Bank of Israel **Research Department**

Inflation Risks in Israel

Michael Gurkov* and Osnat Zohar**

Discussion Paper 2022.19
December 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 Alon Binyamini, Nadav Ben-Zeev and Sigal Ribon for their helpful comments.
The Hebrew version of this paper was published in the book "Monetary Policy in Times of Price Stability", Bank of Israel, July 2022.

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

Inflation Risks in Israel

Michael Gurkov and Osnat Zohar

Abstract

We examine how inflation risks in Israel evolved over time. We find that until 2013, inflation uncertainty was stable, and risks were moderately skewed downwards. However, since 2014, uncertainty decreased, and downside risks to inflation became much more dominant. The model attributes these developments to the decline in the inflation environment, as it is captured by realized inflation and long-term expectations, and to changes in oil prices. However, we cannot rule out that the monetary rate approaching the effective lower bound also contributed to these changes.

Keywords: inflation at risk, density forecasts, quantile regressions, effective lower bound.
JEL Classification: E31, E37, E58.

סיכוני האינפלציה בישראל

מיכאל גורקוב ואסנת זהר

תקציר

אנו בוחנים כיצד התפתחו סיכוני האינפלציה בישראל מאז אימוץ טווח יעד האינפלציה הנוכחי בעזרת מודל להתפלגות תחזית האינפלציה. אנו מוצאים כי עד שנת 2013 מידת אי הוודאות לגבי האינפלציה הייתה יציבה והסיכונים היו מוטים מעט כלפי מטה. עם זאת, מאז 2014 אי הוודאות פחתה אך הסיכונים כלפי מטה הפכו דומיננטיים יותר. המודל מייחס התפתחויות אלו לירידה בסביבת האינפלציה, כפי שהיא נתפסת על ידי האינפלציה בפועל והציפיות לטווח ארוך, וכן לשינויים במחירי הנפט. עם זאת, איננו יכולים לשלול כי גם התקרבותה של הריבית המוניטרית למחסום האפקטיבי שלה תרמה להתפתחויות אלו.

1 Introduction

The adoption of the current inflation target range in Israel (1–3%) in 2003 was followed by a decade of stable inflation with an average of about two percent (Figure 1). However, since the mid-2010s, inflation declined for several years, reaching a negative territory. As illustrated in Figure 1, the decline was accompanied by a decrease in the standard deviation of inflation. These changes in the statistical properties of inflation may indicate that the risks to inflation have changed.

The purpose of this paper is to evaluate how risks to inflation evolved over time and what fundamentals of the economy drove their evolution. To answer these questions, we estimate the distribution of CPI inflation twelve months ahead, conditional on current real and nominal conditions. We adopt a similar approach to Banerjee et al. (2020) and López-Salido and Loria (2020), and employ a conventional Philips curve model to forecast the *distribution* of future inflation using quantile regressions.

The estimated quantiles allow us to characterize central features of the forecast distribution, which capture the main attributes of risks to inflation. First, we examine the dispersion of the forecast distribution, which captures forecast uncertainty. Second, we construct a measure for the distribution skewness, which captures the balance between upside and downside risks to inflation.

We find that in the decade following the adoption of the 1–3% target range, uncertainty was stable, and inflation risks were moderately skewed downwards; namely, the dispersion of the distribution was stable, and the skewness was slightly negative. However, in 2014, uncertainty decreased, and downside risks became more dominant. Namely, the dispersion and skewness of the distribution declined. Since 2016, realized inflation gradually increased, and simultaneously, both these measures gradually returned to their long-term averages.

Examining the explanatory variables in our model, we find that the changes in the distribution of inflation since 2014 are mainly due to the decrease in the inflation environment, as it is reflected in realized inflation and long term expectations, and the changes in oil prices at that time.

In 2014 the monetary rate in Israel reached an all-time low, and non-conventional monetary tools were put into use (e.g., forward guidance and foreign exchange interventions). The fact that the monetary rate was approaching the effective lower bound (ELB) may explain why downside risks to inflation became more dominant in these years, as reacting to lower inflation rates was more limited. To test this hypothesis, we re-estimate our model and account for the ELB period using a dummy variable. While the enhancement in downside risks is more moderate in this estimation, the dummy for the ELB accounts for most of it.

Figure 1: Year-over-Year CPI Inflation Excluding Fruits, Vegetables, and Housing (Monthly, 2004-2019)

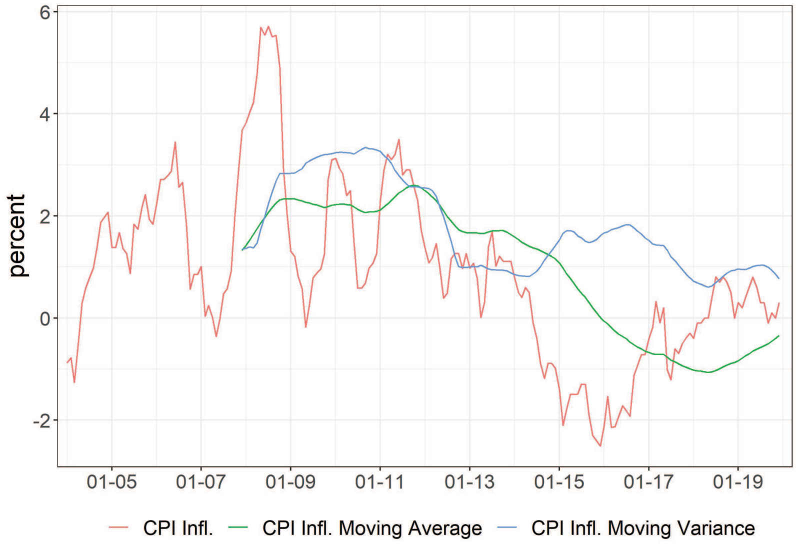

Notes: The figure shows year-over-year CPI inflation rates excluding fruits, vegetables, and housing. It also shows the moving average and variance of inflation rates in a 48-month rolling window.

Furthermore, it accounts for a large share of the decline in uncertainty. Thus, we cannot rule out that the monetary rate approaching the ELB had a substantial effect on the forecast distribution of inflation.

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 restricted models containing (1) only an intercept; (2) an intercept and realized inflation. We evaluate the predictive quality of these forecasts using two measures designed to assess density forecasts. First, following Rossi and Sekhposyan (2019), we analyze the fraction of observations that fall below each quantile (the empirical cumulative distribution of the probability integral transform). Second, we use the quantile R-squared score proposed by Giglio et al. (2016). Both measures indicate that the unrestricted model outperforms the two benchmarks along most of the distribution. Admittedly, the unrestricted model does not improve upon the restricted ones in capturing upside risks to inflation; however, all three models have poor performance in this area. Furthermore, we find that the unrestricted model is the only one that cannot be rejected as capturing the cumulative distribution of inflation using the test by Rossi and Sekhposyan (2019).

The merits of quantile regressions as a tool for risk assessment were proposed by Adrian et al. (2019) who examine the role of macro and financial variables in explaining the forecast distribution of GDP growth. Their model was employed for several economies (Aikman et al., 2018, 2019; Alessandri et al., 2019), and specifically for Israel (Gurkov and Zohar, 2022). The methodology employed in this paper benefits from the same merits, but it has the advantage of being based on more solid theoretical foundations, namely, the Philips curve model. Nonetheless, it is important to note that while the Philips curve model has well-founded theoretical grounds for explaining mean inflation rates, its role in explaining higher moments of the distribution has so far remained unexplored. The empirical work in this paper, which joins evidence of similar models for other countries (Banerjee et al., 2020; López-Salido and Loria, 2020; Tagliabracci, 2020), may form a basis for future theoretical work to identify the mechanisms through which different fundamentals affect the distribution of inflation.

2 Data and Methodology

To examine the fundamental drivers of risks to inflation, we focus on CPI, excluding fruits, vegetables, and housing.[1] The methodology is based on a Philips curve model. However, it is estimated by quantile regressions to study the *distribution* of inflation. For each quantile τ, denote by Q_t^τ the τ-quantile of year-over-year CPI inflation in month t. We estimate the following model, which we name *Inflation at Risk (IaR)*:

$$Q_{t+12}^{\tau} = \beta^{\tau} X_t + \epsilon_t^{\tau}, \tag{1}$$
$$X_t = \left[1, \pi_t, \pi_t^e, \hat{y}_t, oil_t, s_t\right]',$$

where

- π_t is the year-over-year CPI inflation excluding fruits, vegetables, and housing in month t (Figure 1).
- π_t^e is the monthly average of five-year five-year forward breakeven inflation rates.
- $\hat{y}_t$ is a monthly estimate of the output gap (HP filtered State-of-the Economy Index (Marom et al., 2003))
- oil_t is the month-over-month percentage change in Brent Crude oil prices, denominated in USD.

[1] The reason for excluding fruits and vegetables is that this component of the CPI is very volatile and unrelated to the fundamentals of the economy. Excluding housing prices is conducted to overcome the structural change in this component, following the reduction of housing contracts denominated in US dollar since 2007 (Binyamini et al., 2008). Appendix A provides robustness for our main results using headline inflation on a sample starting in 2007.

- s_t is the month-over-month percentage change in the NIS-USD exchange rate.

The selection of explanatory variables generally follows Banerjee et al. (2020). However, we add long-term inflation expectations as in López-Salido and Loria (2020). This variable captures the forward-looking behavior of firms in the economy. We focus on long-term rather than short-term expectations that are more common in standard Philips curve models since we are interested in fundamental drivers of inflation risks. While short-term expectations are by themselves affected by the other fundamentals in our model, long-term expectations mainly capture the nominal anchor as the public perceives it.[2]

The model is estimated on monthly data for the period 2004-2019. The sample begins after the disinflation process in Israel, namely, after the adoption of the 1–3% target range, and ends before the COVID19 crisis. Figure 2 shows the estimated coefficients for each quantile, together with OLS coefficients for comparison.

The coefficients of CPI inflation are positive at around 0.2-0.4 and generally statistically significant, owing to the persistence of inflation. Long-term inflation expectations have a positive effect in most quantiles, but the slope equality test (Bassett and Koenker, 1982) shows that they are significantly different across quantiles (equality is rejected at a 1% significance level).[3] In fact, the coefficients decrease in the quantiles; namely, expectations have a more substantial effect on lower quantiles than on the higher quantiles. As discussed in the following section, this feature bears implications for the expectations' effect on the balance of risks to inflation. Throughout the sample, expectations displayed a downward trend. The estimated coefficients imply that it led to a sharper decline in the lower quantiles of inflation, while the effect on upper quantiles was much smaller. Namely, the decline in expectations is associated with increased downside risks to inflation.

Oil prices also have a more substantial effect on the lower quantiles of inflation than on the higher ones. Thus, a one percent drop in oil prices is reflected in a decrease of about 0.04 percentage points in the median and lower inflation scenarios, with almost no change in the higher quantiles. Namely, large drops in oil price, such as those that occurred in 2008 or 2014, broadened the gap between the upper tail of the forecast distribution and its mid-to-lower part. Thus, they increased forecast uncertainty while enhancing upside risks

[2]It is not straightforward why expectations should be included in a model that aims to explain future inflation. For example, a standard Philips curve model under rational expectations implies that expected inflation should only depend on realized variables (expected inflation "cancels out" from the right-hand side of the equation). However, our model deals with the entire distribution of future inflation and not just its mean. Furthermore, we cannot argue that our model's forecasts are consistent with the market's and that both coincide with rational expectations. Thus, we cannot rule out expected inflation as an explanatory variable in our model.

[3]For the other coefficients, we cannot reject equality across quantiles at any conventional significance level.

Figure 2: Quantile Regression Coefficients

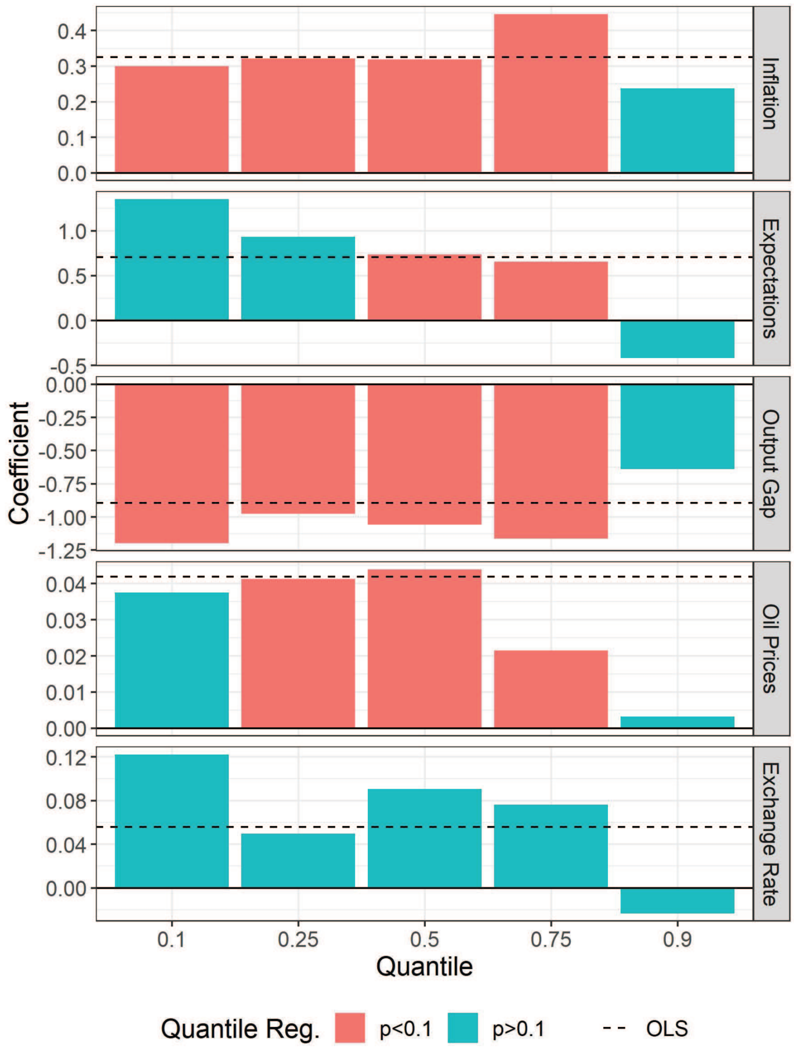

Notes: Each column shows the estimated coefficients from the quantile regression (1). Red bars mark coefficients that are significant at 10% (see Koenker, 1994). Dashed lines show the OLS estimates of Equation (1).

to inflation.

Our estimate for the effect of oil prices on the median scenario (0.04) is similar to the OLS estimate and the pass-through estimated by Kozin (2019). Furthermore, the effect of the exchange rate is also similar to the pass-through in Kozin (2019): a one percent depreciation of the Shekel against the USD increases the median forecast of inflation by 0.08 percentage points.

As for the output gap, its coefficients are *negative* and mostly significant. This result is in contrast to the predictions of standard Philips curve models. To reconcile this contrast, note that the IaR model refers to inflation twelve months ahead, while standard Philips curve models examine higher frequency dynamics of inflation. Figure 13 in the Appendix shows the estimated coefficients of Equation (1) at the one-month horizon, which is similar to a standard Philips curve model. Indeed, we find that the output gap has small and positive coefficients at the shorter horizon, as also found in previous estimates[4]. Thus, the negative coefficients that we find for the twelve-month horizon may originate from the auto-regressive properties of the output gap. In fact, in our sample, the autocorrelation of the output gap with its twelve-month lag is -0.17 ($p = 0.04$). Together, these results imply that while the output gap may not strongly affect inflation in the short run, its cumulative effect over twelve months is significant.

3 Features of the Forecast Distribution

Figure 3 shows quarterly averages of the twenty-fifth, fiftieth, and seventy-fifth quantiles of inflation one year ahead. Not only do the estimated quantiles change over time, but their relative location also varies; namely, the characteristics of the distribution change over time. A notable change occurs at the beginning of 2014 when all three quantiles drop substantially. However, they do not drop to the same extent. First, the upper quantile drops more substantially than the lower one. Consequently, the width of the band that covers fifty percent of the distribution shrinks from about 1.5 percentage points in early 2014 to about one percentage point in mid-2016. If we consider this band to cover the central expected inflation scenarios, then their dispersion decreases in that period. Namely, the central scenarios become more concentrated, and uncertainty about future inflation decreases. Second, the median scenario's drop is the mildest, so the upper quantile approaches the median, and the lower quantile becomes more distant. This development indicates that inflation rates lower than the median become more likely than higher rates, namely, downside risks to inflation

[4]Using various specifications, Box 3.2 in the Bank of Israel Annual Report (2016) finds a coefficient of approximately 0.1 on the output gap, similarly to our estimate at the one-month horizon.

Figure 3: Quantiles of Inflation Twelve Months Ahead (Quarterly Averages)

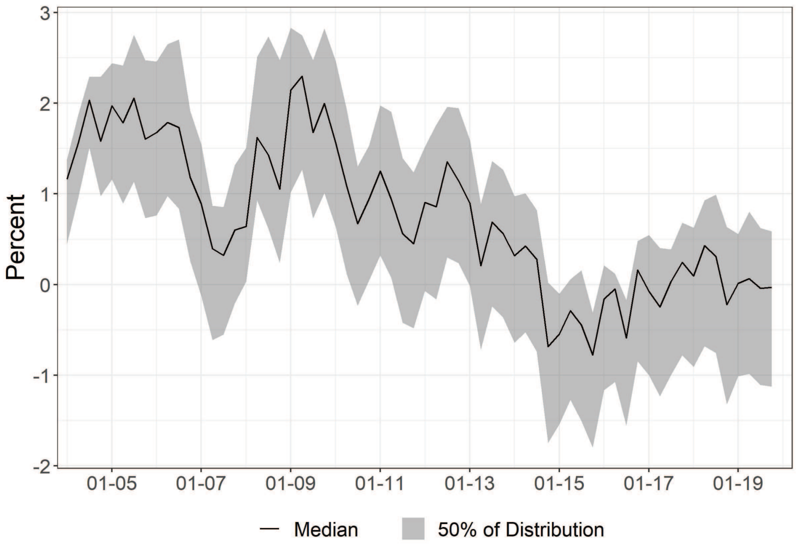

intensify.

Next, we turn to a more formal assessment of the distribution features and their evolution over time. Specifically, we focus on the dispersion of the distribution, which captures forecast uncertainty, and on its skewness, which summarizes the balance between upside and downside risks.

Uncertainty (Dispersion of the Distribution): To evaluate forecast uncertainty, we examine the dispersion of the forecast distribution. Our proxy for the dispersion is the inter-quartile range of the distribution:

$$Dispersion_t \equiv Q_t^{0.75} - Q_t^{0.25}.$$

Balance of Risks (Skewness of the Distribution): To evaluate the balance of upside and downside risks, we examine the skewness of the forecast distribution as captured by the "quartile skewness":

$$Skewness_t = \frac{Q_t^{0.75} + Q_t^{0.25} - 2Q_t^{0.50}}{Dispersion_t}.$$

Figure 4 shows the evolution of the dispersion and skewness over time. During the global financial crisis, both the dispersion and the skewness of the distribution rose. Namely, the crisis was characterized by increased uncertainty and a relative rise in upside risks to inflation. The co-movement of the two indicators is also apparent in 2014 when both declined. In that

Figure 4: High Moments of the Forecast Distribution

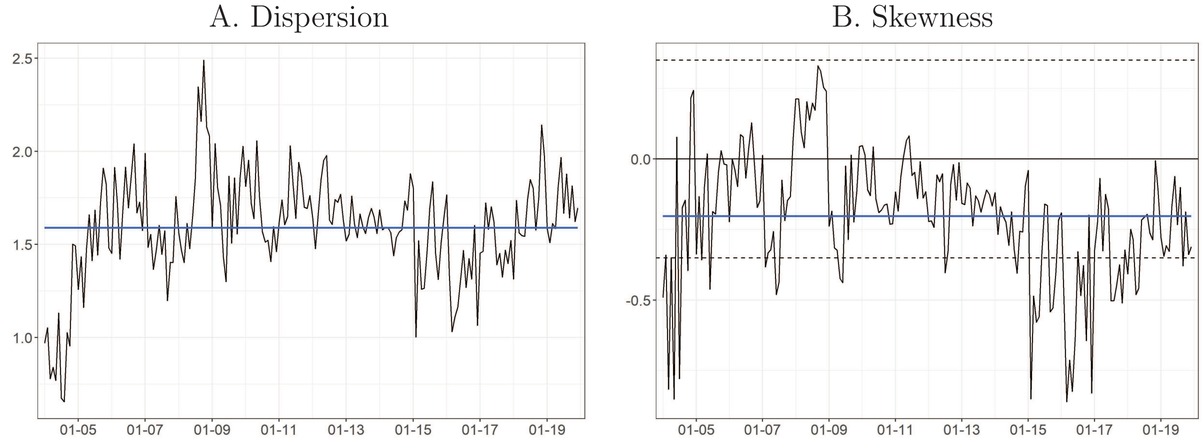

Notes: Blue horizontal lines show sample means. In Panel B, the dashed lines represent $\pm 2SES$, where SES is an estimate of the standard deviation of sample skewness ($SES = \sqrt{\frac{6N(N-1)}{(N-2)(N+1)(N+3)}}$, where N is the sample size).

period, uncertainty decreased together with a relative rise in downside risks.

Figure 5 compares the dispersion and skewness measures from the IaR model to equivalent measures derived from the Bank of Israel's survey of professional forecasters.[5] In Panel A, we show the standard deviation of the forecaster's point estimates, and in Panel B, we show a measure of skewness based on the gap between the mean point forecast and the median. The correlation between the IaR measures and the survey-based measures are about 0.2, both in the dispersion and in the skewness. Both sources indicate an increase in the dispersion and skewness around the 2008 crisis, and a decrease in both during the second decade of the millennium.

To understand what drives forecast uncertainty and risk balance in the model, Figure 6 shows the correlations matrix between the median, dispersion, skewness, and explanatory variables. The dispersion and skewness of the distribution are highly correlated. Namely, in general, higher uncertainty is associated with increased upside risks. Examining the correlations of the distribution features with the explanatory variables shows that the realization of inflation is a primary determinant of the forecast distribution, as its correlation with all the distribution features exceeds 0.4. Changes in oil prices also play an essential role, as they

[5]Professional forecasters report point forecasts for inflation one year ahead. We use these forecasts to construct the measures of dispersion and skewness. Thus, these measures refer to the distribution of point forecasts and not necessarily to the distribution of future inflation. For example, dispersion of point forecasts may originate from forecasters holding different information or using different models. However, there is empirical evidence, supported by theoretical models, that the distribution of point forecasts is related to the distribution of the underlying variable (Zohar (2021) and references therein).

Figure 5: Alternative Measures of Uncertainty and Risk Balance (Quarterly Averages)

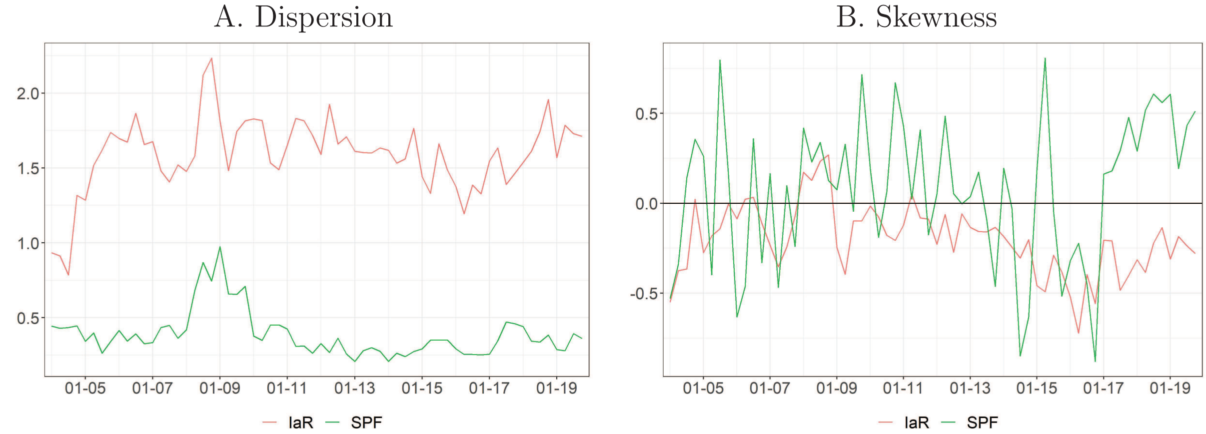

Notes: The figure shows the dispersion and skewness measures from the IaR model alongside equivalent measures based on the Bank of Israel's survey of professional forecasters. Panel A shows the standard deviation of twelve-months-ahead point forecasts. Panel B shows a survey-based skewness measure: $3\big[Mean(f_i) - Median(f_i)\big]/STD(f_i)$, where $\{f_i\}$ is the set of twelve-months-ahead inflation forecasts.

are associated with higher median forecasts and lower uncertainty and skewness.

While Figure 6 summarizes the correlation between the main features of the distribution and the explanatory variables, we can also examine how each variable affected the distribution throughout the sample. Note that

$$E\left(Dispersion_{t+12}\big|X_t\right) = (\beta^{0.75} - \beta^{0.25})X_t,$$

$$E\left(Skewness_{t+12}\big|X_t, Dispersion_{t+12}\right) = \frac{(\beta^{0.75} + \beta^{0.25} - 2\beta^{0.50})X_t}{Dispersion_{t+12}}.$$

Thus, the contribution of the ith variable to the forecast uncertainty is $(\beta_i^{0.75} - \beta_i^{0.25})x_{i,t}$ and its contribution to the skewness (conditional on the uncertainty level) is $\frac{(\beta_i^{0.75}+\beta_i^{0.25}-2\beta_i^{0.50})x_{i,t}}{Dispersion_{t+12}}$.

Figure 7 shows the contribution of each variable to the dispersion and skewness over time.[6] The changes in the dispersion and skewness since 2014 are mainly due to the decrease in the inflation environment, as it is reflected in realized inflation and long-term expectations, and to changes in oil prices.

[6]Note that the contribution of the intercept is not constant over time because each contribution consists of a division by the dispersion in that period.

Figure 6: Correlations between Features of the Forecast Distribution and the Explanatory Variables

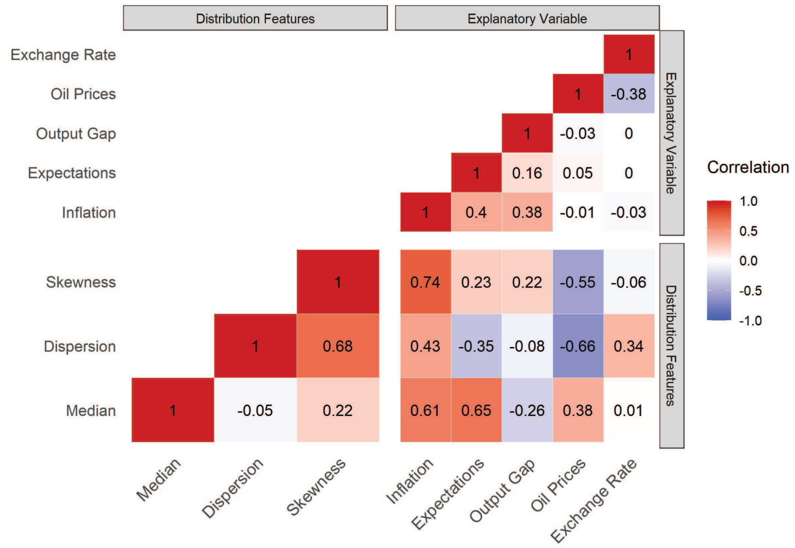

4 Accounting for the Effective Lower Bound of the Monetary Rate

Our baseline estimation shows that the dispersion and skewness of the forecast distribution decreased since 2014, and this is mainly due to the decrease in the inflation environment and changes in oil prices. However, it is important to note that during 2014, the monetary rate in Israel reached an all-time low, and non-conventional monetary tools were employed (e.g., forward guidance and foreign exchange interventions). The fact that the monetary rate was approaching the effective lower bound (ELB) may have played an important role in the change in the forecast distribution. Specifically, it can explain why downside risks to inflation became more dominant in these years since the Monetary Policy Committee was more restrained in reacting to lower inflation rates.

To assess what role the ELB played in generating our results, we estimate a variant of Equation (1) that includes interactions of all variables with a dummy variable for the ELB period. We define the beginning of the ELB period as September 2014, when the monetary rate was lowered to 0.25. At the time, this was the lowest level the monetary rate had ever reached, even though it was later lowered to 0.1. Since the monetary rate was no higher than 0.25 from that point to the end of our sample, we define the ELB period as 09/2014 to 12/2019.

When accounting for the ELB period, we find that the dispersion of the forecast distribution declined since 2014, as in the baseline estimation (Figure 8). However, the skewness

Figure 7: Variable Contribution to Dispersion and Skewness
(Deviations from Mean, Quarterly Averages)

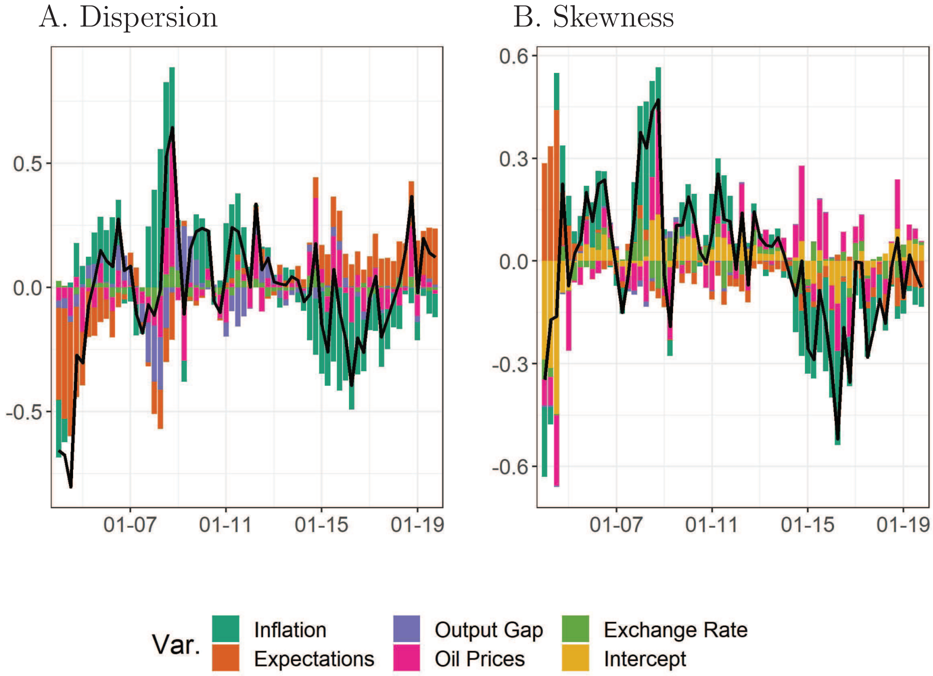

Notes: Stacked bars in Panel A show each variable's contribution to the forecast's dispersion, $(\beta_i^{0.75} - \beta_i^{0.25})x_{i,t}$, minus its average contribution in the sample. In Panel B, they show each variable's contribution to the skewness conditional on the dispersion, $(\beta_i^{0.75} + \beta_i^{0.25} - 2\beta_i^{0.50})\frac{x_{i,t}}{Dispersion_{t+12}}$, minus its average contribution in the sample. The solid line in each panel shows the dispersion and skewness's deviation from their mean.

Figure 8: Accounting for the ELB: Variable Contribution to Dispersion and Skewness (Deviation from Mean, Quarterly Averages)

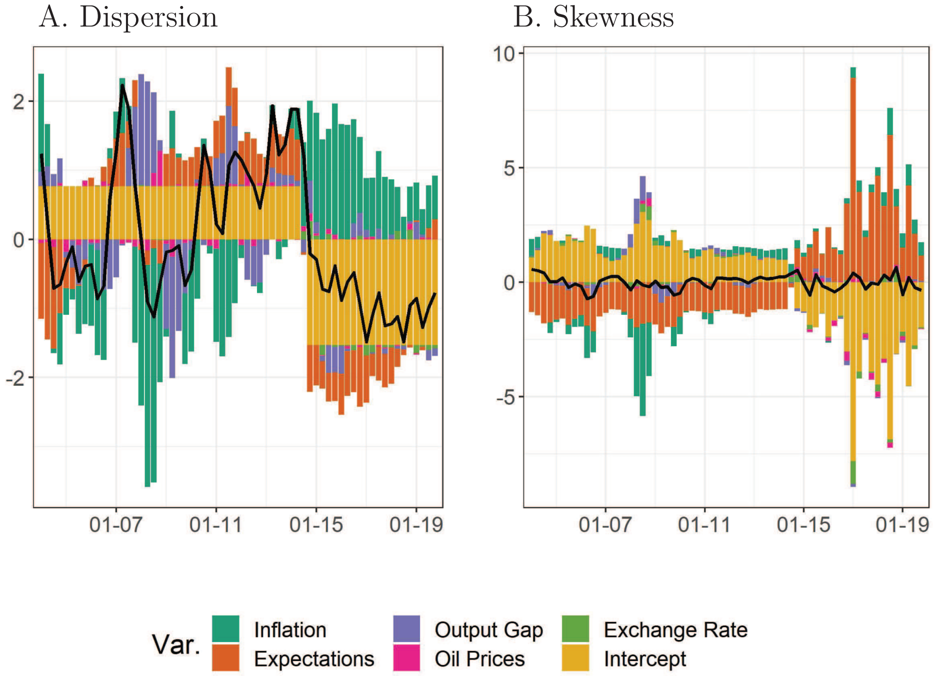

Notes: Stacked bars in Panel A show each variable's contribution to the forecast's dispersion, $(\beta_i^{0.75} - \beta_i^{0.25})x_{i,t}$, minus its average contribution in the sample. In Panel B, they show each variable's contribution to the skewness conditional on the dispersion, $(\beta_i^{0.75} + \beta_i^{0.25} - 2\beta_i^{0.50})\frac{x_{i,t}}{Dispersion_{t+12}}$, minus its average contribution in the sample. We sum the contribution of each variables and its interaction with the ELB dummy. The solid line in each panel shows the dispersion and skewness's deviation from their mean.

of the distribution only mildly decreased in that period.

Figure 8 shows the contribution of each variable to the dispersion and skewness. For each variable, we sum its contribution before the ELB period and after it. Namely, we sum the contributions of the variable and its interaction with the ELB dummy. The figure shows that the inflation environment is still an important contributor to the decrease in the dispersion, but oil prices no longer seem to have an important role. Furthermore, the dummy variable for the ELB period, which is captured by the intercept variable, also contributes substantially to the dispersion's decline.

As for the skewness, it seems that the ELB dummy contributes substantially to its decline. This is expected as the ELB limits the possibility to react to low inflation rates using the monetary rate and thus enhances downside risks. However, it seems that expectations have a positive effect that almost completely offsets the negative effect of the ELB, so overall skewness decreased only mildly.

Altogether, we find that our main results hold, namely, that the dispersion and skewness of the forecast distribution declined since 2014. However, the decline in the skewness is much milder than in the baseline estimation. Furthermore, the dummy variable for the ELB period has a substantial explanatory power for these changes, while the inflation environment and oil prices play only a secondary role. It is important to note that the ELB dummy may be capturing downward trends in the inflation environment or oil prices, thus taking on some of their explanatory power. However, we cannot rule out that the changes in the features of the forecast distribution were caused by the interest rate approaching the ELB.

5 Out-of-Sample Forecast Performance

In this section, we evaluate the model's out-of-sample forecast performance. We compare the IaR model to two restricted models: (1) a model containing only an intercept; (2) a model containing an intercept and realized inflation.

We evaluate the performance of all three models as they would have been used in real-time. We generate out-of-sample forecasts for CPI inflation quantiles on an expanding window, with an initial window of eleven years (Jan-2004 – Jan-2015).[7] Specifically, the IaR forecasts for period t are generated by performing the methodology of Section 2 on the sample ending in period $t-12$ (Figure 9). Similarly, the benchmark forecasts generated by estimating restricted versions of Equation (1) on the same window (Figure 14 in the Appendix).

Next, we evaluate the predictive quality of the models using two measures designed to assess density forecasts. First, following Rossi and Sekhposyan (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, we compute the empirical cumulative distribution of the PITs, which is the percentage of observations that fall below the forecast quantile $\hat{Q}_t^\tau$:

$$\varphi(\tau) \equiv \frac{1}{T-t_0+1} \sum_{t=t_0}^{T} \mathbb{I}_{\{\pi_t < \hat{Q}_t^\tau\}},$$

where $\hat{Q}_t^\tau$ are forecasts of the relevant model for period t (estimated on the sample ending in period $t-12$), t_0 is the first out-of-sample forecast date (Jan-2016 in our case), T is the sample size, and $\mathbb{I}$ is the indicator function. A model is better fitted the closer $\varphi(\tau)$ is to

[7] We choose an initial window that contains the structural change we found in Section 3.

Figure 9: Out-of-Sample Forecasts and Realized CPI Inflation

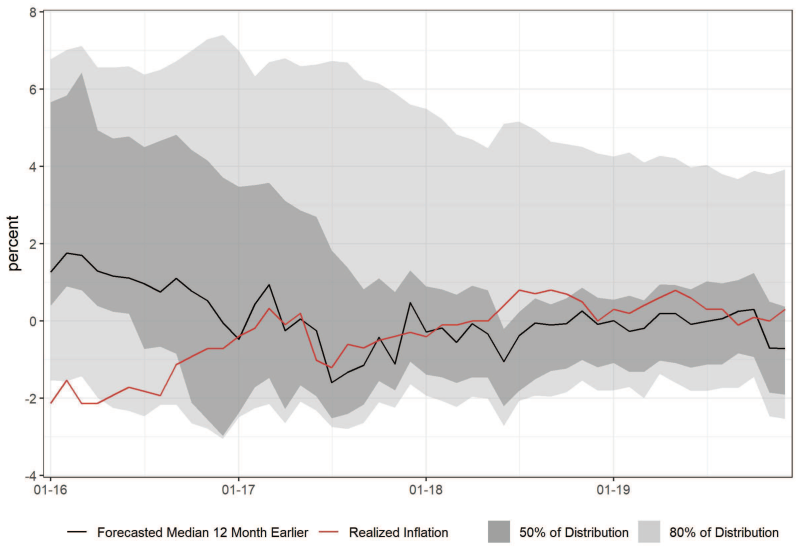

Notes: The figure shows the median out-of-sample forecast of the IaR model estimated twelve months before the marked date (black line) together with the corresponding realized inflation (red line). The grey areas show the 80% and 50% bands based on the IaR model.

the 45-degree line. If the model perfectly fits the empirical cumulative distribution and the sample is large enough, then the fraction of observations falling below quantile τ should be exactly τ, namely, $\varphi(\tau) = \tau$.

Panel A in Figure 10 shows the φ scores of the IaR model and the two benchmark models, together with 95% confidence bands (Rossi and Sekhposyan, 2019). The IaR model is the only one that lies inside the confidence band in all quantiles. Thus, the restricted models are rejected for capturing the true commutative distribution, while the IaR model is not. Furthermore, the IaR model outperforms the two benchmarks in most quantiles. Note that all three models have poor performance in the ninetieth quantile, as a hundred percent of the observations fall below the forecasts for this quantile. Namely, none of the models identifies upside risks well.

The assessment based on the PITs gives a relatively rough estimate of the forecast performance, as it only addresses the question of whether realized inflation 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 and the quantile forecasts, namely, the forecast errors. It compares these errors to those generated by the restricted model containing only an intercept. For each quantile τ, we look at a weighted average of forecast errors $\left|\pi_t - \hat{Q}_t^{\tau}\right|$, and compare it to a similar weighted average of the

Figure 10: Evaluation of Out-of-Sample Forecast Accuracy

A. Probability Integral Transform (PIT)

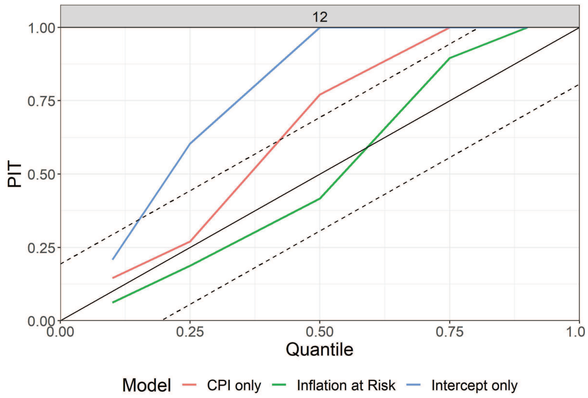

B. Quantile R-Squared

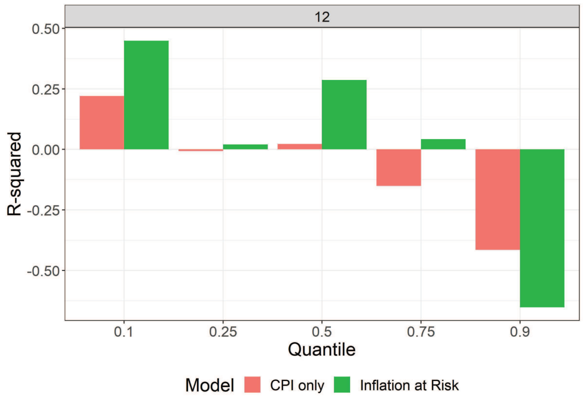

Notes: Panel A depicts the empirical cumulative distribution of the probability integral transform (PIT), i.e., the share of observations that fall below each forecast quantile. The 45-degree black line is added for reference, and dashed lines represent the 95% confidence bands following Rossi and Sekhposyan (2019). Panel B depicts the quantile R-squared scores. In both panels, the scores of the Inflation at Risk model appear alongside a restricted model containing only an intercept and CPI inflation. Panel A also shows the performance of the third model - a restricted model containing only an intercept - which is already inherent in the R-squared score.

forecast errors from the "intercept only" model:

$$R^2(\tau) \equiv 1 - \frac{\sum_{t=t_0}^{T} \left(\pi_t - \hat{Q}_t^{\tau}\right)\left(\tau - \mathbb{I}_{\{\pi_t < \hat{Q}_t^{\tau}\}}\right)}{\sum_{t=t_0}^{T} \left(\pi_t - \hat{c}_t^{\tau}\right)\left(\tau - \mathbb{I}_{\{\pi_t < \hat{c}_t^{\tau}\}}\right)},$$

where $\hat{c}_t^{\tau}$ are the forecasts from the "intercept only" model. Note that forecast errors are weighted according to the quantile in question. For example, the 0.10 quantile R-squared weighs positive forecast errors by 0.10 and negative errors by 0.90, heavily penalizing realizations that fall below the quantile forecast. Higher R-squared values express better forecast performance compared to the "intercept only" benchmark. Conversely, negative values indicate that the model's accuracy falls short of that benchmark.

Panel B in Figure 10 shows R-squared scores of the Inflation at Risk and "CPI only" models. The IaR scores are positive in all quantiles but 0.90, so the model improves upon the "intercept only" benchmark up to this quantile. Furthermore, in the same region, the IaR scores are higher than those of the "CPI only" benchmark. Namely, the IaR model outperforms the two benchmarks in all quantiles but the ninetieth quantile.

6 Conclusion

The decade following the adoption of the current inflation target range (1–3%) was characterized by inflation varying around the middle of the range. However, since 2014 the statistical properties of inflation changed: both its mean level and volatility have declined. The IaR model proposed in this paper allows us to characterize if and how these observed developments reflect changes in the distribution of inflation. Furthermore, the model identifies the fundamentals of the economy that contributed to these changes.

Using the model, we find that the decline in inflation since 2014 was associated with a decline in uncertainty and a relative rise in downside risks to inflation. The decline in long-term expectations and changes in oil prices also played a prominent role in these developments. Other determinants of inflation, namely, the output gap and exchange rate, played only a minor role in the evolution of the distribution of inflation at the time.

We also find that the monetary rate approaching its effective lower bound since 2014 may have contributed to the changes in the risks to inflation. Specifically, it may have enhanced downside risks since reacting to lower inflation rates is more limited around the ELB.

While the main focus of this paper is the characterization of historical risks to inflation, the IaR model can also be used to assess inflation risks regularly. Its main advantage is in offering state-dependent assessments of risks, and the out-of-sample evaluation shows that it performs well compared to other, more parsimonious benchmarks.

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*.

Banerjee, R. N., J. Contreras, A. Mehrotra, and F. Zampolli (2020). Inflation at risk in advanced and emerging market economies. *BIS Working Paper* (883).

Bank of Israel Annual Report (2016). *Box 3.2: The factors affecting the development of inflation in Israel in the short-term*, Chapter 3, pp. 105–111.

Bassett, G. and R. Koenker (1982). Tests of linear hypotheses and l1 estimation. *Econometrica 50*(3), 1577–1583.

Binyamini, A., Z. Eckstein, and K. Flug (2008). The evolution of monetary-policy strategy and exchange-rate regime in Israel.

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

Gurkov, M. and O. Zohar (2022). Growth at risk: Forecast distribution of GDP growth in Israel. *Bank of Israel Discussion Paper No 2022.08.*

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

Kozin, M. (2019). Exchange rate pass-through to prices. *Selected Research and Policy Analysis Notes, Bank of Israel, February 2019*, 39–53.

López-Salido, J. D. and F. Loria (2020). Inflation at risk. *Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System 2020-013.*

Marom, A., Y. Menashe, T. Suchoy, et al. (2003). *The State-of The-economy Index and the Probability of Recession: The Markov Regime-switching Model.* Research Department, Bank of Israel.

Rossi, B. and T. Sekhposyan (2019). Alternative tests for correct specification of conditional predictive densities. *Journal of Econometrics 208*(2), 638–657.

Tagliabracci, A. (2020). Asymmetry in the conditional distribution of euro-area inflation. *Bank of Italy Temi di Discussione (Working Paper) No 1270.*

Zohar, O. (2021). Cyclicality of uncertainty and disagreement. *Bank of Israel Discussion Paper* (2021.09).

Appendix

A Headline Inflation

In the baseline model, we estimated the forecast distribution of CPI since the entrenching of the current inflation target in 2004. However, we used a partial consumer price index, excluding fruits, vegetables, and housing. We excluded fruits and vegetables since their prices are very volatile and unrelated to the fundamentals of the economy. Housing prices were excluded to overcome the structural change in this component, following the reduction of housing contracts denominated in US dollar since 2007 (Binyamini et al., 2008). This omission allowed us to maximize the sample size without directly dealing with the structural change. A large sample improves the estimation of quantile regressions and the identification of risks to inflation.

However, since the Bank of Israel targets headline inflation, this section analyzes its forecast distribution. To deal with the issue of the structural change in housing contracts mentioned above, we begin our sample when the shift toward Shekel-denominated contracts began. Namely, we estimate the model for the period of 2007-2019.

Figure 11 shows that the forecast distribution of headline inflation also exhibited a decline

Figure 11: Headline Inflation: Variable Contribution to Dispersion and Skewness (Deviation from Mean, Quarterly Averages)

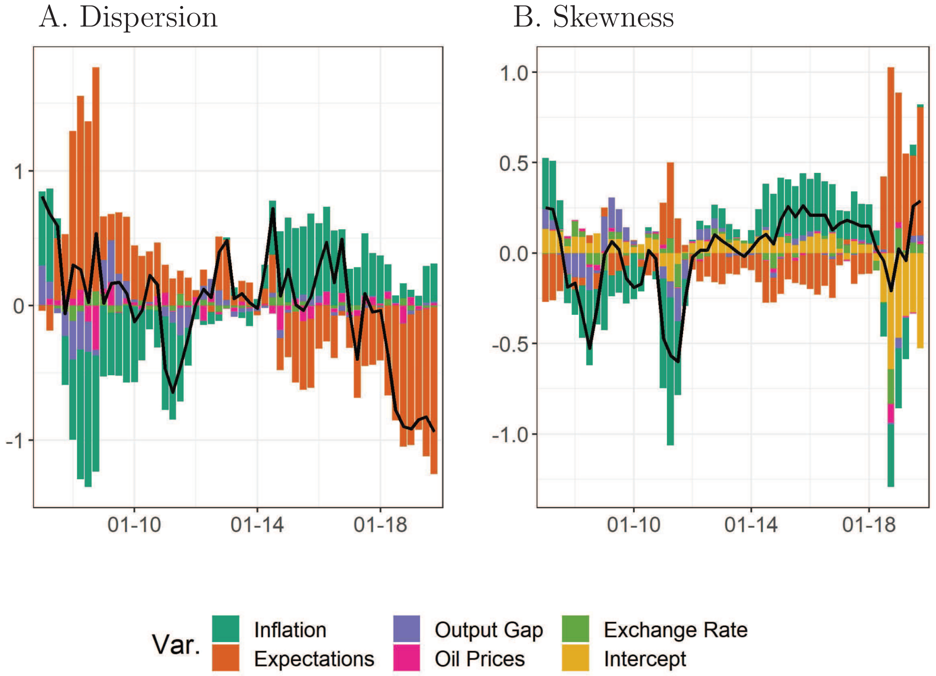

Notes: Stacked bars in Panel A show each variable's contribution to the forecast's dispersion, $(\beta_i^{0.75} - \beta_i^{0.25})x_{i,t}$, minus its average contribution in the sample. In Panel B, they show each variable's contribution to the skewness conditional on the dispersion, $(\beta_i^{0.75} + \beta_i^{0.25} - 2\beta_i^{0.50})\frac{x_{i,t}}{Dispersion_{t+12}}$, minus its average contribution in the sample. The solid line in each panel shows the dispersion and skewness's deviation from their mean.

Figure 12: Headline Inflation: Evaluation of Out-of Sample Forecast Accuracy

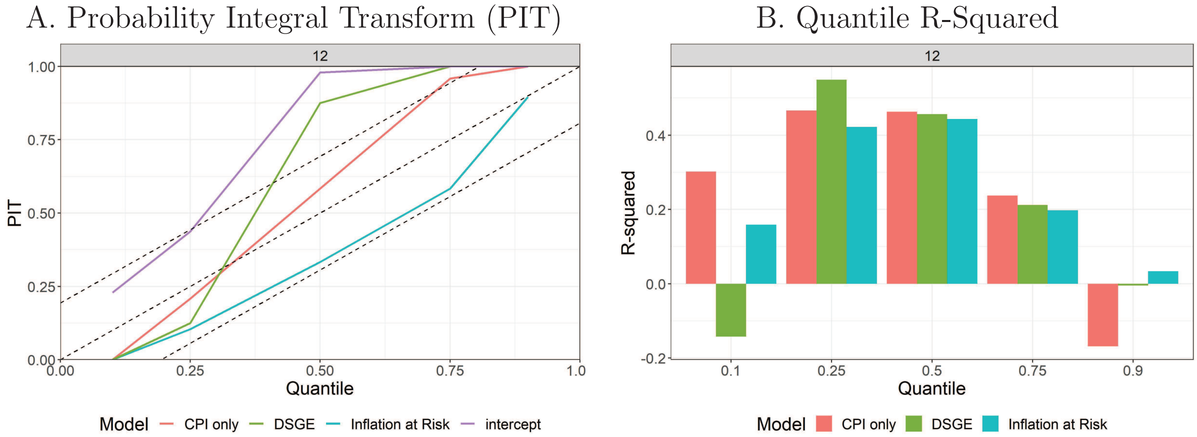

Notes: Panel A depicts the empirical cumulative distribution of the probability integral transform (PIT), i.e., the share of observations that fall below each forecast quantile. The 45-degree black line is added for reference, and dashed lines represent the 95% confidence bands following Rossi and Sekhposyan (2019). Panel B depicts the quantile R-squared scores. In both panels, the scores of the Inflation at Risk model appear alongside a restricted model containing only an intercept and CPI inflation and alongside the DSGE model (this model is quarterly so the confidence interval of Panel A is irrelevant to it, see Footnote 9). Panel A also shows the performance of the forth model - a restricted model containing only an intercept - which is already inherent in the R-squared score.

in the dispersion and skewness since 2014. However, compared to the baseline model, the decline in the dispersion was more persistent, while the decline in skewness was shorter-lived. Furthermore, as in the baseline model, the inflation environment (realized and expected inflation) played an essential role in these changes.

The headline inflation IaR model may be more suitable for regular analysis of risks to inflation as it tracks the Bank of Israel's primary variable of interest. Therefore, we also examined the performance of the out-of-sample forecasts of this model. Panel A in Figure 12 shows that the IaR model's PIT scores lie inside the 95% confidence band in all quantiles, while the restricted models deviate from it. Thus, as in the baseline estimation, the IaR model is the only one that cannot be rejected as capturing the true cumulative distribution of inflation. Furthermore, according to this test, the IaR model outperforms the "intercept only" benchmark in all quantiles. Admittedly, the "CPI only" benchmark outperforms the IaR model in some quantiles.

Panel B in Figure 12 shows that the IaR model outperforms the "intercept only" benchmark, as all the scores are positive. The model performs similarly to the "CPI only" benchmark, with a higher score at the ninetieth percentile and a lower score at the tenth percentile.

Figure 12 also presents the forecast performance of the Bank of Israel's DSGE model (Argov et al., 2012). The distribution of the DSGE forecast is symmetric with a fixed variance, so the quantiles are at a fixed distance from the median forecast.[8] It should be noted that the DSGE model was re-estimated in 2019, so its predictions are in-sample, which gives it an advantage over the other models. Nonetheless, the two tests in Figure 12 show that the IaR model does not substantially fall short of the DSGE model. In fact, it shows better performance in some quantiles (in the PIT test, the IaR model improves upon the DSGE mainly in the upper part of the distribution, and in the R-squared test, it improves in the two extreme tails of the distribution).[9]

Altogether, both forecast performance tests support the use of the distribution generated by the IaR model to forecast headline inflation.

B Additional Figures

[8]The forecast distribution in the DSGE model originates from the various shocks in the model. The assumption in the model is that these shocks are normally distributed, which imposes normality on the forecast distribution of inflation. As a result, the forecast is symmetric, and its standard deviation is fixed and stems from the standard deviation of the shocks.

[9]The DSGE model is a quarterly model, so the number of observations on which we test its forecast is relatively small (20 observations versus 60 of the IaR model and the restricted models). One of the implications is that the confidence band presented in Figure 12.A is irrelevant for the DSGE model. However, even with a confidence band fitted for the number of quarterly observations, the hypothesis that the DSGE model captures the true inflation distribution is rejected.

Figure 13: Quantile Regression Coefficients One and Twelve Months Ahead

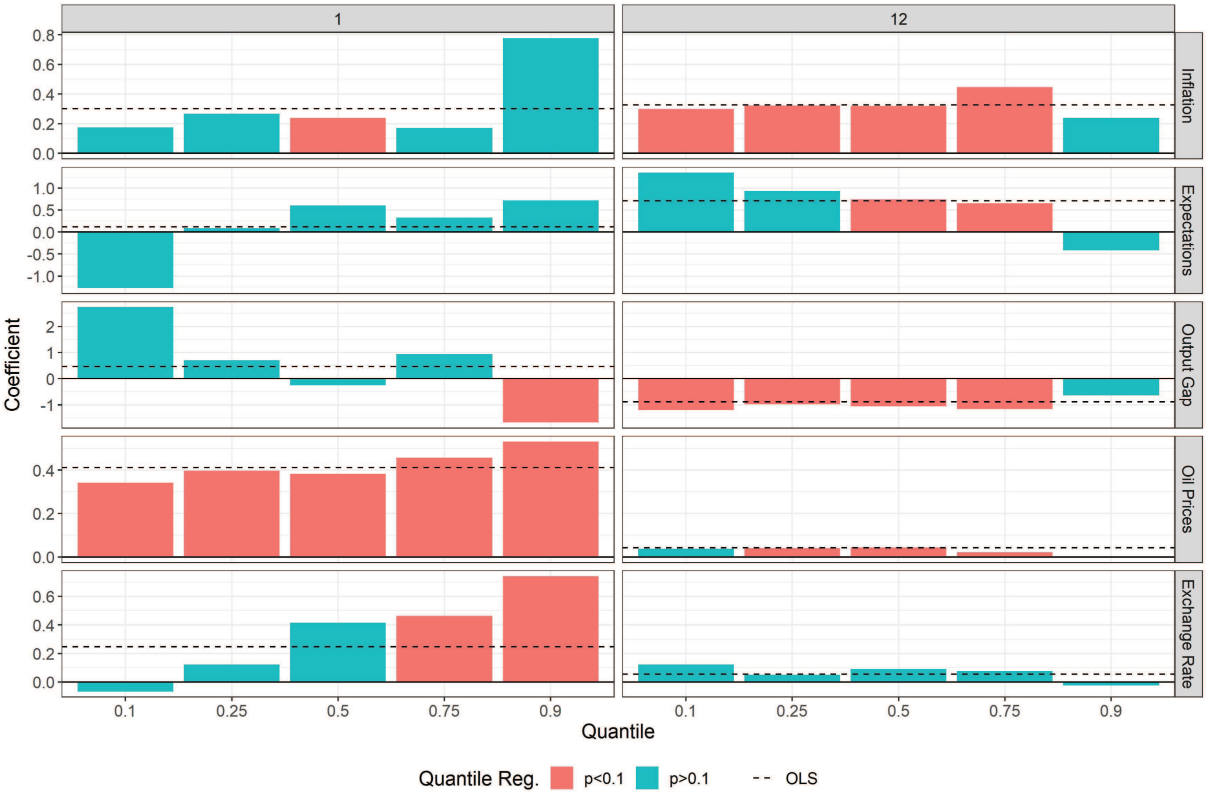

Notes: Each column shows the estimated coefficients from the quantile regression (1) for one or twelve months ahead. For the one-month horizon, the dependent variable in Equation (1) is replaced by monthly inflation in annual terms. Red bars mark coefficients that are significant at 10% (see Koenker, 1994). Dashed lines show the OLS estimates of Equation (1).

Figure 14: Out-of-Sample Forecasts

A. Inflation at Risk

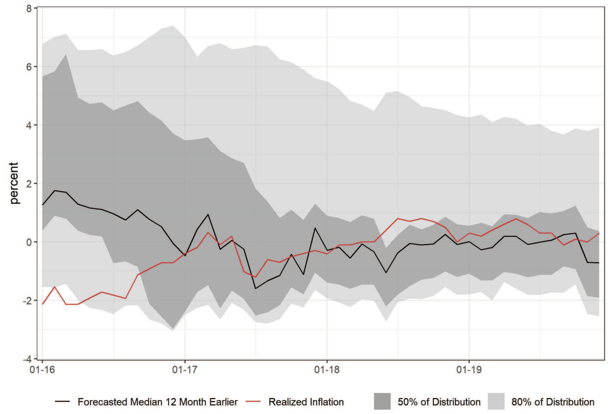

B. "CPI Only"

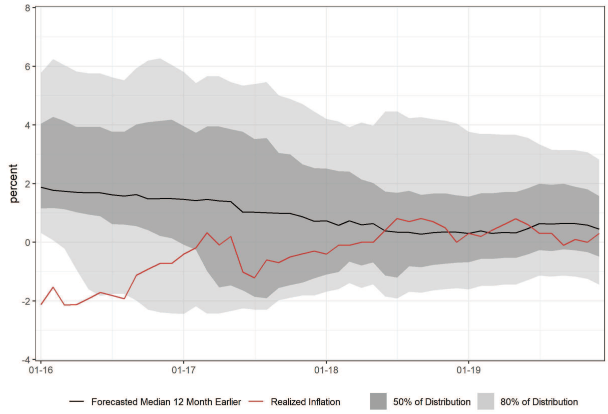

C. "Intercept Only"

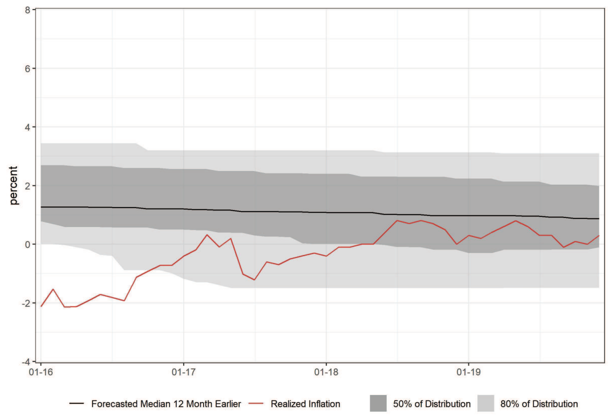

Notes: Each panel shows the median out-of-sample forecast of the respective model estimated twelve months before the marked date (black line), together with bands covering 80% and 50% of the forecast distribution (grey bands). The red line depicts realized CPI inflation.