Bank of Israel **Research Department**

Growth at Risk: Forecast Distribution of GDP Growth in Israel

Michael Gurkov* and Osnat Zohar**

Discussion Paper 2022.08
February 2022

Bank of Israel - http://www.boi.org.il

* Research Department, Bank of Israel - Email: Michael.Gurkov@boi.org.il
** Research Department, Bank of Israel, email: Osnat.Zohar@boi.org.il
We thank Roman Lafarguette for his guidance in the preliminary steps of this work. We also thank Alex Ilek, Alon Binyamini, Yaacov Chen-Zion, Itamar Caspi, participants of the Bank of Israel's Research Department seminar, and participants of the IMF's MCM Policy Forum for their helpful comments.

חטיבת המחקר, בנק ישראל ת״ד 780 ירושלים 91007
Research Department, Bank of Israel. POB 780, 91007 Jerusalem, Israel

Growth at Risk: Forecast Distribution of GDP Growth in Israel

Michael Gurkov and Osnat Zohar

Abstract

We estimate the distribution of future GDP growth given current financial and macroeconomic conditions. The distribution is generally symmetric, indicating that upside and downside risks are balanced. Its dispersion, which captures forecast uncertainty, rises when the median forecast decreases. The model allows us to study the connection between financial variables and GDP growth. We find that accommodative financial conditions, either in the local or the global economy, contribute to downside risks to growth within three years.

Keywords: downside risks, macro-financial linkages, volatility paradox.
JEL Classification: E17, E32, E44, G1.

שימוש במתודלוגית "צמיחה בסיכון" לאמידת התפלגות תחזית הצמיחה בישראל

מיכאל גורקוב ואסנת זהר

תקציר

אנו אומדים את התפלגות הצמיחה העתידית של התוצר כתלות בתנאים הפיננסיים והמקרו-כלכליים הנוכחיים. ההתפלגות הנאמדת בדרך כלל סימטרית, כלומר הסיכונים לצמיחה לרוב מאוזנים. רוחב ההתפלגות, אשר תופס את מידת אי הוודאות בתחזית, נוטה לעלות כאשר התחזית החציונית יורדת. המודל מאפשר בחינה של הקשר בין התנאים הפיננסיים לצמיחת התוצר ולסיכונים לצמיחה. אנו מוצאים כי תנאים פיננסיים מרחיבים, במשק המקומי או העולמי, מגבירים את הסיכון לצמיחה מתונה בטווח של שלוש שנים.

1 Introduction

Assessing risks to future economic activity and understanding their sources is essential for making knowledgeable forecasts. One of the sources whose importance became apparent since the 2008 global crisis is financial vulnerabilities. Considering the linkages between financial conditions and the real economy can help provide a comprehensive assessment of risks to GDP growth. Such an assessment is particularly critical for conducting monetary policy as it is forward-looking. For example, in their meetings, members of the Monetary Committee of the Bank of Israel often refer to uncertainty regarding GDP growth, its sources, and the balance between upside and downside risks.

The purpose of this paper is to provide a methodological tool for assessing risks to growth. To this end, we employ the "Growth at Risk" (GaR) methodology proposed by Adrian et al. (2019). GaR exploits interlinkages between financial and macroeconomic variables to estimate the distribution of future GDP growth. It is based on quantile regressions, which formulate the distribution of future growth as a linear function of current macro-financial conditions.

To motivate this approach, Figure 1 shows a scatter plot of year-over-year GDP growth four quarters ahead against current GDP and credit growth. It also shows the univariate quantile regression lines for the fifth, fiftieth, and ninety-fifth quantiles. The figure illustrates that the distribution of future GDP growth depends on current conditions. While the median forecast of GDP growth is not significantly correlated with either variable, the lower and upper quantiles are. Specifically, growth of either GDP or credit increases the ninety-fifth quantile of future GDP growth and decreases the fifth quantile, indicating a rise in forecast uncertainty.

The estimated GaR model for the Israeli economy indicates an essential contribution of financial variables to downward risks to growth. We find that an increase in asset prices and credit volume is associated with elevated downside risks to GDP growth within twelve quarters. High asset prices and credit volume are indicative of accomadative financial conditions that may stem from low financial risks. The finding that they raise risks in the medium term is consistent with of a phenomenon known in the literature as the "volatility paradox" (Brunnermeier and Sannikov, 2014). The theory postulates that low financial risks contribute to short-term growth in real activity. However, they also raise asset prices and credit and thus contribute to the accumulation of medium-term risks, raising the probability of subdued growth in the future. Several papers find evidence of an adverse effect of credit and asset prices on risks to growth in the medium term (Claessens et al., 2012; Borio, 2014; Mian et al., 2017). Aikman et al. (2018, 2019) find such evidence in GaR models similar to ours for the UK and other developed economies.

Since Israel is a small open economy, we also examine the effect of foreign financial conditions

Figure 1: Univariate Quantile Regressions of Year-over-Year GDP Growth Four Quarters Ahead (1996-2019)

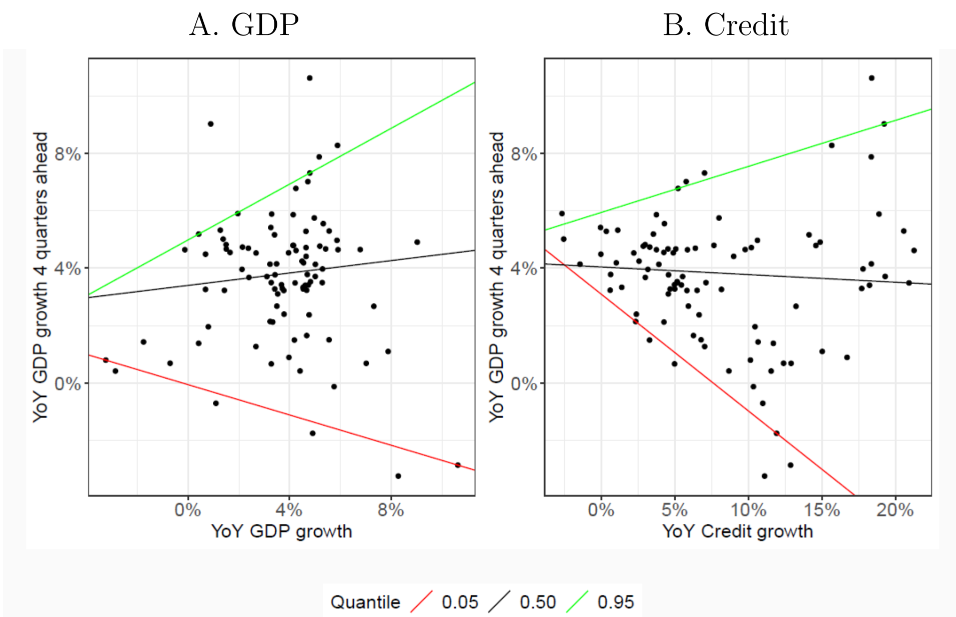

Notes: The figure shows scatter plots of year-over-year real GDP growth four-quarter-ahead against current year-over-year growth of GDP (Panel A) and credit (Panel B). It also shows fitted lines of the univariate quantile regressions for quantiles 0.05, 0.5, and 0.95.

on local growth. We find that accomadative financial conditions abroad stimulate local growth in the short term but increase downside risks twelve quarters ahead (similarly to local financial conditions). The importance of foreign financial conditions to local real growth can be explained by the high degree of openness to trade and financial flows in Israel. It might be that a volatility paradox that occurs abroad permeates the local economy through real channels, such as demand for Israeli exports. Alternatively, it may be that the permeation occurs through financial channels. Namely, accommodative financial conditions abroad contribute to the easing of local financial conditions, which, in turn, stimulate the real economy but also contribute to the accumulation of risks.

In addition to studying macro-financial linkages, the GaR approach allows us to characterize the distribution of future GDP growth. We examine three elements of the distribution. First, we examine downside risks, which refer to the distribution's left tail, and focus on depressed growth scenarios. To examine them, we look at the fifth percentile of the distribution. We find that in the years 2012-2019, downside risks were moderate as the fifth percentile of the distribution was stable at 0-2%. Furthermore, we find that short-term downside risks are mainly driven by global conditions, while medium-term risks are affected by local conditions.

Second, we examine forecast uncertainty, which is defined as the dispersion of the distribution. We proxy uncertainty by the inter-quartile range of the distribution (the gap between the seventy-fifth and the twenty-fifth quantiles). We find that up to eight quarters ahead, the interquartile range tends to rise when the median forecast falls. Namely, lower growth prospects are associated with higher uncertainty.

Third, we look at the balance of upside and downside risks by examining the distribution's skewness. We find that the distribution of growth is generally symmetric. This finding is unique for the Israeli economy, as GaR models for other countries show asymmetries in the distribution (Adrian et al., 2019; Alessandri et al., 2019). We explain this finding by noting that GDP growth was symmetric in our sample compared to other countries that experienced occasional sharp declines in activity and gradual recoveries. The fact that the model captures this central characteristic of growth supports its validity.

The GaR methodology allows some discretion in the selection of explanatory variables. While our baseline specification follows the common practice in the GaR literature, we test the sensitivity of our results to variable choice. Using a sub-sampling bootstrap procedure (Politis and Romano, 1994), we find that the qualitative characterizations of the distribution are robust: (1) the median forecast is negatively correlated with the distribution's dispersion up to eight quarters ahead; (2) the distribution is generally symmetric; (3) expansionary financial conditions are associated with elevated downside risks within three years.

We evaluate the model's forecast performance as it would have been used in real-time by generating out-of-sample forecasts on an expanding window. The model's performance is compared to two alternative benchmarks: (1) out-of-sample forecasts of a restricted model including only an intercept; (2) the density forecast derived from the Bank of Israel's DSGE model (Argov et al., 2012). The DSGE model is currently the primary tool for assessing risks to GDP growth in the Bank of Israel. However, it was re-estimated in 2019, so its forecasts are in-sample, giving it some advantage over the other two models.

We evaluate the predictive quality of these forecasts using two measures designed to assess density forecasts. First, following Adrian et al. (2019), we analyze the fraction of observations that fall below each quantile (the empirical cumulative distribution of the probability integral transform). According to this measure, the GaR model weakly outperforms the other two models in most horizons and quantiles. Second, we use the quantile R-squared score proposed by Giglio et al. (2016). Admittedly, the GaR model's relative performance is less favorable when assessed with this measure. It outperforms the DSGE model at horizons of eight and twelve quarters, especially in the distribution tails, but generally performs similarly to the intercept benchmark.

Nonetheless, the GaR model has important advantages over the other two benchmarks. It provides a state-dependent assessment of risks and allows analysis of their sources. A simple model with an intercept captures the historical distribution of growth and provides a risk assessment that is independent of the current state. The DSGE's forecast distribution is based on standard deviations of shocks in the model. In that sense, the implied assessment of risks represents average historical risks. In contrast, the GaR model uses historical data to determine elasticities, but the real-time forecast and risk assessment depend on current macro and financial conditions.

The rest of the paper is organized as follows. Section 2 describes the GaR methodology; Section 3 describes the data and presents estimation results; Section 4 discusses the role of the financial variables in our model; Section 5 analyzes the in-sample properties of the distribution of future GDP growth; Section 6 tests the sensitivity of these findings to variable selection; Section 7 tests the out-of sample forecast performance of the model.

2 Methodology

The GaR model generates forecasts for GDP growth. However, the emphasis is not on a point estimate but rather on the distribution of future growth. The model was proposed by Adrian et al. (2019), and the International Monetary Fund developed a methodology and infrastructure for its use (Prasad et al., 2019). The methodology consists of three steps:

1. **Variable partition:** the collection of explanatory variables is divided into a small number of groups.

2. **Dimension reduction:** a single common factor is extracted from each group of variables using the principal component method.

3. **Quantile regressions:** phase 2 factors are used as explanatory variables in quantile regressions for GDP growth at different horizons: For each quantile $\tau = 0.05$,0.25,0.50, 0.75,0.95, and forecast horizon $h = 1, 4, 8, 12$, we denote the τ-quantile of year-over-year GDP growth in period $t + h$ by $Q^{\tau}_{t,h}$. The estimated model is:

$$Q^{\tau}_{t,h} = \beta^{\tau}_{h} X_t + \epsilon^{\tau}_{t,h}, \tag{1}$$

   where X_t is the vector of factors from step 2 (including an intercept). Details about the estimation method of quantile regressions appear in the Appendix.[1]

[1] Due to the linear structure of quantile regressions, the estimated quantiles may sometimes cross. Namely, there may be a period t and forecast horizon h in which the quantile estimators $\hat{Q}^{\tau}_{t,h} = \hat{\beta}^{\tau}_{t,h} X_t$ are not monotone in τ. We address this issue, which is known as "quantile crossing", using to the method of Chernozhukov et al. (2010).

3 Data and Estimation Results

For our application to Israel, we use 17 macro and financial variables, which we divide into four groups (Table 1). Our rationale is to differentiate between macro and financial variables and between domestic and external variables. The groups are defined as follows:

1. **Domestic Macro:** the group contains the domestic uses of output (private consumption, public consumption, and real investment), together with the quarterly change in the unemployment rate.

2. **Domestic Financial Conditions:** a group that represents the local financial conditions. It includes the Bank of Israel's policy rate, which is the fundamental determinant of financing costs. It also includes asset prices (house and stock prices) and the private sector's credit.

3. **External Macro:** the group contains growth in Israeli exports and imports, as well as in OECD imports – the main trading partners of Israel.

4. **External Financial Conditions:** the group includes asset prices and interest rates in the US and euro area, together with commodity prices.

We estimate the first principal components on quarterly data between 1990Q1-2019Q4. Not all the series go back to 1990, but we use the longest available sample to exploit all the information about the variables' common movement.[2] Panel B in Table 1 shows the first principal component of each group. The Appendix shows the loadings of each factor and the share of the sample variance explained by it.

The quantile regressions are estimated on the common sample of our data set, which is 1996Q1-2019Q4. Figure 2 shows the estimated coefficients, for each forecast horizon. The Domestic Macro factor includes the larger share of output. Thus, it is not surprising that its coefficients at the one-quarter horizon are positive, as our model predicts year-over-year GDP growth which, by construction, shows high persistence from one quarter to the next. However, at the four-quarter horizon, the coefficients are negative, indicating that GDP growth is mean-reverting. At the longer horizons (eight and twelve quarters), the coefficients are positive in the left tail and negative in the right tail. Namely, a rise in the Domestic Macro factor is associated with moderation in both upside and downside risks.

The external factors, both macro and financial, seem to have a substantial and significant effect on GDP growth. This result is in line with the fact that Israel is a small open economy, characterized by a high level of openness to trade and financial flows.

[2]We found that using the shortest common sample of all variables in this stage hampers the out-of-sample forecast performance of the model.

Table 1: Explanatory Variables and Partition to Groups

A. Variables and Partition to Groups

Domestic Macro	Domestic Financial Conditions	External Macro	External Financial Conditions
Private cons.	BOI rate	Exports	Federal funds rate
Real investment	Credit	Imports	Eonia
Public cons.	House price index	OECD imports	S&P500
Unemployment	TA125		Eurostoxx600
			Brent oil price
			Non-energy cmdty

B. First Principal Components

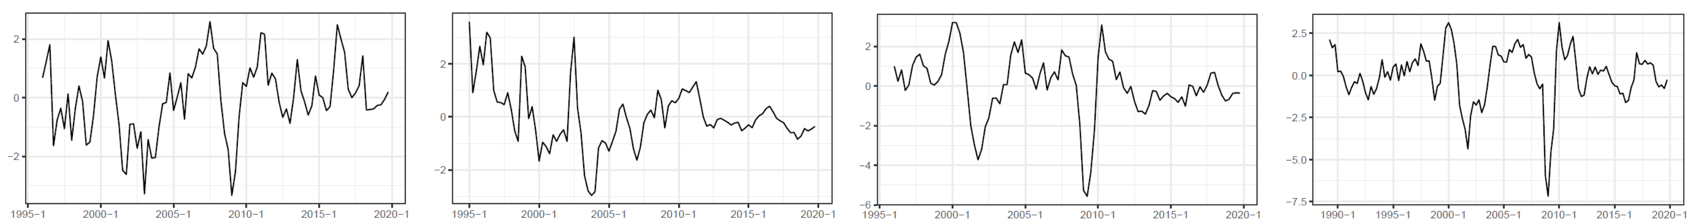

Notes: Panel A lists the explanatory variable divided into four groups. We converted daily and monthly data to quarterly averages. Subsequently, non-stationary series were converted to rates of change or first differences (see Table 3 in the Appendix). According to the common practice, all series were normalized before the first common factor was estimated. These factors are depicted in Panel B.

4 Financial Vulnerabilities

Focusing on the two financial conditions factors in Figure 2 (external and domestic), we see that they negatively affect growth at the twelve-quarter horizon in the left tail of the distribution. Namely, a rise in either of these factors, which captures easing financial conditions, intensifies downside risks in the medium term. These findings are consistent with a phenomenon known in the literature as the "volatility paradox" (Brunnermeier and Sannikov, 2014). It refers to a process in which low risks contribute to short-term growth but also stimulate the accumulation of risks in the medium-to-long term.

We find two manifestations of the volatility paradox in our setting, one stemming from financial conditions abroad and the other from local conditions. The first manifestation is through the External Financial Conditions factor. A rise in this factor is indicative of expansionary financial conditions, consistent with a lower risk environment. Indeed, up to four quarters ahead, this has a stimulating effect on local growth, as indicated by the positive coefficients. However, at longer horizons, the coefficients become negative, especially in the left tail of the distribution. Namely, an environment of low risks today enhances risks in the medium term, consistent with the volatility paradox theory.

To understand how financial conditions abroad affect the local economy, keep in mind that Israel is a small economy characterized by a high degree of openness to trade and financial flows. Thus, it might be that the coefficients of the External Financial Conditions capture a volatility paradox that occurs abroad and permeates the local economy through real channels, such as demand for Israeli exports. Alternatively, it may be that the permeation occurs through financial channels. Namely, accommodative financial conditions abroad contribute to the easing of local financial conditions, which, in turn, stimulate the real economy but also contribute to the accumulation of risks.

Direct evidence of a "local volatility paradox" can be found in the Domestic Financial Conditions factor. A rise in this factor enhances risks in the long term, as indicated by the negative coefficients of Domestic Financial Conditions at the left tail of the distribution twelve quarters ahead. This finding is consistent with several papers that find that growth in credit or asset prices enhances the severity of worst-case growth scenarios at the medium-term (Claessens et al., 2012; Borio, 2014; Mian et al., 2017). Admittedly, we find no evidence that domestic financial conditions have a stimulating effect on growth in the short term, but this could be due to the high correlation between local and foreign financial conditions. Namely, it may be that the positive effect in the short term is captured by the External Financial Conditions factor (recall that the coefficients capture the partial effect of each factor after controlling for the others).

Figure 2: Quantile Regression Coefficients

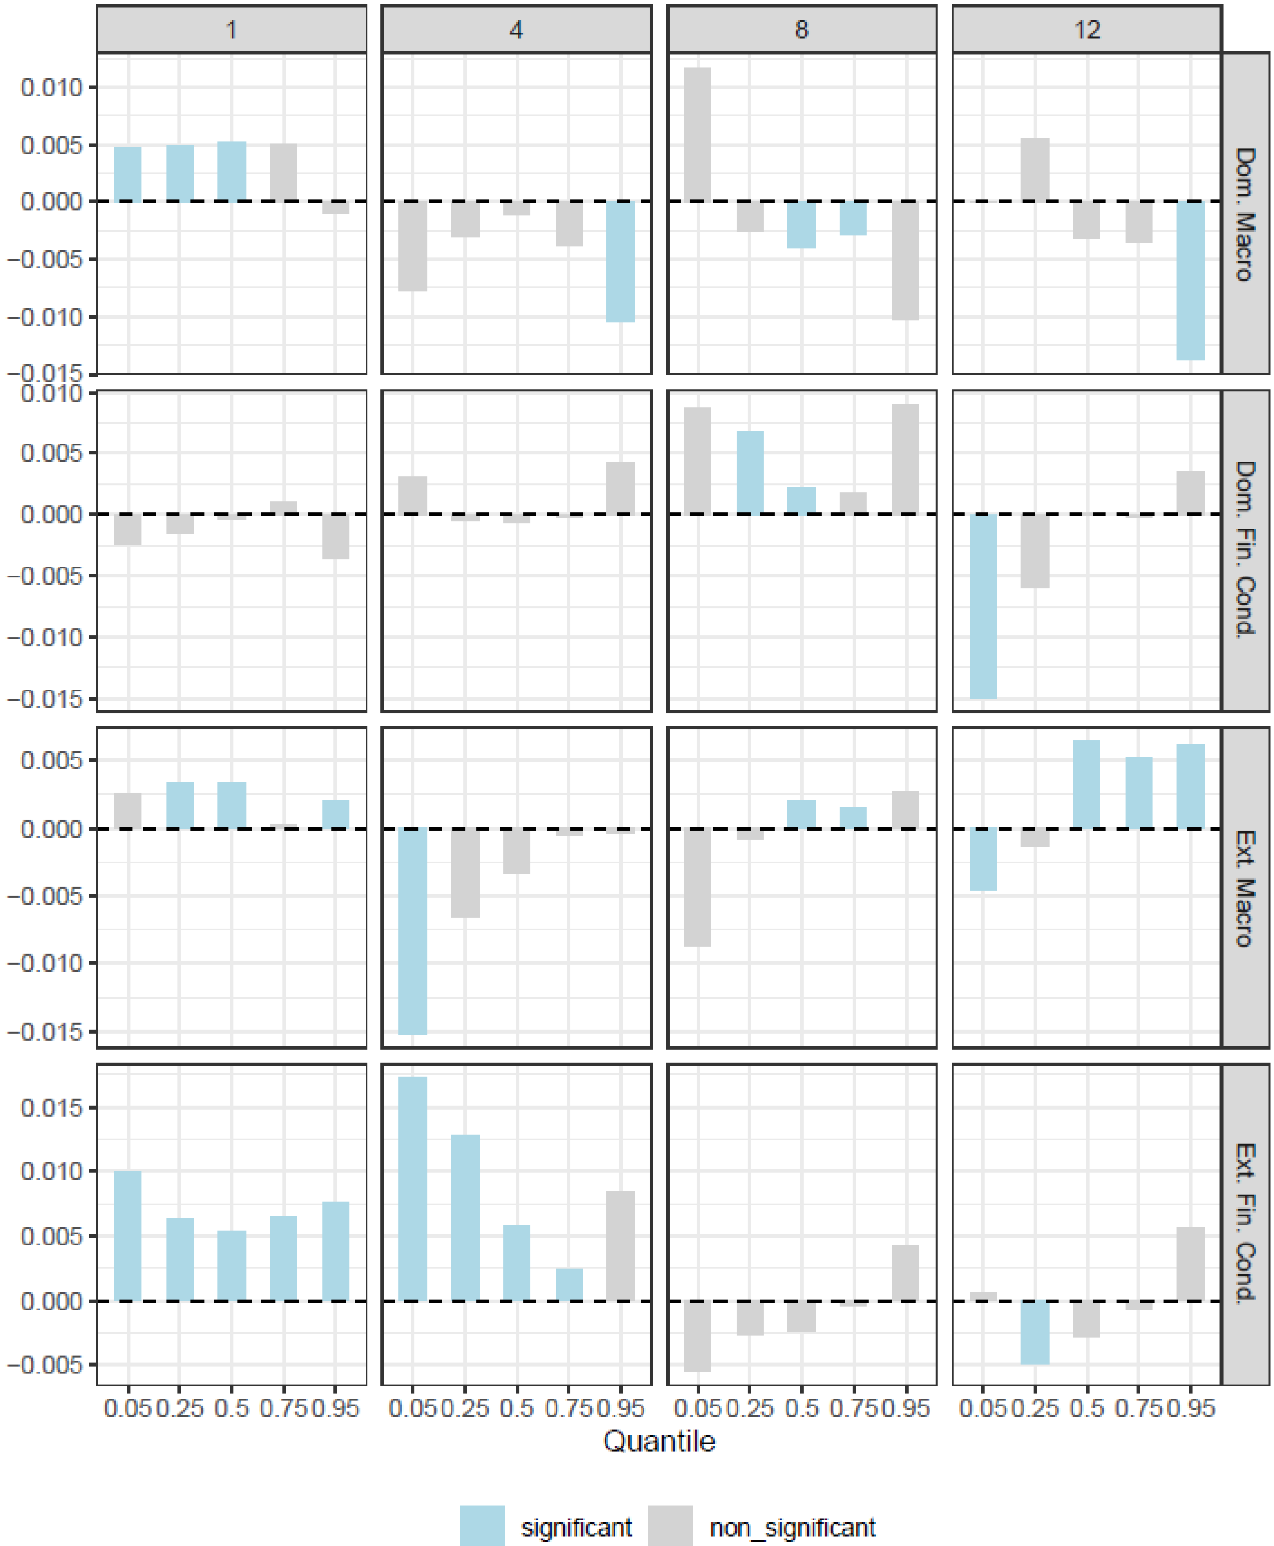

Notes: Each column shows the estimated coefficients from the quantile regression (1), at a specific forecast horizon. Blue bars mark coefficients that are significant at 10% (see Koenker, 1994).

Figure 3: Growth at Risk - Fifth Percentile Forecast One and Twelve Quarters Ahead

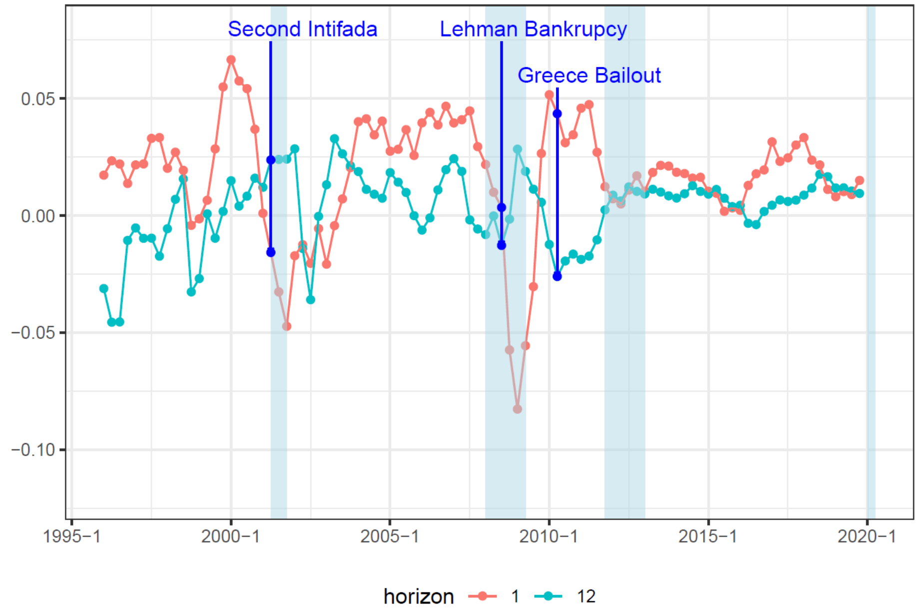

Notes: The figure shows the estimated 5th percentile of GDP growth in forecast horizons 1 and 12 (fitted values of $Q^{0.05}_{t,1}$ and $Q^{0.05}_{t,12}$ from Equation (1)). Shaded areas indicate recessions in the US or Euro area (NBER and CEPR contractions), and blue dots mark major crisis-related events.

5 Features of Risks to Growth and Uncertainty

In this section, we examine the characteristics of the distribution of GDP growth in our sample. Specifically, we study three elements of the distribution. First, to explore downside risks to growth, we focus on the forecast of the fifth percentile. Second, we examine forecast uncertainty by looking at the dispersion of the distribution. Third, we study the skewness of the distribution to test whether risks are generally balanced.

5.1 Downside Risks to Growth

The GaR model enables examining downside risk to growth by monitoring the forecast for the fifth percentile. Figure 3 shows the development of this forecast one and twelve quarters ahead ($\hat{Q}^{0.05}_{t,1}$ and $\hat{Q}^{0.05}_{t,12}$). The short horizon estimate is much more volatile than that of the long horizon, and it shows sharp decreases before and during crises. The forecast for downside risks three years ahead was stable at 0%-2% since 2012.

Figure 4 shows a breakdown of the factors' contributions to the fifth percentile forecast in each horizon. Up to four quarters ahead, external conditions (both macro and financial) contribute significantly to the development of the fifth percentile. At the longer horizons (eight and twelve quarters), the domestic factors become dominant in explaining risks to growth.

Figure 4: Decomposition of the Fifth Percentile

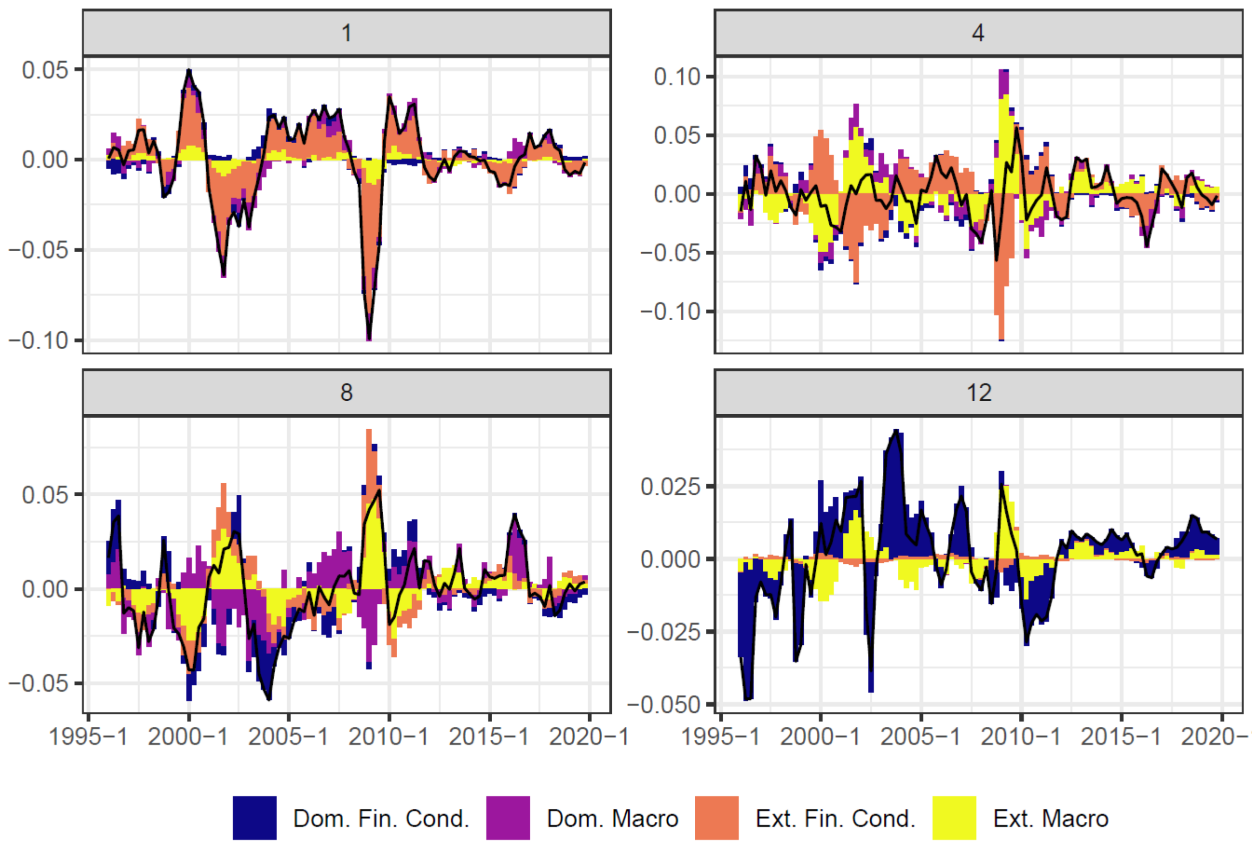

Notes: The figure shows a breakdown of the fifth percentile's forecast to the factors' contributions in each forecast horizon. In each panel, the black line shows the forecast's deviation from its mean, and the colored bars show the contributions to the deviations from the mean.

In particular, at the twelve-quarter horizon, the Domestic Financial Conditions factor is the predominant contributor, consistent with the view that it reflects an accumulation of long-term risks. It follows that the anchoring of long-term downside risks in recent years is due to stable interest rates and growth of credit and asset prices.

5.2 Uncertainty

To estimate forecast uncertainty, we examine the dispersion of the distribution. At each horizon h, our proxy for the dispersion is the interquartile range of the distribution:

$$IQR_{t,h} \equiv \hat{Q}^{0.75}_{t,h} - \hat{Q}^{0.25}_{t,h}.$$

Figure 5 shows the relationship between the median forecast $\hat{Q}^{0.5}_{t,h}$ and the dispersion proxy $IQR_{t,h}$. Up to eight quarters ahead, the relationship between the two variables is significantly

Figure 5: Connection between the Median and the Dispersion of the Predictive Distributions at Different Horizons

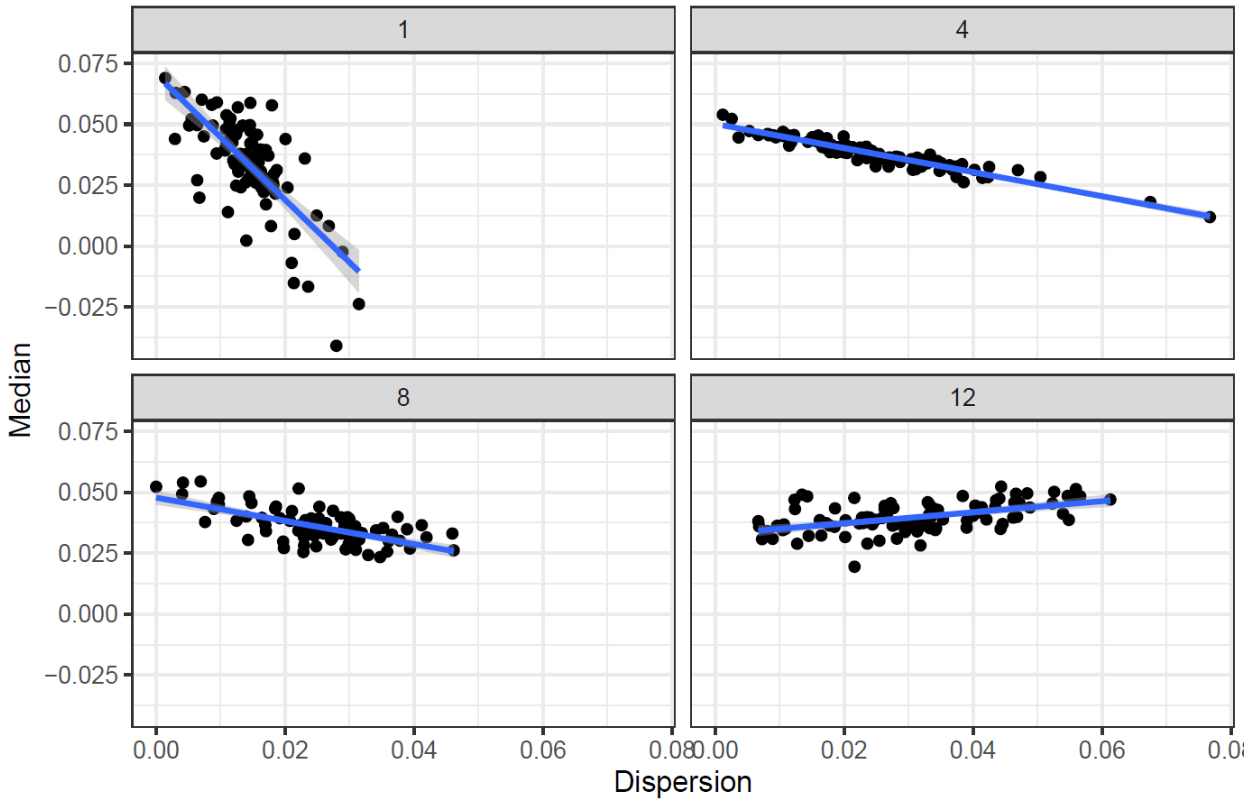

negative, and strongest one quarter ahead. That is, low median forecasts are associated with higher uncertainty.

5.3 Skewness

A natural question is whether upside and downside risks to growth are balanced. Namely, is the distribution generally symmetric or rather skewed? To test this, we examine the forecast quartile skewness.[3] Panel A of Figure 6 shows the development of this estimate, one quarter ahead. Throughout the sample, the skewness is low and not significant.[4] A similar picture arises for longer forecast horizons and using alternative estimates for skewness.[5] That is, the

[3]Quartile skewness is defined as $QS_{t,h} = \left(\frac{Q^{0.75}_{t,h}+Q^{0.25}_{t,h}}{2} - Q^{0.50}_{t,h}\right)\left(\frac{Q^{0.75}_{t,h}-Q^{0.25}_{t,h}}{2}\right)^{-1}$.

[4]The significance test should be taken with a grain of salt as it is based on the standard error of sample skewness ($SES = \sqrt{\frac{6N(N-1)}{(N-2)(N+1)(N+3)}}$, where N is the sample size), rather than a forecast distribution skewness. Another significance test for the skewness is conducted using bootstrap estimation in Section 6.

[5]We examined similar estimates based on the 5th and 95th percentiles. We also examined the gap between the mean forecast, estimated with OLS, and the median forecast $\hat{Q}^{0.50}_{t,h}$.

Figure 6: Skewness of the Forecast Distribution and Historical Growth

A. Quartile Skewness of the Forecast Distribution One-Quarter Ahead

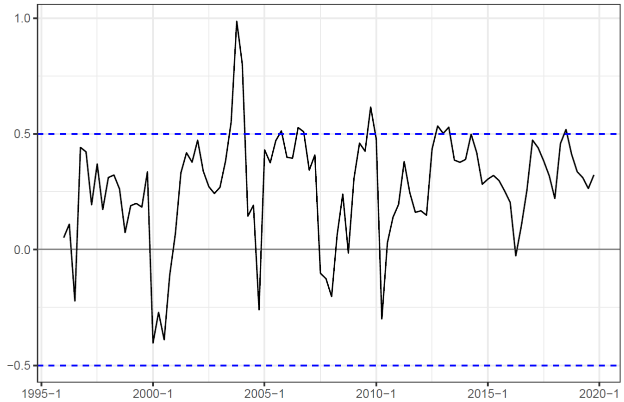

B. Actual Year-over-Year GDP Growth and Its Skewness - Israel, US, and Euro Area

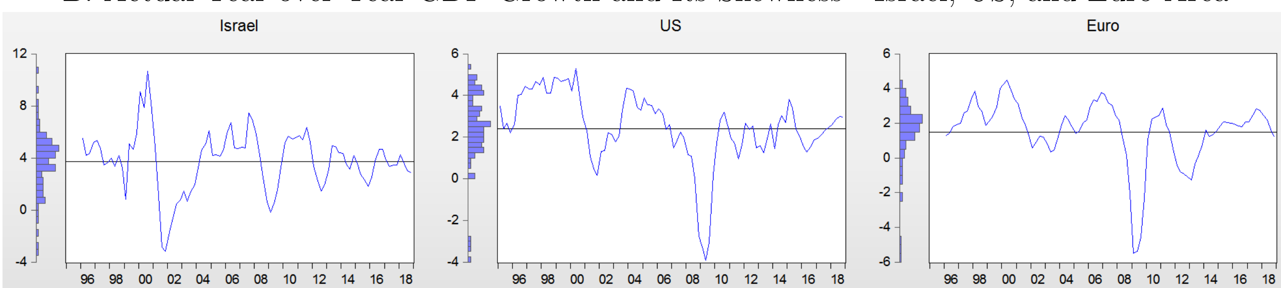

skew = -0.3 | skew = -1.4*** | skew = -1.8***

Notes: The significance of the estimates in Panels A and B is based on the SES estimate of the standard deviation of sample skewness (Footnote 4). In Panel A, the dashed lines represent $\pm 2SES$, and in Panel B, the significance test is based on this standard error.

model indicates that the distribution is generally symmetric. This finding is consistent with the fact that GDP growth in Israel was symmetric throughout the sample. The fact that the model captures the symmetry of the data supports our modeling choices and indicates that the model captures the central features of the true data generating process. It is interesting to note that the symmetry property is unique to the Israeli economy, as GDP growth in other economies, such as the US and the euro area, was negatively skewed in that period (Figure 6.B).

6 Sensitivity of the Results to Variable Selection

While our selection of explanatory variable follows the convention in the GaR literature, we acknowledge that our results might be sensitive to this choice. In this section, we perform a bootstrap procedure to test whether the results from Section 5 are robust to variable selection.

Table 2: Structure of Groups for the Bootstrap Analysis

Dom. Macro	Dom. Fin.	Ext. Macro	Ext. Fin.
A. Fixed Variables			
Private cons.	Credit	Exports	Federal funds rate
Real investment	BOI rate	Imports	Eonia
B. Alternating Variables			
Public cons.	House price index	OECD imports	S&P500
Unemployment	TA125		Eurostoxx600
			Brent oil price
			Non-energy cmdty

Notes: The table shows the division of the groups for the bootstrap procedure. Each iteration includes all the variables from Panel A and a combination of variables from Panel B.

In the spirit of Politis and Romano (1994), we repeat our analysis using different subsets of variables. However, we do not go over all subsets or draw them randomly, but instead, we put some structure on the sub-sampling procedure. First, we determine a minimal set of eight variables included in each iteration and go over all the different combinations of the other nine variables (Table 2). The reason for that is to rule out specifications that contain only immaterial variables. Second, we fix the partition to groups, so variables do not switch groups at different iterations. For example, each iteration that includes unemployment includes it in the Domestic Macro group. We impose this restriction to keep the interpretation of each factor and maintain the GaR framework, as our partition generally follows the common convention in the literature.

First, we show that our measure of downside risks is robust. In Figure 7, Panel A shows the mean series of fifth percentile forecasts generated by the sub-sampling procedure. The figure also shows the ninety percent coverage band, which is generally narrower than two percentage points.

Second, Panels B-D in Figure 7 show that three of our main results are robust to variable selection:

1. **Median-dispersion negative correlation:** Panel B shows the mean estimate of the correlation between the dispersion of the distribution and its median at each forecast horizon. The confidence interval covers ninety percent of the iterations. We find that the correlation is significantly negative up to eight quarters ahead, consistent with our

baseline results.

2. **Symmetric distribution:** Panel C shows the mean of the quartile skewness estimates, together with ninety percent coverage bands. Up to eight quarters ahead, the skewness measure is generally low and insignificant. Thus, we find no evidence that the distribution is significantly skewed in these horizons and conclude that the distribution is broadly symmetric. Admittedly, at the twelve-quarter horizon, the skewness shows several episodes of significant negative skewness.

3. **Financial vulnerabilities:** Panel D shows the mean coefficients of the financial factors at the twelve-quarter horizon, together with ninety percent coverage bands. The mean coefficients of both financial factors are negative at the left tail of the distribution (the domestic factor's effect is stronger and more significant than the external factor). These results confirm that expansionary financial conditions generate risks to growth in the medium term, as discussed in Section 4.

7 Forecast Accuracy

We evaluate the performance of the model as it would have been used in real-time. We generate out-of-sample forecasts for GDP quantiles on an expanding window, starting with a window size of 40 observations. Specifically, the t-period forecasts $\hat{Q}^{\tau}_{t,h}$ ($t \geq 40$) are generated by performing the methodology of Section 2 on the sample ending in period t.

Next, we evaluate the model's out-of-sample predictive quality using two measures designed to assess density forecasts. First, following Adrian et al. (2019), we analyze the fraction of observations that fall below each quantile. Second, we use the quantile R-squared score proposed by Giglio et al. (2016).

The first forecast evaluation is based on the probability integral transform (PIT). For each quantile τ and horizon h, we compute the empirical cumulative distribution of the PITs, which is the percentage of observations that fall below the forecast quantile $\hat{Q}^{\tau}_{t,h}$:

$$\varphi_{\tau,h} \equiv \frac{1}{T-h} \sum_{t=1}^{T-h} \mathbb{I}_{\left\{y_{t+h} \leq \hat{Q}^{\tau}_{t,h}\right\}},$$

where $\hat{Q}^{\tau}_{t,h}$ are forecasts of the GaR model estimated in period t, T is the sample size, and $\mathbb{I}$ is the indicator function. A model is better fitted the closer $\varphi_{\tau,h}$ is to the 45-degree line. If the model perfectly fits the empirical distribution, then the fraction of observations falling below quantile τ should be exactly τ, namely, $\varphi_{\tau,h} = \tau$.

Figure 7: Sub-Sampling Bootstrap Results

A. Fifth Percentile of the Forecast Distribution

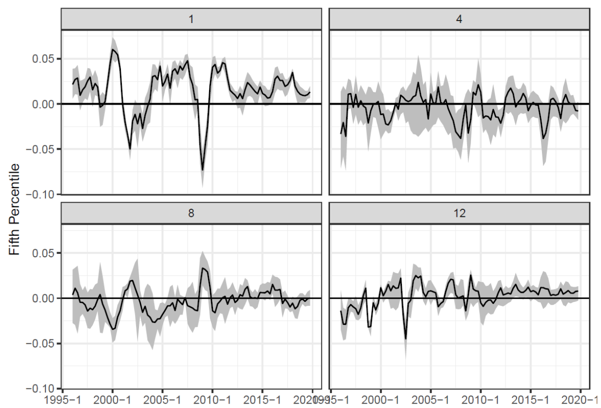

B. Correlation between the Median and the Dispersion of the Forecast Distribution

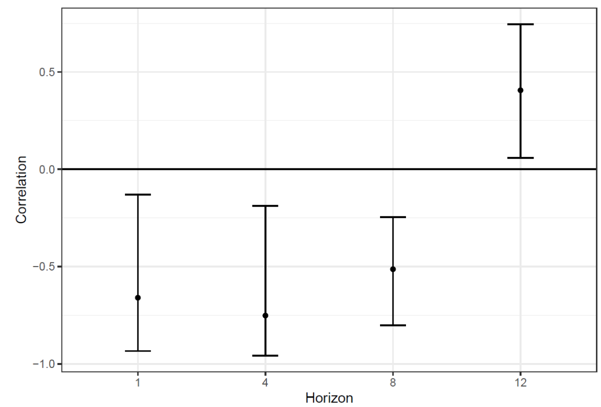

C. Quartile Skewness

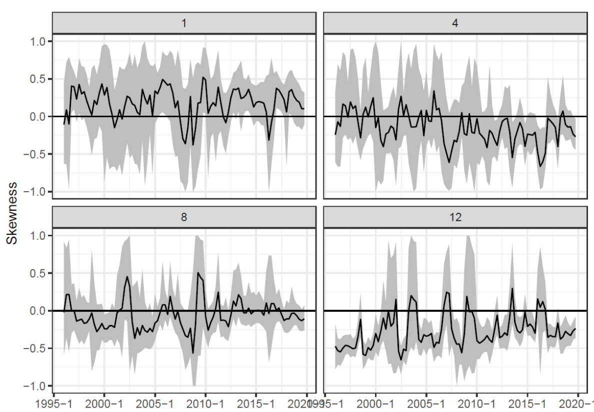

D. Financial Factors' Coefficients

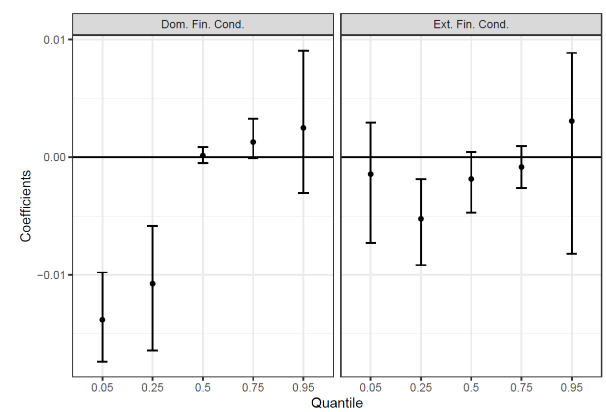

Notes: The figure shows results from a sub-sampling bootstrap with 512 iterations. Each iteration contains all the fixed variables listed in Table 2 and a combination of the other variables. In Panels A and C, the black line shows the mean estimate, and the grey area covers ninety percent of the iterations. In Panels B and D, the dots show the mean estimate, and the bars cover ninety percent of the iterations.

Panel A in Figure 8 shows the $\varphi_{\tau,h}$ scores of out-of-sample GaR forecasts. It also shows scores of two alternative benchmarks: (1) out-of-sample forecasts of a restricted model including only an intercept; (2) the density forecast derived from the Bank of Israel's DSGE model (Argov et al., 2012). The DSGE forecast distributions are symmetric, with a constant variance. Thus, the quantiles are at a fixed distance from the median forecast. One should note that the DSGE model was re-estimated in 2019, so its forecasts are in-sample, giving it some advantage over the other two models. Panel A in Figure 8 shows that the GaR model performs similarly or better than the other two models in most horizons and quantiles (the cumulative distribution of the PITs is closest to the 45-degree line).

The assessment based on the PITs gives a relatively rough estimate of the forecast performance, as it only addresses the question of whether realized growth fell above or below the forecast of a specific quantile. The second measure we use, quantile R-squared (Giglio et al., 2016), places more emphasis on the distance between the realized value y_{t+h} and the quantile forecast $\hat{Q}^{\tau}_{t,h}$, namely, the forecast error. It compares these errors to the ones generated by a restricted model containing only a intercept. For each quantile τ and forecast horizon h, we look at a weighted average of forecast errors $\left|y_{t+h} - \hat{Q}^{\tau}_{t,h}\right|$, and compare it to a similar weighted average of the forecast errors from the restricted model with an intercept:

$$R^2_{\tau,h} \equiv 1 - \frac{\sum_{t=1}^{T-h} \left(y_{t+h} - \hat{Q}^{\tau}_{t,h}\right)\left(\tau - \mathbb{I}_{\left\{y_{t+h} < \hat{Q}^{\tau}_{t,h}\right\}}\right)}{\sum_{t=1}^{T-h} \left(y_{t+h} - \hat{c}^{\tau}_{t,h}\right)\left(\tau - \mathbb{I}_{\left\{y_{t+h} < \hat{c}^{\tau}_{t,h}\right\}}\right)},$$

where $\hat{c}^{\tau}_{t,h}$ are the intercept forecasts. Note that forecast errors are weighted according to the quantile in question. For example, the fifth quantile R-squared weighs positive forecast errors by 0.05 and negative errors by 0.95, heavily penalizing realizations below the quantile forecast. Higher values of $R^2_{\tau,h}$ indicate better forecast performance than the intercept benchmark. Negative values indicate that the performance of the model falls short of the benchmark.

In Figure 8, Panel B shows R-squared scores of the GaR model compared to the DSGE model (since the R-squared is constructed relative to the intercept benchmark, it is excluded from the figure). Up to four quarters ahead, the GaR model is inferior to the DSGE model. However, at horizons eight and twelve, it outperforms the DSGE model, especially in the distribution tails. Admittedly, both models do not show substantially better performance than the intercept benchmark four to twelve quarters ahead, as indicated by the low and sometimes negative R-squared scores.

Although the GaR model's relative performance is less favorable when assessed with the R-squared scores than with the PIT scores, it has material advantages over the other two benchmarks. It provides a state-dependent assessment of risks and allows analysis of their

Figure 8: Evaluation of Forecast Accuracy

A. Probability Integral Transform (PIT)

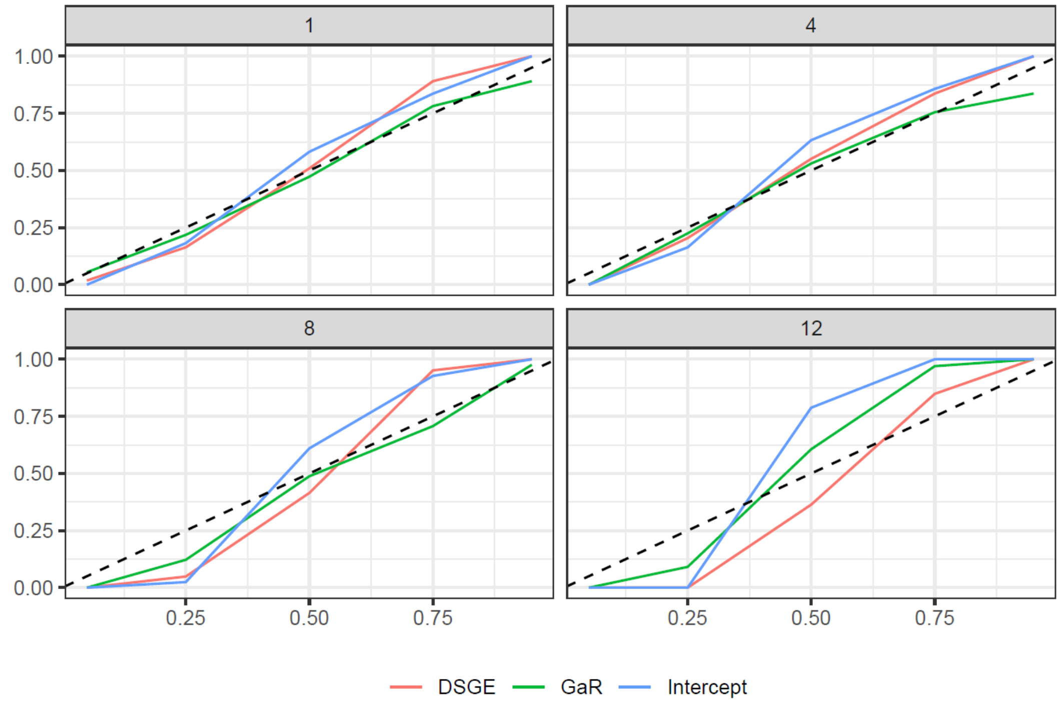

B. Quantile R-Squared

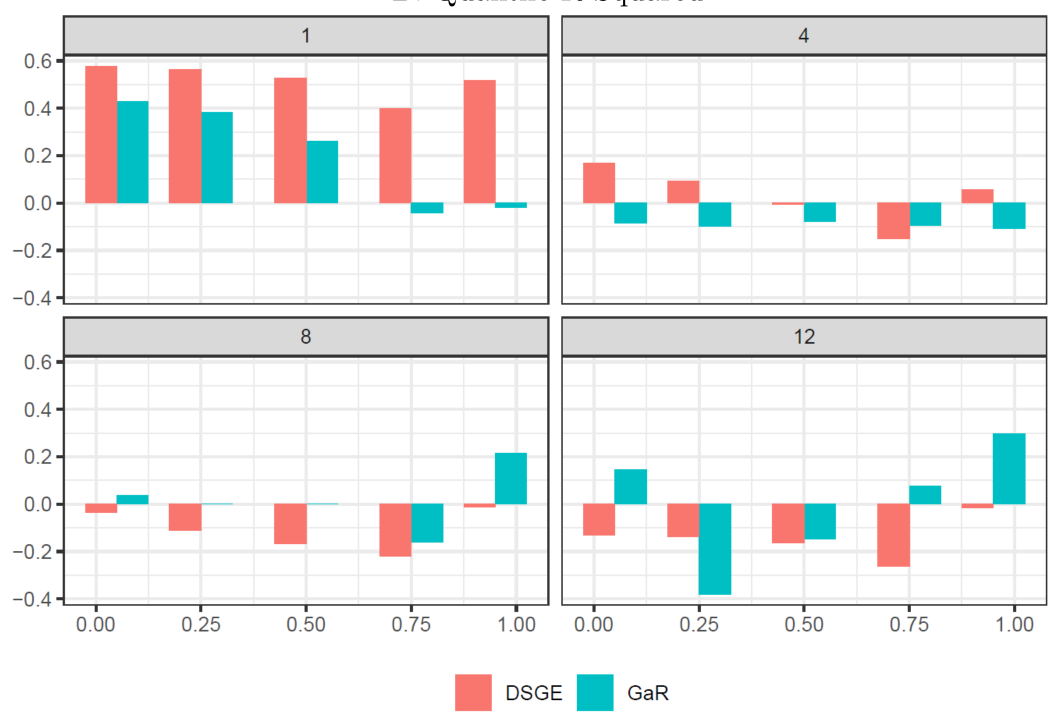

Notes: In Panel A, each figure depicts the empirical cumulative distribution of the probability integral transform (PIT) for a specific forecast horizon, i.e., the share of observations that fall below each forecast quantile. The 45-degree dashed line is added for reference. In Panel B, each figure depicts the quantile R-squared scores at a specific forecast horizon. In both panels, the GaR model scores appear alongside those of the Bank of Israel's DSGE model. Panel A also shows the performance of a model containing only an intercept, which is already inherent in the R-squared score.

sources. A simple model with an intercept captures the historical distribution of growth and so provides an assessment that is independent of the current state. The DSGE's density forecast is based on standard deviations of shocks which were estimated on historical data. In that sense, the implied assessment of risks represents average historical risks. In contrast, the GaR model uses historical data to determine elasticities, but the real-time forecasts and risk assessment depend on current macro and financial conditions.

8 Conclusion

We estimate quantiles of future GDP growth as a linear function of real and financial variables. We characterize the in-sample properties of the implied distribution and show that it is generally symmetric. However, the distribution is not constant over time. For example, up to eight quarters ahead, its dispersion increases when the median forecast decreases, indicating that elevated uncertainty is associated with lower expected growth. We also characterize links between financial variables and GDP growth. Consistent with previous findings in the literature, accommodative financial conditions are associated with elevated risks in the medium term. These results are robust to variable selection.

The GaR approach has several merits. First, quantile regressions' linear and parsimonious nature allows timely and coherent interpretation of the results, making the GaR approach very appealing for practical uses. For example, it allows regular assessment of downside risks by tracking the evolution of the fifth percentile of GDP forecasts. Second, it sets empirical foundations for modeling macro-financial interactions in Israel. The empirical findings in this paper can guide the design and estimation of structural macro-financial models in future research.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review 109*(4), 1263–89.

Aikman, D., J. Bridges, S. Burgess, R. Galletly, I. Levina, C. O'Neill, and A. Varadi (2018). Measuring risks to UK financial stability. Bank of England working papers 738, Bank of England.

Aikman, D., J. Bridges, S. Hacioglu Hoke, C. O'Neill, and A. Raja (2019, sep). Credit, capital and crises: a GDP-at-Risk approach. Bank of England working papers 824, Bank of England.

Alessandri, P., L. D. Vecchio, and A. Miglietta (2019, oct). Financial Conditions and 'Growth at Risk' in Italy. Temi di discussione (Economic working papers) 1242, Bank of Italy, Economic Research and International Relations Area.

Argov, E., E. Barnea, A. Binyamini, E. Borenstein, D. Elkayam, and I. Rozenshtrom (2012). Moise: A dsge model for the israeli economy. *Bank of Israel Discussion Paper*.

Borio, C. (2014). The financial cycle and macroeconomics: What have we learnt? *Journal of Banking & Finance 45*, 182–198.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104*(2), 379–421.

Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica 78*(3), 1093–1125.

Claessens, S., M. A. Kose, and M. E. Terrones (2012). How do business and financial cycles interact? *Journal of International economics 87*(1), 178–190.

Giglio, S., B. Kelly, and S. Pruitt (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics 119*(3), 457–471.

Koenker, R. (1994). Confidence intervals for regression quantiles. In *Asymptotic statistics*, pp. 349–359. Springer.

Mian, A., A. Sufi, and E. Verner (2017). Household debt and business cycles worldwide. *The Quarterly Journal of Economics 132*(4), 1755–1817.

Politis, D. N. and J. P. Romano (1994). The stationary bootstrap. *Journal of the American Statistical Association 89*(428), 1303–1313.

Prasad, M. A., S. Elekdag, M. P. Jeasakul, R. Lafarguette, M. A. Alter, A. X. Feng, and C. Wang (2019). *Growth at Risk: Concept and Application in IMF Country Surveillance.* International Monetary Fund.

Appendix

A Quantile Regressions

Quantile regressions are used to estimate the dependent variable's distribution as a linear function of the explanatory variables. In our case, for quantile τ and forecast horizon h, we estimate the model:

$$Q^{\tau}_{t,h} = \beta^{\tau}_{h} X_t + \epsilon^{\tau}_{t,h},$$

where $Q^{\tau}_{t,h}$ is the τ-quantile of year-over-year GDP growth in period $t+h$, and X_t is the vector of factors estimated in step 2 as described in Section 2 (the vector includes an intercept). The coefficients β^{τ}_{h} are estimated by solving the minimization problem:

$$\min_{\beta} \sum_{t=1}^{T-h} \left(y_{t+h} - \beta X_t\right) \left(\tau - \mathbb{I}_{\{y_{t+h} < \beta X_t\}}\right),$$

where y_{t+h} is GDP growth in period $t+h$ and $\mathbb{I}$ is the indicator function. For example, the coefficient of quantile 0.9 percent one quarter ahead ($\tau = 0.9, h = 1$) is estimated by measuring the distances of all points y_{t+1} from the line βX_t, where points above the line receive a weight of 0.9, and the rest receive a weight of 0.1. We are looking for a coefficient β that minimizes this sum.

B Additional Tables and Figures

Table 3: Summary Statistics

Variable	N	Mean	St. Dev.	Min	Max
GDP (YoY growth)	97	3.694	2.242	−3.245	10.632
Private consumption (YoY growth)	96	0.041	0.023	-0.015	0.112
Public consumption (YoY growth)	96	0.025	0.019	-0.028	0.064
Real investment (YoY growth)	96	0.030	0.059	-0.068	0.158
Unemployment rate (1st diff.)	138	-0.019	0.489	-1.614	1.399
Exports (YoY growth)	96	0.064	0.086	-0.173	0.287
Imports (YoY growth)	96	0.051	0.072	-0.160	0.188
OECD imports (YoY growth)	123	0.052	0.049	-0.162	0.154
Eonia (1st diff.)	102	−0.064	0.315	−1.799	0.701
Federal funds rate (1st diff.)	102	−0.055	0.394	−1.418	0.560
S&P500 (YoY growth)	99	0.086	0.160	−0.402	0.409
Eurostoxx600 (YoY growth)	99	0.062	0.191	−0.438	0.500
Brent oil price (YoY growth)	99	0.098	0.353	−0.525	1.296
Non-energy cmdty (YoY growth)	99	0.026	0.154	−0.335	0.446
BOI rate (1st diff.)	102	−0.161	0.781	−2.507	3.776
Credit (YoY growth)	98	0.078	0.059	−0.027	0.213
House prices (YoY growth)	98	0.045	0.064	−0.083	0.206
TA125 (YoY growth)	99	0.110	0.227	−0.462	0.814

Figure 9: Loadings of the First Principal Components and the Share of the Sample Variance Explained by Them

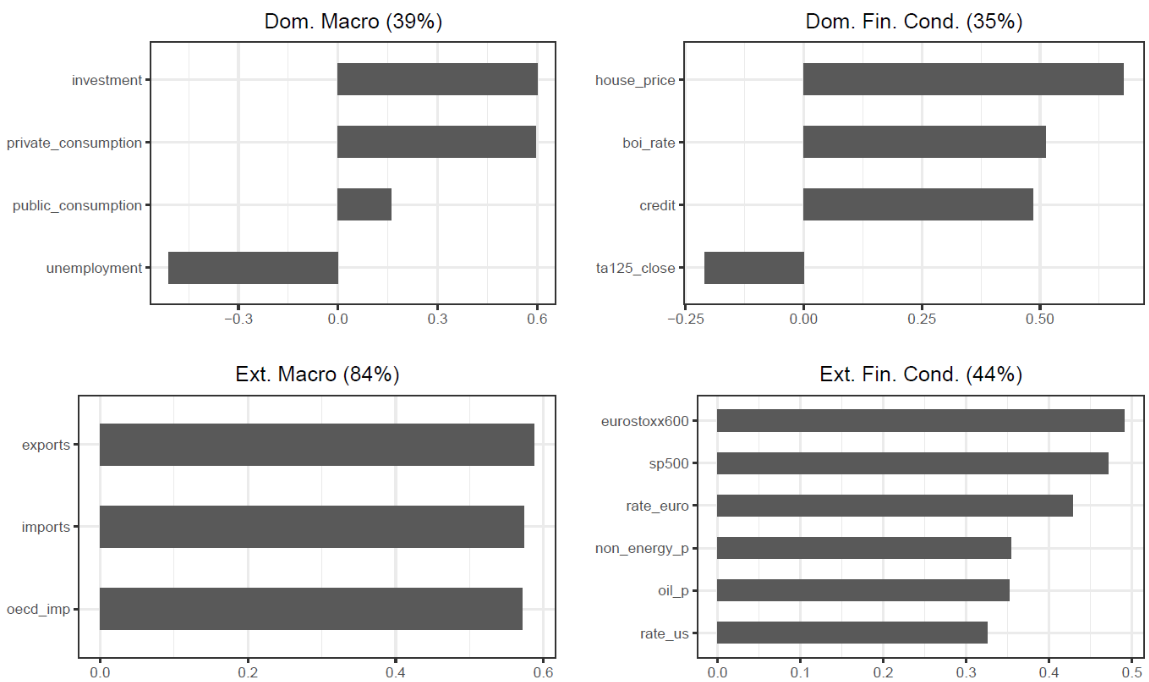

Note: Each bar plot depicts the loadings of the first principal component of the respective group. The share of the common sample variance explained by this factor is reported in parenthesis.

Figure 10: Quantile Forecasts Over Time

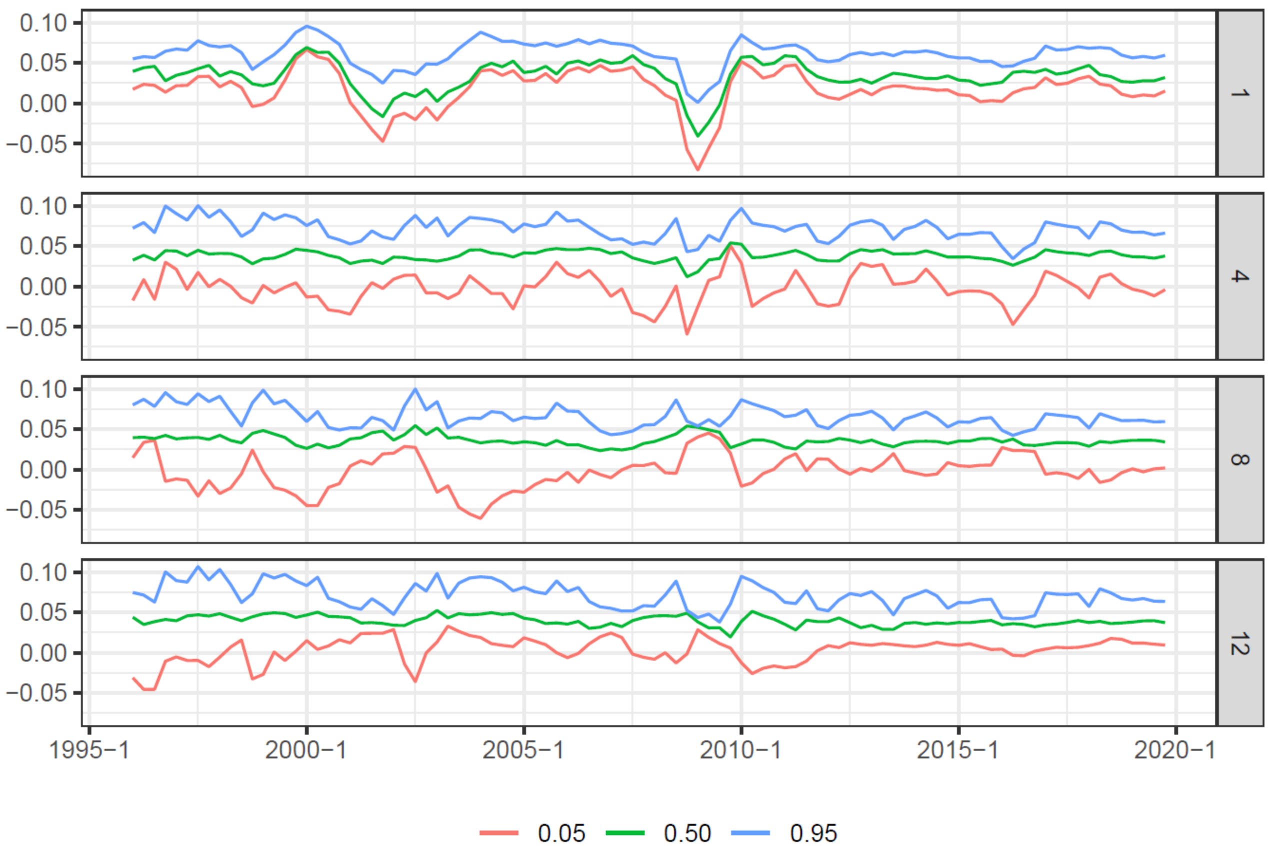

Notes: Each graph shows the in-sample forecasts for quantiles 0.05, 0.50 and 0.95, for a specific horizon.