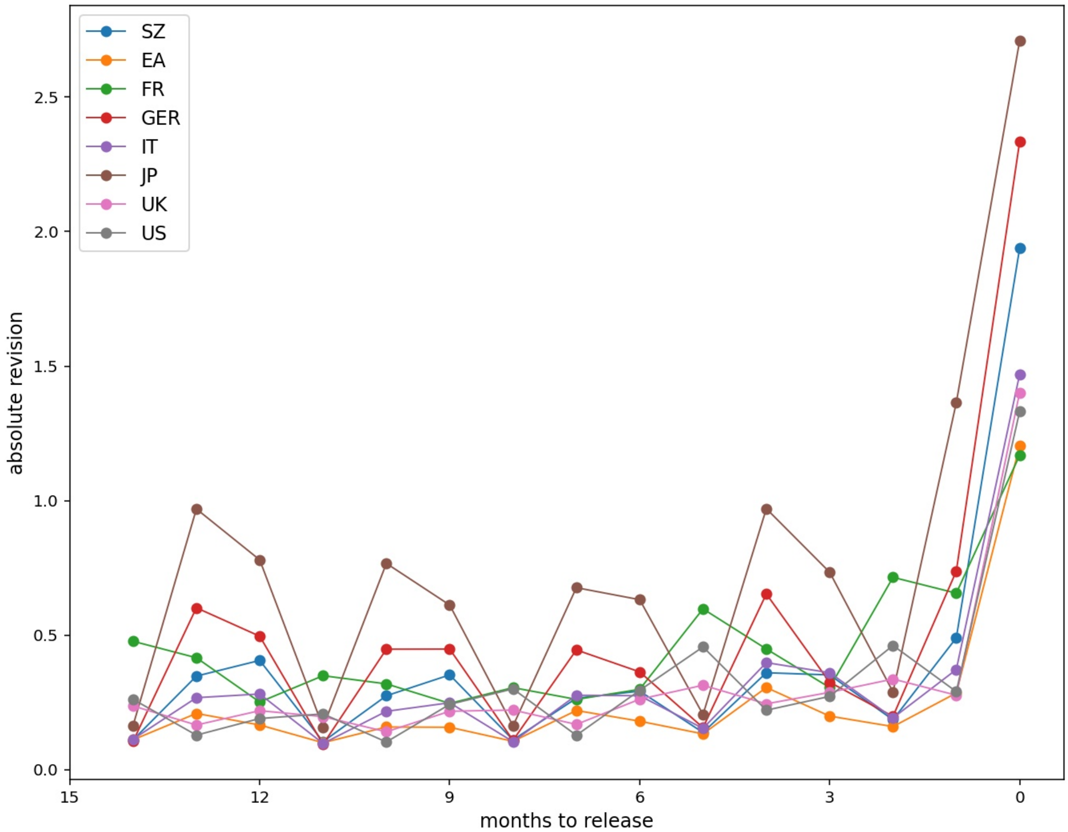

Notes: Each line represents the mean absolute forecast revision as the target quarter is approached. The x-axis represents months prior to the end of the target quarter, while 0 indicates the difference between the forecast made in the last month of the target quarter and the first official GDP release.

5 Recommendation to the applied economist and conclusions

The validation for the U.S. and the results for the other economies illustrate that our proposed disaggregation methodology performs well. In fact, while we were expecting that the methodology performs well for longer forecasting horizons, we were genuinely surprised by the performance for near-term horizons.

With hindsight, we see two reasons for the good performance. First, GDP forecasting is difficult, i.e., even with a very large information set, it is difficult to forecast the quarterly growth pattern beyond a general tendency that is well reflected in the smooth imputed forecasts. Second, while forecasting the quarterly growth pattern is difficult, taking the available quarterly data on board when producing annual forecasts is very important because of carry-over effects (see Tödter (2010)). Hence, given that available quarterly GDP data are incorporated into the annual projections and that underlying growth quarterly growth forecasts are smooth, our imputation methodology seems to come close to the underlying projections.

A corollary of the second point is that it would be wrong to conclude from our results that forecasting at the quarterly frequency is futile as the imputations of annual projections perform as well as quarterly projections. The imputations only perform well because the annual projections make use of the available quarterly data.

For the applied economist looking for quarterly growth projections that are regularly updated reflecting incoming information, we recommend the following: use the annual CE projections that are released every month but append the latest longer-term forecasts (two and more years ahead) that are published once every six months (or every three months depending on your CE subscription). Apply then our imputation with the AR(2) law of motion. This yields very competitive quarterly GDP projections for long projection periods. Of course, the methodology can be applied to other variables, such as the CPI or the unemployment rate. Depending on the nature of the variable (e.g., stock vs. flow), one will have to adjust equation 1. Appendix A provides examples.

References

Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A., & Friedman, B. M. (1997). Systematic monetary policy and the effects of oil price shocks. *Brookings Papers on Economic Activity*, *1997*(1), 91–157.

Bäurle, G., & Steiner, E. (2015). How do individual sectors respond to macroeconomic shocks? a structural dynamic factor approach applied to swiss data. *Swiss Journal of Economics and Statistics (SJES)*, *151*(III), 167-225.

Chow, G. C., & Lin, A.-l. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. *The Review of Economics and Statistics*, 372–375.

Chow, G. C., & Lin, A.-L. (1976). Best linear unbiased estimation of missing observations in an economic time series. *Journal of the American Statistical Association*, *71*(355), 719–721.

Croushore, D. (1993). Introducing: the survey of professional forecasters. *Business Review-Federal Reserve Bank of Philadelphia*, *6*, 3.

Croushore, D., & Stark, T. (2001, November). A real-time data set for macroeconomists. *Journal of Econometrics*, *105*(1), 111-130.

Cuche, N. A., & Hess, M. K. (2000). Estimating monthly gdp in a general kalman filter framework: Evidence from switzerland. *Economic & Financial Modelling*, 153.

Dovern, J., Fritsche, U., & Slacalek, J. (2012, 11). Disagreement Among Forecasters in G7 Countries. *The Review of Economics and Statistics*, *94*(4), 1081-1096.

Fernandez, R. B. (1981). A methodological note on the estimation of time series. *The Review of Economics and Statistics*, *63*(3), 471–476.

Hadzi-Vaskov, M., Ricci, L. A., Werner, A. M., & Zamarripa, R. (2021). Patterns in imf growth forecast revisions: A panel study at multiple horizons. *CEPR Discussion Paper No. DP16153*.

International Monetary Fund. (2017). *Quarterly national accounts manual* (Tech. Rep.). International Monetary Fund (IMF).

Knüppel, M., & Vladu, A. L. (2016). Approximating fixed-horizon forecasts using fixed-event forecasts. *Deutsche Bundesbank Discussion Papers*(28/2016).

Lenz, C. (1999). *Interpolation von Quartalsdaten aus Jahresdaten.* Unpublished Manuscript.

Litterman, R. B. (1983). A random walk, markov model for the distribution of time series. *Journal of Business & Economic Statistics*, *1*(2), 169–173.

Loungani, P. (2001). How accurate are private sector forecasts? cross-country evidence from consensus forecasts of output growth. *International Journal of Forecasting*, *17*(3), 419–432.

McKenzie, R. (2006). Undertaking revisions and real-time data analysis using the oecd main economic indicators original release data and revisions database. *OECD Statistics Working Papers*(2006/02).

Stark, T. (2010). Realistic evaluation of real-time forecasts in the Survey of Professional Forecasters. *Research Rap Special Report*(May).

Stock, J. H., & Watson, M. W. (1998). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association*, *93*(441), 349-358.

Stock, J. H., & Watson, M. W. (2007, February). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, *39*(s1), 3-33.

Tödter, K.-H. (2010). How useful is the carry-over effect for short-term economic forecasting? *Discussion Paper Series 1: Economic Studies*(2010,21).

A The state-space system

To write out the state-space system, some preparatory definitions are helpful. Let t denote a quarterly time index and $q(t)$ a function that returns the quarter of time period t and $ye(t)$ returning the year of period t, e.g., if $t =$ Q2 2000, $q(t) = 2$ and $ye(t) = 2000$. We denote quarterly GDP with y_t^{obs}. If quarterly GDP is not observed (because it is in the forecasting period), it simply takes the value $missing$. We denote the annual forecasts with y_t^A. Finally, we define $\tilde{y}_t^A$ as follows:

$$\tilde{y}_t^A = \begin{cases} y_t^A & \text{if } q(t) = 4 \text{ and forecast available for } ye(t). \\ missing & \text{otherwise.} \end{cases}$$

x_t denotes the unobserved state variable (implied quarterly GDP growth in our case).

A.1 Measurement equations

Generally, the measurement equation takes the following form:

$$\begin{bmatrix} \tilde{y}_t^A \\ y_t^{obs} \end{bmatrix} = C \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} + \begin{bmatrix} e_t \\ 0 \end{bmatrix},$$

where the form of C depends on the nature of the variable for which we have the annual projection.

In the case of disaggregated annual growth, C is

$$C = \begin{bmatrix} \frac{1}{4} & \frac{2}{4} & \frac{3}{4} & \frac{4}{4} & \frac{3}{4} & \frac{2}{4} & \frac{1}{4} \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

in the case of Q4/Q4 growth (as, e.g., used in the FOMC's Summary of Economic Projections) or stock variables for which the average level is forecasted (e.g., the average unemployment rate during a given year), C reads

$$C = \begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix},$$

and in the case of stock variables for which the annual forecast refers to the level of the variable during the last quarter of the year (which is often used for interest rates or unemployment rates), C reads

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

A.2 Transition equations

In the case of a random walk, the transition equation reads

$$\begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} + \begin{bmatrix} v_t \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

In the AR(2) case, it reads

$$\begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} = \begin{bmatrix} \phi_0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \phi_1 & \phi_2 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} + \begin{bmatrix} v_t \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

If one is disaggregating inflation forecasts, one may also find it worth modeling the law of motion as an exponentially weighted moving average. The reasoning being that trend inflation (see, e.g., Stock and Watson (2007)) is often seen as one of the toughest benchmarks when forecasting inflation and trend inflation is closely related to the exponentially weighted moving average. In this case, the transition equation would read as

$$\begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} = \begin{bmatrix} \phi_0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} \bar{\omega}_1 & \bar{\omega}_2 & \bar{\omega}_3 & \bar{\omega}_4 & \bar{\omega}_5 & \bar{\omega}_6 & \bar{\omega}_7 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ x_{t-1} \\ x_{t-2} \\ x_{t-3} \\ x_{t-4} \\ x_{t-6} \\ x_{t-7} \end{bmatrix} + \begin{bmatrix} v_t \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix},$$

with $\bar{\omega}_i$ defined as

$$\bar{\omega}_i = \frac{\omega^i}{\sum_{k=1}^{7} \omega^k}$$

and ω is a parameter to be estimated. Of course, the length of the state vector could be extended if one wants to allow for a slower decay, i.e., ω close to 1.

B Additional figures

Figure 10: Evolution of forecast errors over time evaluated w.r.t. the first GDP release

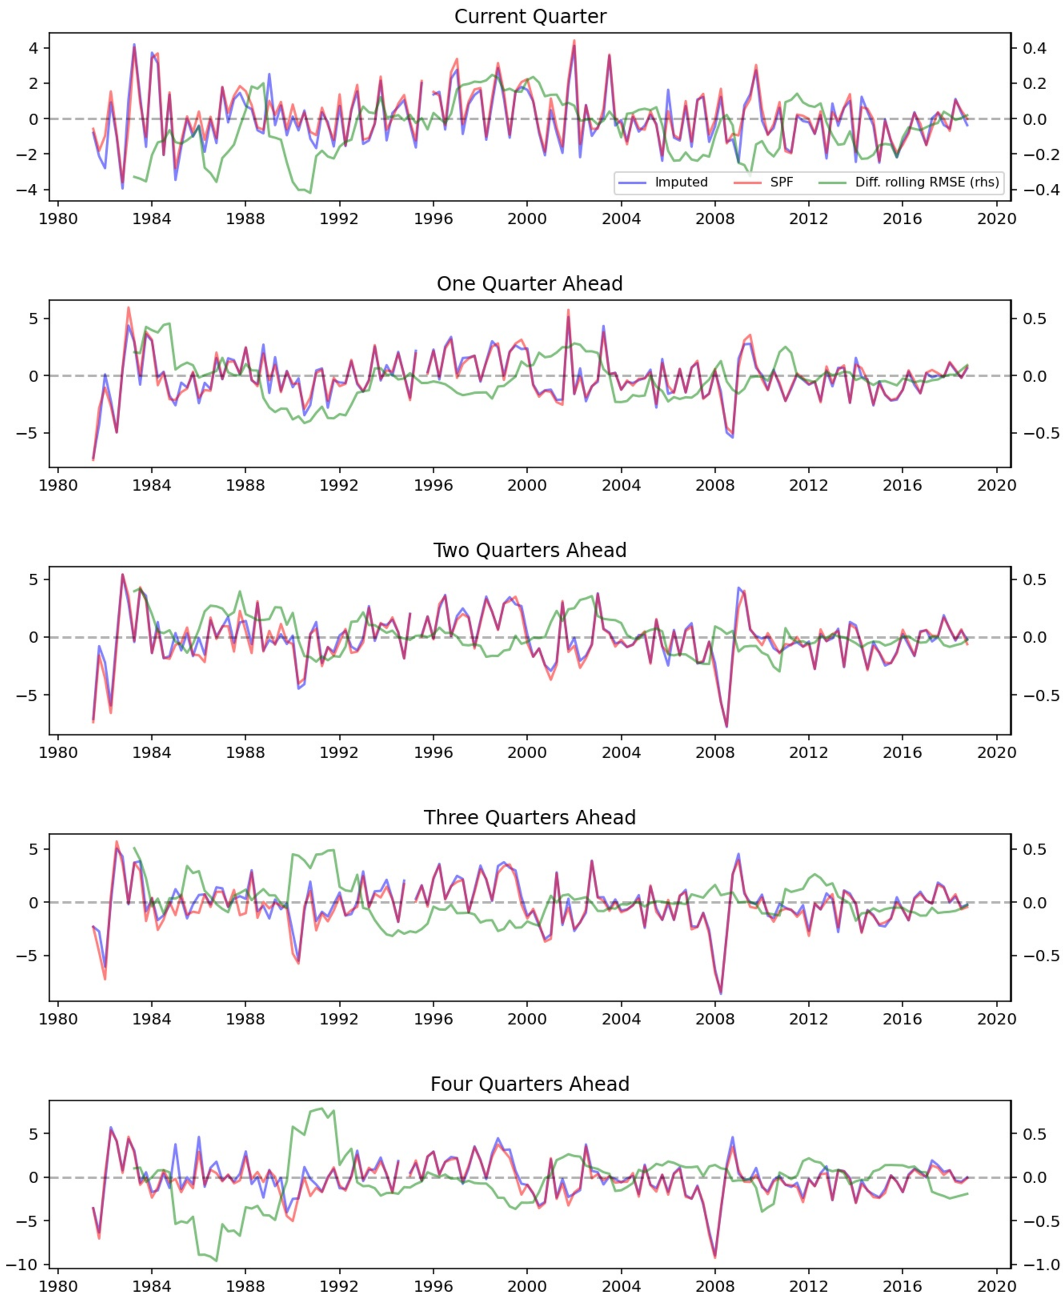

Notes: This figure is analogous to Figure 2 in the main text, with final GDP being replaced by the first GDP release.

Figure 11: Revision pattern including the 95th percentile and the maximum

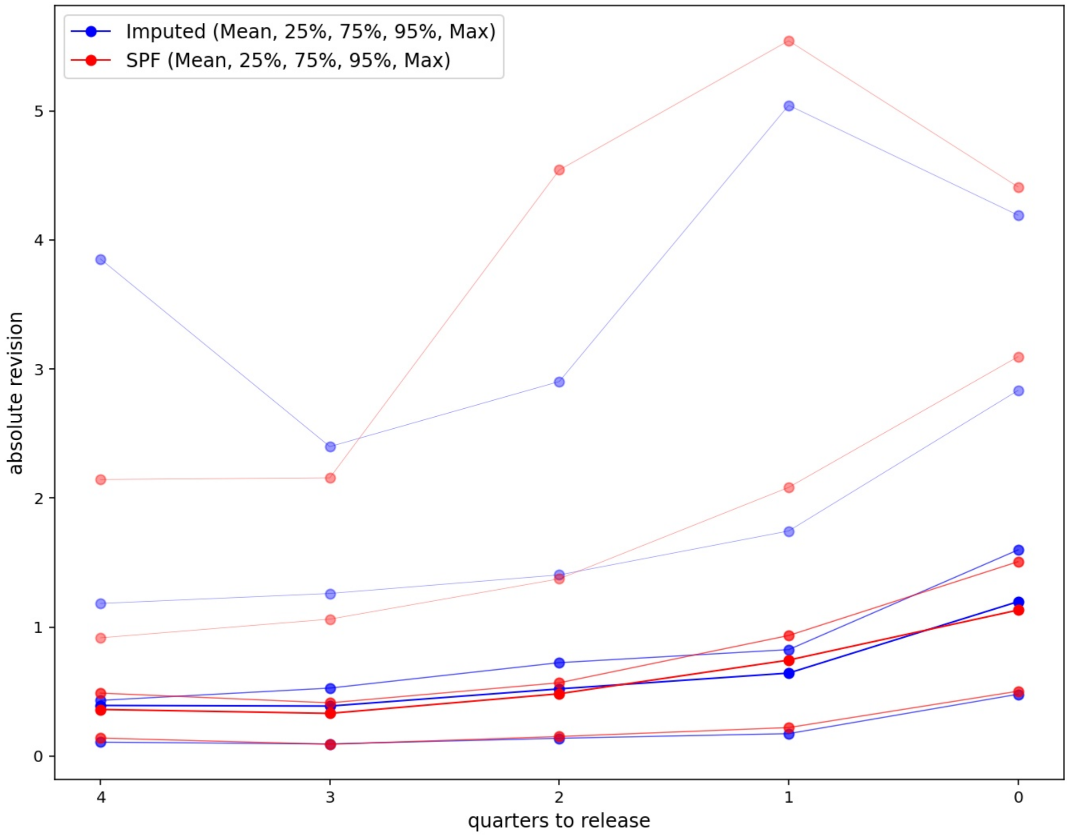

Notes: This figure is analogous to Figure 3, but it additionally presents the 95th percentile and the maximum of the distribution of the absolute forecast revision from one quarter to the next.

Figure 12: Quarterly revision pattern across economies

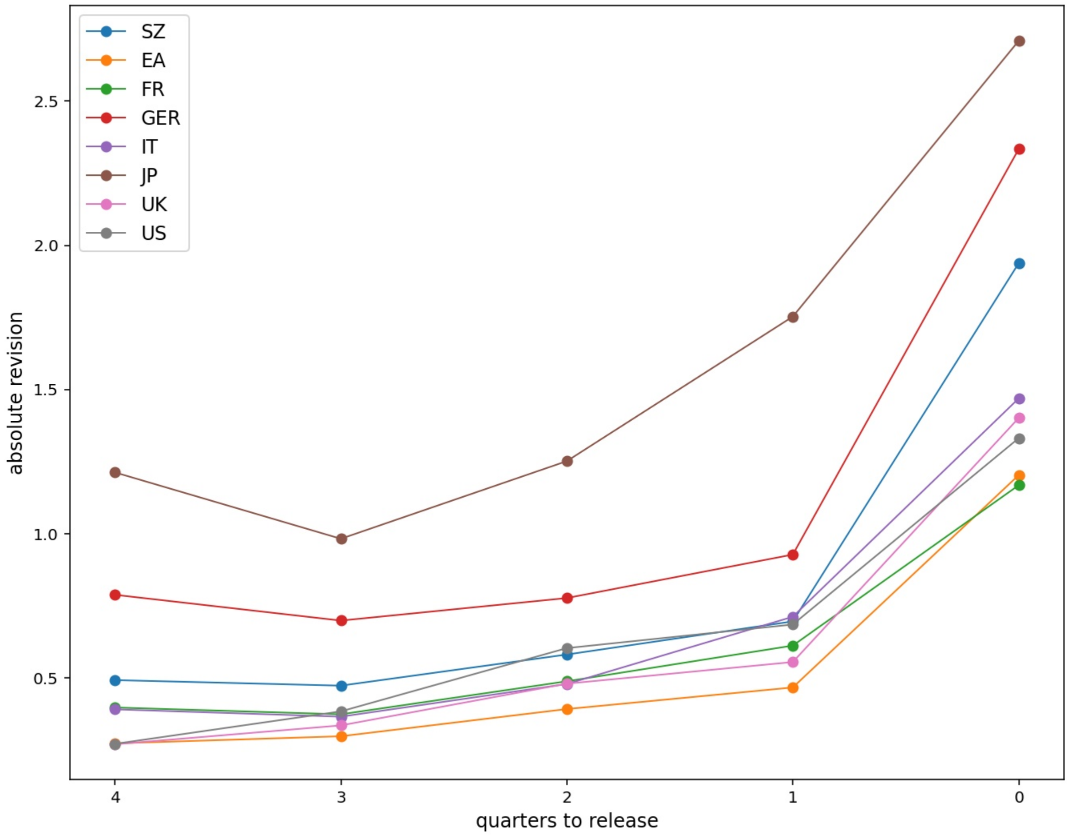

Notes: Each line represents the mean absolute forecast revision as the target quarter is approached. The figure is in quarterly frequency to smooth out the within-quarter seasonality present in Figure 9. The x-axis represents quarters prior to the end of the target quarter. The revisions are those that occur from the first month of a quarter to the first month of the next quarter. Time 0 represents the difference between the forecast made in the last month of the target quarter with the first GDP release.

Recent SNB Working Papers

2022-05 Jason Blunier, Christian Hepenstrick:
What were they thinking? Estimating the quarterly forecasts underlying annual growth projections

2022-04 Lena Lee Andresen:
The influence of financial corporations on IMF lending: Has it changed with the global financial crisis?

2022-03 Alain Galli, Rina Rosenblatt-Wisch:
Analysing households' consumption and saving patterns using tax data

2022-02 Martin Indergand, Eric Jondeau, Andreas Fuster:
Measuring and stress-testing market-implied bank capital

2022-01 Enrique Alberola, Carlos Cantú, Paolo Cavallino, Nikola Mirkov:
Fiscal regimes and the exchange rate

2021-20 Alexander Dentler, Enzo Rossi:
Shooting up liquidity: the effect of crime on real estate

2021-19 Romain Baeriswyl, Samuel Reynard, Alexandre Swoboda:
Retail CBDC purposes and risk transfers to the central bank

2021-18 Nicole Allenspach, Oleg Reichmann, Javier Rodriguez-Martin:
Are banks still 'too big to fail'? – A market perspective

2021-17 Lucas Marc Fuhrer, Matthias Jüttner, Jan Wrampelmeyer, Matthias Zwicker:
Reserve tiering and the interbank market

2021-16 Matthias Burgert, Philipp Pfeiffer, Werner Roeger:
Fiscal policy in a monetary union with downward nominal wage rigidity

2021-15 Marc Blatter, Andreas Fuster:
Scale effects on efficiency and profitability in the Swiss banking sector

2021-14 Maxime Phillot, Samuel Reynard:
Monetary Policy Financial Transmission and Treasury Liquidity Premia

2021-13 Martin Indergand, Gabriela Hrasko:
Does the market believe in loss-absorbing bank debt?

2021-12 Philipp F. M. Baumann, Enzo Rossi, Alexander Volkmann:
What drives inflation and how? Evidence from additive mixed models selected by cAIC

2021-11 Philippe Bacchetta, Rachel Cordonier, Ouarda Merrouche:
The rise in foreign currency bonds: the role of US monetary policy and capital controls

2021-10 Andreas Fuster, Tan Schelling, Pascal Towbin:
Tiers of joy? Reserve tiering and bank behavior in a negative-rate environment

2021-09 Angela Abbate, Dominik Thaler:
Optimal monetary policy with the risk-taking channel

2021-08 Thomas Nitschka, Shajivan Satkurunathan:
Habits die hard: implications for bond and stock markets internationally

2021-07 Lucas Fuhrer, Nils Herger:
Real interest rates and demographic developments across generations: A panel-data analysis over two centuries

2021-06 Winfried Koeniger, Benedikt Lennartz, Marc-Antoine Ramelet:
On the transmission of monetary policy to the housing market

2021-05 Romain Baeriswyl, Lucas Fuhrer, Petra Gerlach-Kristen, Jörn Tenhofen:
The dynamics of bank rates in a negative-rate environment – the Swiss case

SCHWEIZERISCHE NATIONALBANK
BANQUE NATIONALE SUISSE
BANCA NAZIONALE SVIZZERA
BANCA NAZIUNALA SVIZRA
SWISS NATIONAL BANK