Bank of Japan Working Paper Series

New Facts about Firms' Inflation Expectations: Short- versus Long-Term Inflation Expectations

Yosuke Uno*
yousuke.uno@boj.or.jp

Saori Naganuma**
saori.naganuma@boj.or.jp

Naoko Hara***
naoko.hara@boj.or.jp

No.18-E-15
October 2018

Bank of Japan
2-1-1 Nihonbashi-Hongokucho, Chuo-ku, Tokyo 103-0021, Japan

* Research and Statistics Department
** Research and Statistics Department (currently at the Financial System and Bank Examination Department)
*** Research and Statistics Department (currently at the Institute for Monetary and Economic Studies)

New Facts about Firms' Inflation Expectations: Short- versus Long-Term Inflation Expectations*

Yosuke Uno† Saori Naganuma‡ Naoko Hara§

August 21, 2018

Abstract

In this paper, using large-scale firm-level micro-data from the *Tankan* survey we examine firms' inflation expectations at different time horizons. Our principle findings are twofold. First, with regard to long-term expectations, a number of firms offer no forecasts. Second, and more importantly, the frequency of forecast revisions is higher for longer time horizons.

JEL codes: E31, E58

Keywords: long-term inflation expectations, frequency of forecast revisions

*We thank Ryohei Hisano, Kosuke Aoki, Masahiro Hori, Shigehiro Kuwabara, Toshitaka Sekine, Koji Nakamura, Hibiki Ichiue, Sohei Kaihatsu, Satoshi Ito, Koki Inamura, Ichiro Muto, Toshinao Yoshiba, Takushi Kurozumi, and Ryo Kato for useful comments. All remaining errors are ours. This paper does not necessarily reflect the views of the Bank of Japan.

†Research and Statistics Department

‡Research and Statistics Department (currently at the Financial System and Bank Examination Department)

§Research and Statistics Department (currently at the Institute for Monetary and Economic Studies)

1 Introduction

In this paper, using the same data set as in Uno et al. (2018)—large-scale firm-level micro-data from the *Tankan* survey—we examine firms' inflation expectations at different time horizons. Inflation expectations at different time horizons are traditionally categorized and discussed vis-à-vis whether they are "short-term" or "long-term." On the one hand, as in Clark and Davig (2008) and Yellen (2015), some researchers emphasize the importance of long-term inflation expectations in inflation dynamics. For example, Yellen (2015) argues, by means of graphical comparison, in favor of a direct link between the long-term inflation expectations formed by households and professional forecasters and estimates of the long-term inflation trend. By contrast, Fuhrer (2012) finds that long-term inflation expectations have no direct influence on inflation. According to Fuhrer, inflation is driven by short-term expectations; long-term expectations influence inflation only indirectly via any impact on short-term expectations.

In this paper, we report some findings on firms' inflation expectations at different time horizons that appear to contradict the arguments of Clark and Davig (2008) and Yellen (2015). Our principle findings are twofold. First, with regard to long-term expectations, a number of firms offer no forecasts. Specifically, although only 15 percent of Japanese firms in our *Tankan* survey data set answer "don't know" for forecasts 1-year ahead, no less than 40 percent of firms respond "don't know" for forecasts 5-years ahead. This suggests that more than 40 percent of firms would also answer "don't know" for longer-term forecasts such as for 10-years ahead. Second, and more importantly, the frequency of forecast revisions rises with the length of the time horizon. This finding runs counter to the view propounded by Clark and Davig (2008) and Yellen (2015), as it is unlikely that more frequently revised long-term expectations would be able to capture the long-term infla-

tion trend. Since the use of aggregated data, as in most studies of inflation dynamics, prevents researchers from directly observing individual firm behavior with regard to forecast revisions, the use of micro-data to explore the frequency of forecast revision is therefore a distinctive contribution of our approach here.

This paper proceeds as follows. In Section 2, we present our key findings concerning firms' inflation expectations at different time horizons. In Section 3, we discuss the reasons for the more frequent revisions of longer-term inflation expectations that we observed. Section 4 concludes.

2 Findings

We set out to explore the features of firms' inflation expectations and how these vary by time horizon. In this section, we document our findings about firms' inflation expectations at different time horizons.

2.1 Data

In this paper, as in Uno et al. (2018), we use large-scale firm-level micro-data from the *Tankan* survey gathered by the Bank of Japan's Research and Statistics Department. In this subsection, we briefly describe our *Tankan* data.[1]

The *Tankan* collects information on firms' inflation expectations with regard to two sets of prices: general prices and output prices. The question regarding general prices is phrased as follows: "What are your institution's expectations on the annual percent change in general prices (as measured by the consumer price index) for one year ahead, three years ahead, and five years ahead, respectively? Please select the range nearest to your own expectation from the options below." Note that the question explicitly refers

[1]For basic statistics of our sample, see Uno et al. (2018).

to the consumer price index (CPI). Firms are provided with ten options starting with (1) around +6% or higher, (2) around +5% (+4.5%∼+5.4%), (3) around +4% (+3.5%∼+4.4%), and continuing in one-percentage point intervals until (10) around −3% or lower. If respondents have no clear view on the outlook for general prices, they are asked to select one of the following three reasons: (11) uncertainty over the future outlook is high, (12) not really conscious of inflation fluctuations because they should not influence the strategy of the institution, and (13) other.[2]

Regarding output prices, the question is: "Relative to the current level, what are your institution's expectations of the rate of price change in your mainstay domestic product or service for one year ahead, three years ahead, and five years ahead, respectively? Please select the range nearest to your own expectation from the options below." Respondents are again provided ten options to choose from, starting with (1) around +20% or higher, (2) around +15% (+12.5%∼+17.4%), and continuing in five-percentage point intervals to (9) around −20% or lower, and (10) don't know. Regarding their outlook for output prices three and five years ahead, unlike in the case of general prices, firms are asked to provide their outlook for *cumulative* changes in their output prices relative to the current level.

Each quarterly *Tankan* survey covers around 10,000 firms selected from a population of approximately 210,000 firms with paid in capital of at least 20 million yen. The question on inflation expectations is included in the survey from March 2014 onwards. Our observation period therefore runs from March 2014 to the latest survey available at the time of writing, September 2017. The *Tankan* data is an unbalanced panel data in which there is a core of firms that have been surveyed since the start (May 1974), while some have dropped out because their capital fell below 20 million yen and/or went

[2] A sample form of the *Tankan* questionnaire is available on the Bank of Japan's website at http://www.boj.or.jp/en/statistics/outline/exp/tk/extk01.htm/

bankrupt. To preserve the sample size, firms are newly added once in two or three years. In our sample, about 1,000 firms are added in March 2015. The panel structure of the *Tankan* survey allows us to track the individual firms' expectation formation over time.

2.2 Fraction of firms responding "don't know" by time horizon

As decsribed in Section 2.1, in the *Tankan* survey, respondents are required to choose a range closest to their own forecast from a number of options including "don't know." In this subsection, we simply document the fraction of firms responding "don't know" by time horizon. Figure 1 presents the evolution of fraction of answers with "don't know" in our *Tankan* survey.

Figure 1: Fraction responding "don't know" by time horizon

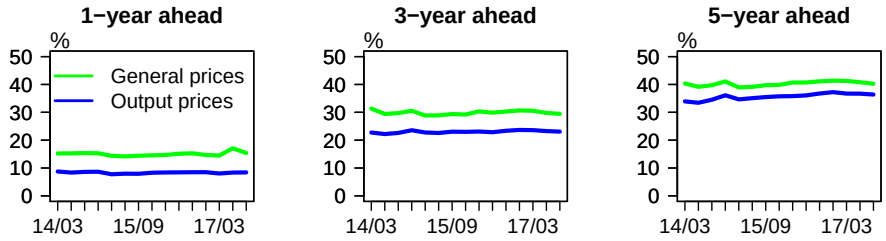

As clearly observed in Figure 1, the longer the inflation expectations horizon, the higher the proportion of respondents answering "don't know." For instance, for general prices (the green line in Figure 1), the fraction of firms answering "don't know" rises from about 15 percent for forecasts 1-year ahead to about 40 percent for forecasts 5-years ahead. Even when forecasting their own output prices (the blue line in Figure 1), about 35 percent of firms answer "don't know" for 5-year horizons. Naturally, this implies that a great number of firms, perhaps significantly more than 40 percent will answer "don't know" for longer-term horizons such as 10-years ahead. Another important observation is that the fraction responding "don't know" remains

stable over time, even for longer (5-year) horizons. This suggests that firms' ability to forecast price changes at such horizons is unlikely to be especially sensitive to the economic environment, which varied considerably during the period of the survey.

2.3 Term structure

Next, we examine the term structure of firms' inflation expectations. Table 1 shows the relative size of forecasts at different time horizons: 1-, 3-, and 5-years ahead. This is a simple way to capture the term structure of firms' inflation expectations.

As mentioned in Section 2.1, forecasts for output prices 3- and 5-years ahead indicate cumulative changes relative to current levels, so forecasts at different horizons are not readily comparable. In following observations are, therefore, restricted to general prices.

First, the most common term structure is where "1 = 3 = 5"; the proportion of firms whose responses fit this pattern is the highest. For instance, 39 percent of firms predict the same year-on-year rates of change in the CPI at 1-, 3-, and 5-year horizons. This tendency is more prominent in large firms: 50 percent of large firms expect the same rate at all three time horizons. This finding about term structure is consistent with Uno et al. (2018), where it is noted that firms' inflation expectations vary considerably depending on firm size. Viewed in terms of "the inflation expectations curve," the 1-year-to-5-year curves for 39 percent of all firms and 50 percent of large firms are seen to be *flat*, though it should be noted that these flat "curves" trace different expected inflation levels.

Second, focusing on 3- and 5-year horizons, 71 percent of firms forecast the same year-on-year rates both 3- and 5-years ahead. That is, for 71 percent of firms "the inflation expectations curve" conforms to one of the

following three shapes: "$1 = 3 = 5$," "$1 > 3 = 5$," and "$1 < 3 = 5$." A significant majority of firms thus expect CPI inflation to become constant after three years and to remain stable for a while afterward. Put differently, the 5-year forecasts of 71 percent of Japanese firms contain information that is no different from the 3-year forecasts. Again from the point of view of heterogeneity with respect to firm size, while 70 percent of small and medium-sized firms expect to see the same CPI inflation rate both 3- and 5-years ahead, the equivalent figure for large firms is 77 percent.

Third, the proportion of firms with an upward sloping or "normal" curve (i.e., "$1 < 3 < 5$") is less than 10 percent; meanwhile, a similarly small fraction of firms exhibit an "inverted" curve ("$1 > 3 > 5$"). As Aruoba (2016) argues, variations in the shape of the curve potentially contain a wealth of information on public perception of monetary policy. Unfortunately, our sample comprises only fifteen quarters of time series data, which is insufficient to analyze changes in the shape of firms' inflation expectations curves over time. Such an examination of changes in term structure over time would be an interesting avenue for future research once more data has been compiled.

Table 1: Term structure of inflation expectations

	$1=3=5$	$1=3>5$	$1=3<5$	$1>3=5$	$1>3>5$	$1>3<5$	$1<3=5$	$1<3>5$	$1<3<5$
	General prices								
All firms	0.39	0.04	0.08	0.09	0.02	0.01	0.23	0.04	0.09
Large firms	0.50	0.04	0.07	0.08	0.02	0.01	0.19	0.04	0.06
S&M firms	0.37	0.04	0.09	0.09	0.03	0.01	0.24	0.04	0.09
Manu 1	0.40	0.04	0.08	0.08	0.02	0.01	0.23	0.04	0.09
Manu 2	0.41	0.04	0.09	0.08	0.02	0.01	0.23	0.04	0.09
Non-Manu	0.38	0.04	0.09	0.09	0.03	0.01	0.24	0.04	0.09

Notes. "$1=3=5$" means that inflation expectations at 1-year, 3-year, and 5-year horizons are the same. Large firms are defined as firms with capital of at least 1 billion yen; S&M firms—small- and medium-sized firms—are defined as firms with capital of at least 20 million yen but less than 1 billion yen. Manu 1, Manu 2, and Non-Manu denote manufacturing (basic materials), manufacturing (processing), and non-manufacturing, respectively.

2.4 Frequency of forecast revisions by time horizon

The frequency of forecast revisions provides fruitful information to model firms' inflation expectations formation. For example, Andrade and Le Bihan (2013) among others estimate the frequency of forecast revisions using micro-data to distinguish between a sticky information model and a noisy information model. In this subsection, we document estimates of the frequency of forecast revisions by time horizon, as a stepping stone to building a comparable model to investigate the formation of short- and long-term inflation expectations.

As discussed in Uno et al. (2018), using *Tankan* data, we can estimate the frequency of forecast revisions in two ways: the first is appropriate for "fixed-target horizon" forecasts, and is denoted by λ in Uno et al. (2018); the second is appropriate for "rolling horizon" forecasts, and is denoted by λ' in Uno et al. (2018).

In this subsection, we employ λ' as a measure of the frequency of forecast revisions. This is because λ can be calculated only for the 1- and 3-year horizons due to the short length of our *Tankan* data, whereas λ' can be calculated for all three time horizons. As Carroll (2003) shows, to interpret λ' theoretically, we need some assumptions on firms' beliefs about the inflation process. In practice, however, when forecasts at different time horizons show a high degree of correlation, the differences between λ and λ' are limited. Since, as Table 1 shows, 71 percent of firms forecast the same year-on-year inflation rates at both 3- and 5-year horizons, it seems that expectations at these horizons are indeed highly correlated. Thus, although its theoretical foundations are weak compared with λ, given the observed strong correlation between forecasts 3- and 5-years ahead, λ' is sufficiently useful as an alternative measure of the frequency of forecast revisions.

Table 2: Frequency of forecast revisions by time horizons

	General prices			Output prices		
	1-year	3-year	5-year	1-year	3-year	5-year
All firms	0.305	0.338	0.350	0.168	0.230	0.250
	(0.303,0.308)	(0.335,0.341)	(0.347,0.353)	(0.166,0.170)	(0.227,0.232)	(0.247,0.253)
Large firms	0.238	0.254	0.252	0.119	0.163	0.170
	(0.232,0.244)	(0.247,0.261)	(0.244,0.260)	(0.114,0.123)	(0.157,0.168)	(0.164,0.177)
S&M firms	0.320	0.354	0.367	0.180	0.243	0.264
	(0.317,0.323)	(0.351,0.357)	(0.363,0.370)	(0.177,0.182)	(0.240,0.245)	(0.261,0.268)
Manu 1	0.299	0.329	0.340	0.171	0.231	0.251
	(0.295,0.303)	(0.325,0.334)	(0.334,0.345)	(0.168,0.174)	(0.227,0.235)	(0.246,0.255)
Manu 2	0.287	0.321	0.332	0.150	0.216	0.236
	(0.282,0.292)	(0.315,0.327)	(0.326,0.339)	(0.147,0.154)	(0.212,0.221)	(0.231,0.242)
Non-Manu	0.309	0.344	0.356	0.167	0.231	0.251
	(0.306,0.312)	(0.340,0.348)	(0.351,0.360)	(0.164,0.169)	(0.227,0.234)	(0.247,0.255)

Notes. Frequencies are quarter. The table reports the mean frequency with 95% bootstrap ($B = 500$) confidence interval in brackets. Large firms are defined as firms with capital of at least 1 billion yen; S&M firms—small- and medium-sized firms—are defined as firms with capital of at least 20 million yen but less than 1 billion yen. Manu 1, Manu 2, and Non-Manu denote manufacturing (basic materials), manufacturing (processing), and non-manufacturing, respectively.

Table 2 presents our estimates of the frequency of forecast revisions by time horizon. Strikingly, the frequency of forecast revisions is higher for both general prices and output prices as the time horizon lengthens. Looking at "All firms," our estimates of the frequency of forecast revisions for general prices at the 1-, 3-, and 5-year horizons are 0.305, 0.338, and 0.350, respectively. The bootstrap confidence intervals suggest that the differences in our estimates across time horizons are statistically significant at the 95 percent level. This is also the case for output prices, where the estimates for the corresponding forecasts are 0.168, 0.230, and 0.250, respectively.

For both general prices and output prices, differences between the estimated frequencies of forecast revision at the 1- and 3-year horizons are larger than the differences between those at the 3- and 5-year horizons. This is in line with the observation in Section 2.3 that for 71 percent of Japanese firms there is no difference in the information contained in the forecasts 5-years and 3-years ahead. When comparing forecasts at 1- and 3-year horizons for both general prices and output prices, forecasts at the 3-year horizon are seen to be more frequently revised for all sectors and firm sizes. For instance, for large firms, the estimated frequency of revision for forecasts of general prices 3-years ahead is 0.254, which is statistically significantly higher than the comparable estimate for the 1-year horizon of 0.238.

This finding runs counter to the view presented in Clark and Davig (2008) and Yellen (2015). In contrast with their analysis, it is unlikely that more frequently revised long-term expectations observed here would be able to capture the long-term inflation trend in the way that they suggest. Furthermore, since using aggregated data prevents researchers from directly observing individual firm behavior with regard to forecast revisions, the use of micro-data to explore the frequency of forecast revision is therefore a distinctive contribution of our approach here. Note that the more frequent

revisions do not necessarily imply larger fluctuations in aggregated long-term inflation expectations. In fact, as reported in Uno et al. (2018), the forecast means at 3- and 5-year horizons moved smoothly during our sample period.

3 On the more frequent revisions of long-term inflation expectations

In this section, we attempt to explain the more frequent revisions of long-term inflation expectations observed in Section 2.4.

3.1 A simple model

In this subsection, we provide an explanation of the more frequent revisions of long-term inflation expectations based on a simple reduced-form model of inflation dynamics.

As mentioned in Section 2.1, the *Tankan* survey asks respondents to select an expected inflation range from a number of options instead of expressing their view freely. Given this responding style, the more frequent revisions of long-term inflation expectations imply that the magnitude of the change in long-term inflation expectations is larger than that in short-term inflation expectations. In light of this relationship between the frequency and magnitude of forecast revisions, in the following, we examine how the magnitudes of changes in long- and short-term inflation expectations are related using a simple reduced-form model of inflation dynamics.

Assume that the inflation rate π_t follows an AR(1) process:

$$\pi_t = (1-\alpha)\tilde{\pi} + \alpha\pi_{t-1} + \epsilon_t, 0 < \alpha < 1, \epsilon_t \sim N(0, \sigma^2), \forall t \tag{1}$$

where $\tilde{\pi}$ denotes the steady state value of inflation, and ϵ_t denotes a shock

to inflation at time t. The inflation rate at time $t+h$ is as follows:

$$\pi_{t+h} = (1-\alpha^{h+1})\tilde{\pi} + \alpha^{h+1}\pi_{t-1} + \sum_{j=0}^{h} \alpha^j \epsilon_{t+h-j}. \tag{2}$$

The magnitude of the shock at time t on the inflation rate at time $t+h$ is then

$$\frac{\partial \pi_{t+h}}{\partial \epsilon_t} = (1-\alpha^{h+1})\frac{\partial \tilde{\pi}}{\partial \epsilon_t} + \alpha^h. \tag{3}$$

The first term on the right hand side indicates the indirect effect of the shock at time t on the inflation rate at time $t+h$ via the steady state value of inflation; the second term on the right hand side shows the direct effect of the shock at time t on the inflation rate at time $t+h$. Assuming that firms form inflation expectations in line with this simple model, the magnitude of the forecast revision is equal to $E[\partial \pi_{t+h}/\partial \epsilon_t \mid \Omega_t]$, where Ω_t represents all information available at time t, $E[\cdot \mid \cdot]$ denotes the conditional expectation operator.

As easily confirmed, with a small α—low persistency in inflation—, the magnitude of the forecast revision $E[\partial \pi_{t+h}/\partial \epsilon_t \mid \Omega_t]$ relates to the magnitude of the change in the steady state value of inflation $\partial \tilde{\pi}/\partial \epsilon_t$ simply. That is, with low persistency in inflation, the change in long-term inflation expectations with a large h is not significantly different from the change in short-term inflation expectations with a small h, regardless of the magnitude of the change in the steady state value of inflation. Conversely, with a large α—high persistency in inflation—, the change in long-term expectations may diverge from that in short-term expectations, depending on the size of the change in the steady state value of inflation. In particular, when the change in the steady state value of inflation is sufficiently large, the impact of the shock on long-term expectations may exceed that on short-term expectations, implying more frequent revisions of forecasts at longer horizons.

Thus, based on this simple reduced-form model, the more frequent revisions of long-term inflation expectations shown in Section 2.4 imply that, since inflation may be considered to have displayed strong inertia during the sample period, the steady state value of inflation must therefore have shifted substantially. This, however, still leaves us with a vital question that remains to be answered: why did the steady state value change substantially in this period?

3.2 Re-anchoring?

The more frequent revisions of long-term inflation expectations may be usefully explained by deploying the concept of "anchoring." The term "anchored" is defined by Bernanke (2007) to mean "*relatively insensitive to incoming data*." Based on such an understanding, it is clear that frequent revisions of longer-term inflation expectations would not take place if those expectations were strongly anchored. However, if firms' revisions were being made in response to new information from the Bank of Japan, this could provide a different explanation for our finding. Following Fujiwara (2005) and Hattori et al. (2016) who argue that CPI forecasts by the BOJ affect those of professional forecasters, we propose a natural experiment to test whether firms are in fact responding to new information from the BOJ.

The Bank of Japan announces the Policy Board members' projections for the CPI in its regularly released *Outlook for Economic Activity and Prices*. The projection period is three years, and the three-year window rolls forward every April. For example, in April 2016, the three-year window for the projection shifts to FY 2016 through FY 2018. This means that new information on the CPI inflation forecast for 3-years ahead arrives every April. At that time, if firms receive no new information on the CPI forecast for 1-year ahead, the BOJ's announcement of its projections can prompt

revisions only of their forecasts for 3-years ahead.

Fortunately for our purposes here, during the period from 2014 and 2016, the BOJ's projections for CPI inflation 1-year ahead have ranged between 0.5% and 1.3%, as Table 3 indicates. As detailed in Section 2.1, the *Tankan* survey provides respondents with ranges, such as "+0.5% ∼ +1.4%" for the option "around +1%," implying that the observed BOJ projection range of 0.5% to 1.3% would not have generated any incentive for firms to revise their forecasts of CPI inflation at the 1-year horizon. This circumstance provides us with a natural experiment for examining firms' responses to new information from the BOJ. Specifically, we employ a difference-in-differences approach which compares the frequency of forecast revisions at the 1- and 3-year horizons between the March and the June survey.

Table 3: CPI projections by the BOJ

	1-year ahead		3-year ahead	
	Jan.	Apr.	Jan.	Apr.
2014	1.3	1.3	–	2.1
2015	1.0	0.8	–	1.9
2016	0.8	0.5	–	1.9

Notes. Figures indicate the medians of the Policy Board members' forecasts announced in the *Outlook for Economic Activity and Prices*.

Table 4 presents the results of our natural experiment. The differences between the mean frequencies of forecast revision at the 1- and 3-year horizons are statistically significantly different from zero both in the March and the June survey. Also, the estimated frequency of forecast revision at the 3-year horizon in the June survey is statistically significantly larger than in the March survey. Consequently, the difference-in-differences is 0.018. This implies that firms have indeed responded to the new information from the BOJ released every April. Note that although the BOJ's forecasts for

3-years ahead have ranged between 1.9% and 2.1% (see Table 3), the mean of firms' forecasts at that horizon has declined from 1.66% to 1.05% during the period—see Table 1 in Uno et al. (2018). This suggests that firms have perceived new information from the BOJ differently.

Table 4: Results of difference-in-differences analysis

	1-year ahead	3-year ahead	Difference
March survey	0.094	0.117	0.023
	(0.090,0.099)	(0.112,0.122)	
June survey	0.088	0.129	0.040
	(0.085,0.092)	(0.125,0.133)	
Change	−0.006	0.012	0.018

Notes. The table reports the mean frequency of forecast revisions with 95% bootstrap ($B = 500$) confidence interval in brackets. We calculate the frequency with which a revision is made only to the forecast for 1-year (3-years) ahead without updating the forecast for 3-years (1-year) ahead.

The results shown in Table 4 contradict the argument in Afrouzi et al. (2015), where it is suggested that firms' managers are uninformed about the central bank's policy and form their inflation expectations using micro-information based on their own shopping experiences. In contrast, our finding indicates that firms exploit macro-information to revise their expectations. In light of the aggressive monetary policy during our sample period, we interpret the responses by firms observed above as a phenomenon characterizing the process of "re-anchoring" long-term inflation expectations.

4 Conclusion

In this paper, we use the same large-scale firm-level micro-data from the *Tankan* survey as in Uno et al. (2018), to examine firms' inflation expectations at different time horizons.

Our principle two findings appear to contradict the arguments of Clark

and Davig (2008) and Yellen (2015) emphasizing the importance of long-term inflation expectations in inflation dynamics. First, with regard to long-term expectations, a number of firms offer no forecasts. Specifically, although only 15 percent of Japanese firms in our *Tankan* survey answer "don't know" for forecasts 1-year ahead, no less than 40 percent of firms respond "don't know" for forecasts 5-year ahead. This suggests that more than 40 percent of firms would answer "don't know" for longer-term forecasts such as 10-year ahead forecasts. Second, more importantly, the frequency of forecast revisions is higher as the time horizon is longer. In general, the more frequently revised long-term expectations are not likely to be able to capture the long-term trend in inflation.

References

[1] Afrouzi, Hassan, Saten Kumar, Olivier Coibion, and Yuriy Gorodnichenko (2015), "Inflation targeting does not anchor inflation expectations: Evidence from firms in New Zealand." *Brookings Papers and Economic Activity*, 2015.

[2] Andrade, Phillip and Herve Le Bihan (2013), "Inattentive professional forecasters." *Journal of Monetary Economics*, 60(8), 967-982.

[3] Bernanke, Ben S. (2007), "Inflation expectations and inflation forecasting." Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, MA, July 10, 2007.

[4] Carroll, Christopher D. (2003), "The epidemiology of macroeconomic expectations." In Lawrence E. Blume and Steven N. Durlauf, eds., *The Economy as an Evolving Complex System, III: Current Perspectives and Future Directions*, Oxford University Press, 2005.

[5] Clark, Todd E. and Troy Davig (2008), "An empirical assessment of the relationships among inflation and short- and long-term expectations." *FRB Kansas City Research Working Papers*, RWP 08-05, Federal Reserve Bank of Kansas City.

[6] Fuhrer, Jeff (2012), "The role of expectations in inflation dynamics." *International Journal of Central Banking*, 8, 137-165.

[7] Fujiwara, Ippei (2005), "Is the central bank's publication of economic forecasts influential?" *Economic Letters*, 89, 255-261.

[8] Hattori, Masazumi, Steven Kong, Frank Packer, and Toshitaka Sekine (2016), "The effects of a central bank's inflation forecasts on private

sector forecasts: Recent evidence from Japan." *Bank of Japan Working Paper Series*, No. 16-E-11, Bank of Japan.

[9] Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve." *Quarterly Journal of Economics*, 117(4), 1295-1328.

[10] Reis, Ricardo (2006), "Inattentive producers." *Review of Economics Studies*, 73(3), 793-821.

[11] Uno, Yosuke, Saori Naganuma, and Naoko Hara (2018), "New facts about firms' inflation expectations: Simple tests for a sticky information model." *Bank of Japan Working Paper Series*, No. 18-E-14, Bank of Japan.

[12] Yellen, Janet L. (2015), "Inflation dynamics and monetary policy." Speech at the Philip Gamble Memorial Lecture, University of Massachusetts, Amherst, Amherst, MA, September 24.