Bank of Japan Working Paper Series

New Facts about Firms' Inflation Expectations: Simple Tests for a Sticky Information Model

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

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

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

No.18-E-14
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: Simple Tests for a Sticky Information Model*

Yosuke Uno† Saori Naganuma‡ Naoko Hara§

August 21, 2018

Abstract

In this paper, we use a large dataset based on firm-level microdata from the *Tankan* survey to examine firms' inflation expectations. We first present two basic findings: (i) firms' inflation expectations are downwardly rigid at zero, and (ii) differences in firms' inflation expectations are larger across firm sizes than across sectors. We then report three findings which are in line with predictions of the simple sticky information model proposed by Mankiw and Reis (2002). First, in each period, a number of firms do not revise their expectations. Second, the frequency of forecast revisions is constant over time. Third, our estimates of the frequency of forecast revisions based on the *Tankan* survey are much smaller than those in previous studies and are much closer to the value that Mankiw and Reis (2002) assumed in their simulation exercises.

JEL codes: E31, E37

Keywords: inflation expectations, frequency of forecast revisions, sticky information model

*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

Expectations of future inflation rates are regarded as an important determinant of current inflation. Particularly crucial in inflation dynamics are the inflation expectations of firms, which, after all, to a large part are responsible for setting prices. Yet, until not long ago, there were relatively few empirical studies on firms' inflation expectations. The reason is that it has been relatively more difficult to obtain data on the inflation expectations of firms than those of households and professional forecasters.

In recent years, however, there have been a number of studies trying to overcome this data constraint. For instance, Coibion et al. (2015) and Afrouzi et al. (2015) conducted a large-scale survey on firms' inflation expectations in New Zealand and analyzed the detailed firm-level data they obtained. They found that firms' average forecasts of inflation have been systematically higher than actual inflation, suggesting that forecast errors of firms' inflation expectations are predictable and sticky. Although they conducted some follow-up surveys of firms from the first wave, the sample size in each wave differs considerably and the intervals between surveys are not constant.

Kaihatsu and Shiraki (2016) performed an analysis using firms' implied forecasts for the GDP deflator—derived from the difference between firms' forecasts for nominal GDP and those for real GDP—as a proxy for firms' inflation expectations. They find that the distribution of inflation expectations and responses to shocks differed significantly across firm characteristics. While the panel data used in their study is well-balanced and their time series data spans a longer period than the data used by Coibion et al. (2015) and Afrouzi et al. (2015), the measure of inflation expectations is calculated indirectly.

Against this background, the aim of the current paper is to examine how

firms formulate their inflation expectation using a dataset that is a substantially larger and more well-balanced panel of direct quantitative information about firms' inflation expectations than the datasets used in previous studies. Specifically, we use the confidential firm-level micro-data underlying the *Short-term Economic Survey of Enterprises in Japan* (*Tankan*), which is a quarterly survey conducted by the Bank of Japan's Research and Statistics Department. The *Tankan* survey covers over 10,000 firms and provides quantitative information on firms' inflation expectations. The panel structure of the *Tankan* survey allows us to observe individual firms' expectation formation over time. Specifically, by using this novel panel dataset, we can directly calculate the frequency of firms' forecast revisions. To the best of our knowledge, no studies so far have examined the frequency of firms' forecast revisions based on micro-data due to data constraints. We use our estimates of the frequency of forecast revisions to test a set of predictions of the simple sticky information model developed by Mankiw and Reis (2002) and Reis (2006).

This paper contributes to the literature in two respects. First, it presents two basic findings on firms' inflation expectations in Japan: (i) firms' inflation expectations are downwardly rigid at zero, and (ii) differences in firms' inflation expectations are larger across firm sizes than across sectors. The latter finding is particularly noteworthy given that there is large heterogeneity across sectors in price changes, meaning that forecast revisions are not simply related to price changes.

The second contribution of the paper is that it relates the observed frequency of forecast revisions to a simple sticky information model. Specifically, we obtain the following three findings, which are in line with the predictions of sticky information models. First, in each period, a certain fraction of firms leave their expectations unchanged. As argued by Reis

(2006), the simple sticky information model assumes that forecast revision is costly, so that even when new information arrives rational firms revise their expectations only infrequently. Note that this finding contrasts with the theoretical prediction of the noisy information model proposed by Sims (2003). Second, the frequency of forecast revisions is constant over time. Again, this is consistent with the simple sticky information model, in which a constant fraction of firms revise their expectations in each period. Third, our estimates of the frequency of forecast revisions based on the *Tankan* survey are much smaller than those for households and professional forecasters reported in previous studies using micro-data and are much closer to the value assumed by Mankiw and Reis (2002) in their simulation exercises.

This paper proceeds as follows. Section 2 provides an overview of the related literature. Section 3 describes the data used in this paper, while Section 4 reports some basic findings on firms' inflation expectations. Next, Section 5 derives a set of testable predictions of the simple sticky information model proposed by Mankiw and Reis (2002) and Reis (2006), while Section 6 present our results of testing these predictions. Section 7 concludes.

2 Related Literature

There are two well-known models to describe firms' inflation expectations formation. One is the sticky information model proposed by Mankiw and Reis (2002) and Reis (2006), which assumes that acquiring information is costly, so that it is optimal for firms to keep their expectations unchanged. Firms that do not revise their expectations at time t do not use information available at time t, implying that their expectations are based on past information before time t. Therefore, forecast errors of these non-revisers can be predicted by using information available at time t. Consequently, in the sticky information model, the aggregate forecast errors at time t are also

predictable.

The second model is the noisy information model proposed by Sims (2003). The model assumes that firms are limited in their capacity to acquire and process information. Since it is impossible for firms to formulate their expectations using all information at time t, the model suggests that the aggregate forecast error is predictable given all information available at time t.

While the two models have different predictions regarding the response of firms' information acquisition to shocks, the speed of convergence of normalized forecast errors to shocks, and so on, Coibion and Gorodnichenko (2015) show that the relationship between average year-ahead inflation forecast errors across agents and average forecast revisions in the two models can be expressed as the same equation. This means that one cannot simply distinguish the two models using aggregate data. However, the key difference between the two models lies in the assumption regarding whether there are firms that do not revise their expectations at time t. Therefore, whether there indeed are such firms can be directly observed using micro-data.

Until recently, micro-data on firms' inflation expectations has been sparse. Therefore, as an alternative, some researchers have used micro-data on professional forecasters' inflation expectations. Although professional forecasters are not explicitly included in both the sticky information model and the noisy information model, as among the most informed economic agents in the economy, they can serve as a benchmark for economic agents' inflation expectations formation. Using micro-data from surveys of professional forecasters, Andrade and Le Bihan (2013) and Dovern et al. (2014) measured the share of professional forecasters revising their expectations, i.e., the frequency of forecast revisions. They found that there is indeed a fraction of professional forecasters whose expectations remain unchanged; at the

same time, however, they highlight that forecasts are revised much more frequently than assumed by Mankiw and Reis (2002) in their simulation exercises.

Similarly, Pfajfar and Santoro (2013) and Hori and Kawagoe (2013), using household micro-data for the United States and Japan, find that there is also a fraction of households who do not revise their expectations, which is in line with Carroll's (2003) epidemiological model that is consistent with the intuition behind the sticky information model of Mankiw and Reis (2002). Pfajfar and Santoro (2013) and Hori and Kawagoe (2013) find that like the results for professional forecasters, the estimates for the frequency of forecast revisions of households are also much larger than the plausible parameter value.

Thus, previous empirical studies using micro-data on professional forecasters and households have shown that non-revisers do exist, which is in line with the predictions of the sticky information model developed by Mankiw and Reis (2002) and Reis (2006). Yet, at the same time, the previous empirical studies find that the share of agents revising their forecasts at time t is much larger than the value assumed by Mankiw and Reis (2002) in their simulation exercises. This means that the evidence from micro-data on professional forecasters and households is *qualitatively* supportive of the sticky information model, but not *quantitatively*.

In contrast with professional forecasters and households, it has been more difficult to obtain data on the inflation expectations of firms. Against this background, a number of researchers have tried to overcome such data constraints using different approaches. Coibion et al. (2015) and Afrouzi et al. (2015), for example, conducted a large-scale survey of firms' inflation expectations in New Zealand between September 2013 and August 2015 and found that forecast errors of firms' inflation expectations are predictable

and sticky.[1] Meanwhile, Richards and Verstraete (2016) used micro-data from a quarterly survey conducted by the Bank of Canada. While the size of the sample in a cross-sectional dimension is relatively small with about 100 firms, it is relatively large in a time-series dimension. Taking advantage of the fact that the data cover a relatively long period, they use the data to examine firms' inflation expectations to find that expectations are not fully rational but also are not simply adaptive either. Taking a different approach, Kaihatsu and Shiraki (2016) used firms' implied forecasts of the GDP deflator—derived from the difference between firms' forecasts of nominal GDP and those of real GDP—as a proxy for firms' inflation expectations. Specifically, they used the *Annual Survey of Corporate Behavior* conducted by the Cabinet Office of Japan and found that the distribution of inflation expectations and responses to shocks differed significantly across firm characteristics. Finally, Inamura et al. (2017), using the same data as this paper, report that firms' inflation expectations seem to differ across firms of different sizes.[2]

Thus, unlike the studies on professional forecasters' and households' inflation expectations, previous studies using firms' micro-data do not examine the frequency of firms' forecast revisions. Therefore, whether there are firms that do not revise their expectations at time t and, if so, how many do not change their forecasts is still an open question.

[1] The sample sizes of the survey by Coibion et al. (2015) and Afrouzi et al. (2015) are 3,153 firms in the first wave, 714 firms in the second wave, 1,607 firms in the third wave, 1,257 firms in the fourth wave, and 50 firms in the fifth wave.

[2] Other studies using micro-data on firms' inflation expectations are those by Bryan et al. (2015) and Cloyne et al. (2016). However, these studies do not focus on the formation mechanism of firms' inflation expectations.

3 Data

The data used for the analysis in this paper consist of the confidential micro-data gathered by the Bank of Japan's Research and Statistics Department for the *Tankan* survey, a large-scale firm-level survey. This section provides an overview of the *Tankan* data and discusses how our data differs from the data used in previous studies.

3.1 *Tankan* data

The *Tankan* survey consists of three parts. The first part focuses on qualitative data representing firms' subjective judgment on business conditions, their capacity utilization, and so on. The second part consists of quantitative data providing firms' projections of their annual sales, fixed investment, etc. Finally, the third part also provides quantitative data, focusing on firms' inflation expectations for one year ahead, three years ahead, and five years ahead. In this paper, we use this third part of the *Tankan* survey.

The *Tankan* gathers 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 to the consumer price index (CPI). Firms are provided with ten options starting with (1) around +6% or higher, (2) around +5%, (3) around +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.[3] In the analysis that follows, we exclude firms that chose options (11), (12), or (13). We should also note that the *Tankan* survey asks firms to respond "excluding the effects of institutional changes such as changes in the consumption tax rate."

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%, and continuing in five-percentage point intervals to (9) around −20% or lower, and (10) don't know. Firms that chose option (10) are excluded from our analysis. 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 are an unbalanced panel in which there is a core of firms that have been surveyed since the start (May 1974), while some have dropped out because they fell below 20 million yen in paid in capital and/or went bankrupt. To maintain the sample size, firms are newly added every two or three years. In our sample, about 1,000 firms were added in

[3] 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/

March 2015. The panel structure of the *Tankan* survey allows us to track individual firms' expectations formation over time.

3.2 Differences from data in previous studies

In this subsection, we compare our *Tankan* data with survey data used in previous studies, focusing on the format in which responses are provided and what kind of prices respondents should focus on in their responses with regard to inflation expectations.

The *Consumer Confidence Survey* used by Hori and Kawagoe (2013) to examine inflation expectations of households in Japan, like the *Tankan* survey, provides respondents with various options to choose from. Specifically, the *Consumer Confidence Survey* asks: "During the next 12 months, do you think that prices of goods and services that you frequently purchase will go down, up, or remain the same?" Respondents are asked to choose among the following answers: (1) -10% or more, (2) -10% to -5%, (3) -5% to -2%, (4) -2% or less, (5) around 0%, (6) 2% or less, (7) 2% to 5%, (8) 5% to 10%, (9) 10% or more, and (10) don't know. It should be noted that this survey asks about the prices of frequently-purchased items rather than the CPI.

In contrast to the *Tankan* survey and the *Consumer Confidence Survey*, in the survey conducted by Coibion et al. (2015) and Afrouzi et al. (2015) in New Zealand, respondents—in answering to the question, "During the next twelve months, by how much do you think prices will change overall in the economy?"—are allowed to express their views freely instead of choosing from a number of options. Similarly, the *Survey of Professional Forecasters* conducted quarterly by the European Central Bank (ECB) of approximately 90 professional forecasters and used in Andrade and Le Bihan's (2013) study, asks respondents to provide their views on the inflation rate in the euro area

in their own words. It should be noted that the ECB survey explicitly refers to the expected year-on-year rate of change in the Harmonised Index of Consumer Prices.

Meanwhile, the University of Michigan's monthly *Surveys of Consumers* used in Pfajfar and Santoro's (2013) studies combines both response formats. That is, respondents can choose from options as well as give their opinion in their own words. Specifically, presented with the question, "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?," respondents are provided with the following four options: (1) go up, (2) stay the same, (3) go down, and (4) don't know. Respondents that choose either option (1) or (3) are asked to fill in a quantitative answer expressed in percentage terms.[4]

These differences in response formats potentially are important when examining respondents' forecast revisions. That is, in the *Tankan* and the *Consumer Confidence Survey*, even if respondents revise their forecast, if the forecast remains within the range provided by the options, their answer will remain unchanged. This means that these surveys potentially underestimate the frequency with which respondents revise their forecasts, while such revisions would be picked up in surveys that allow respondents to provide forecasts in their own words, such as the ECB's *Survey of Professional Forecasters* and the University of Michigan's *Surveys of Consumers*. In addition, it should be noted that only the ECB's *Survey of Professional Forecasters* like the *Tankan* explicitly asks about inflation expectations with regard to the consumer price index.

[4]The Michigan survey includes a question not only about the next twelve months but also about the next five to ten years, and the answers are to be provided in the same format.

4 Basic findings

In this section, we present some basic findings about firms' inflation expectations.

4.1 Basic statistics

Table 1 reports the basic statistics of our sample. We begin by focusing on general prices in the upper part of Table 1, which shows the basic statistics for the different surveys during our observation period from March 2014 to September 2017. The table indicates that the average of one-year ahead inflation expectations in the last survey of our observation period was 0.69 percent, which is lower than the average inflation expectations for three and five years ahead.[5]

The cross-sectional standard deviation for one, three, and five years ahead in the September 2017 survey were 0.96 percent, 1.09 percent, and 1.26 percent, respectively. Like the mean, the standard deviation is smaller for shorter-term forecasts.

The degree of skewness for one, three, and five years ahead in the September 2017 survey was 1.46, 0.89, and 0.70, respectively. This means that the right side of the distribution was more long-tailed for all time horizons. Moreover, the skewness of one-year ahead expectations increased rapidly during our observation period, rising from 0.81 in March 2014 to 1.46 in September 2017. To show this sharp increase in skewness more clearly, Figure 1 shows the histogram of inflation expectations over time. As can be seen, the shape of the histogram gradually changes, with the number of firms with positive inflation expectations gradually decreasing, the number of firms expecting zero inflation increasing substantially, and the number

[5] We examine how different inflation expectation levels and frequencies of forecast revisions are at different time horizons in our companion paper (Uno et al. (2018)).

Table 1: Basic statistics

		1-year			3-year			5-year		
time	N	Mean	SD	SKW	Mean	SD	SKW	Mean	SD	SKW
General prices										
Mar.14	6129	1.48	1.35	0.81	1.66	1.42	0.72	1.72	1.60	0.58
Jun.14	6230	1.48	1.28	0.85	1.64	1.38	0.73	1.70	1.56	0.67
Sep.14	6152	1.49	1.32	0.87	1.63	1.40	0.70	1.69	1.57	0.61
Dec.14	5993	1.39	1.28	0.93	1.62	1.36	0.68	1.67	1.52	0.61
Mar.15	6723	1.35	1.25	1.05	1.56	1.30	0.74	1.64	1.47	0.63
Jun.15	6666	1.32	1.19	1.09	1.52	1.25	0.70	1.59	1.43	0.61
Sep.15	6569	1.16	1.17	1.11	1.40	1.24	0.78	1.48	1.40	0.64
Dec.15	6533	1.03	1.11	1.16	1.30	1.20	0.77	1.38	1.35	0.64
Mar.16	6416	0.84	1.10	1.24	1.12	1.17	0.82	1.21	1.33	0.74
Jun.16	6385	0.72	1.02	1.32	1.05	1.15	0.74	1.14	1.31	0.57
Sep.16	6318	0.64	0.99	1.31	0.99	1.11	0.77	1.05	1.28	0.60
Dec.16	6268	0.66	0.99	1.48	1.00	1.11	0.90	1.06	1.28	0.70
Mar.17	6269	0.70	0.96	1.47	1.03	1.08	0.87	1.07	1.25	0.76
Jun.17	6289	0.75	0.97	1.40	1.07	1.09	0.86	1.12	1.26	0.76
Sep.17	6315	0.69	0.96	1.46	1.03	1.09	0.89	1.06	1.26	0.70
Output prices										
Mar.14	6788	1.13	3.73	0.56	1.84	5.14	−0.10	2.08	6.40	−0.36
Jun.14	6818	1.14	3.50	0.61	2.00	5.06	−0.12	2.30	6.29	−0.37
Sep.14	6680	1.06	3.46	0.42	1.80	4.98	−0.09	2.12	6.17	−0.36
Dec.14	6500	1.01	3.33	0.60	1.79	4.82	−0.17	2.07	5.92	−0.42
Mar.15	7193	0.98	3.37	0.46	1.81	4.70	−0.12	2.16	5.79	−0.36
Jun.15	7112	1.00	3.20	0.60	1.81	4.61	−0.15	2.07	5.80	−0.39
Sep.15	7035	0.72	3.12	0.44	1.53	4.57	−0.22	1.77	5.68	−0.47
Dec.15	6979	0.59	2.95	0.41	1.36	4.41	−0.25	1.63	5.54	−0.54
Mar.16	6942	0.43	2.89	0.40	1.09	4.34	−0.21	1.33	5.43	−0.50
Jun.16	6884	0.31	2.73	0.19	0.91	4.22	−0.32	1.13	5.26	−0.55
Sep.16	6784	0.28	2.69	0.41	0.88	4.14	−0.24	1.01	5.24	−0.51
Dec.16	6712	0.36	2.68	0.59	0.93	4.18	−0.25	1.07	5.23	−0.53
Mar.17	6752	0.39	2.68	0.33	0.98	4.11	−0.42	1.10	5.18	−0.58
Jun.17	6722	0.48	2.74	0.37	1.04	4.14	−0.37	1.17	5.17	−0.60
Sep.17	6726	0.50	2.77	0.63	1.06	4.19	−0.31	1.16	5.22	−0.59

Notes. SD and SKW denote the standard deviation and the skewness.

of firms with negative inflation expectations remaining largely unchanged. This evolution of the histogram suggests that firms' inflation expectations are downwardly rigid at zero. This result is line with Kamada's (2013) observation that households' inflation expectations in Japan are also downwardly rigid. This means that downward rigidity in inflation expectations can be regarded as a feature common to both households and firms in Japan.

Figure 1: Histogram of inflation expectations over time

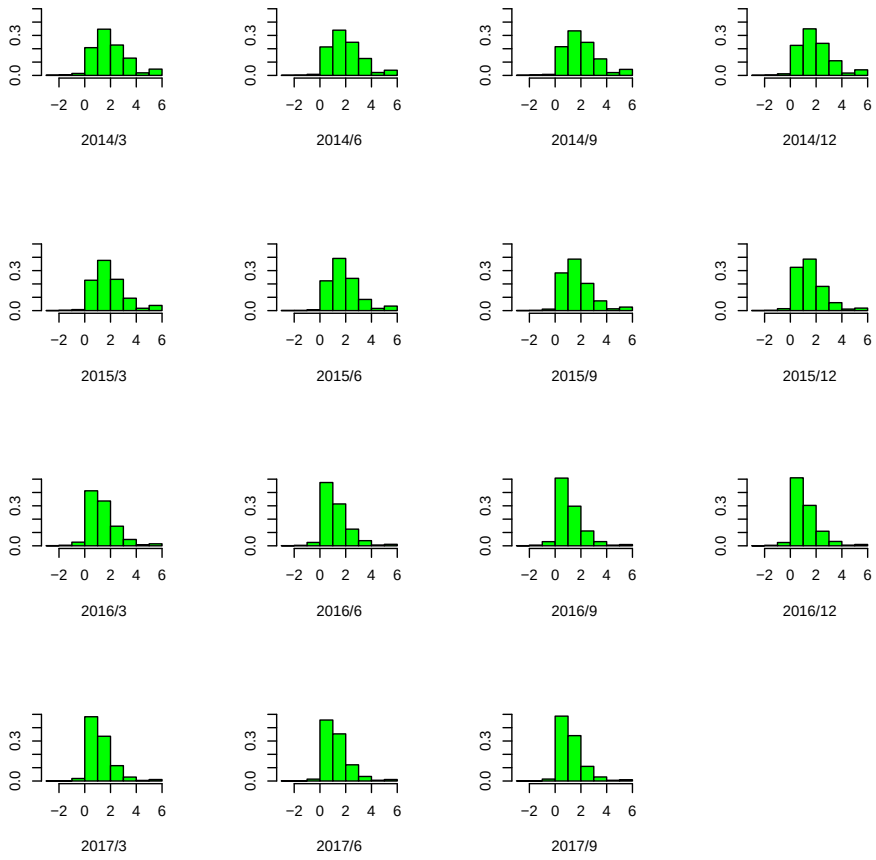

We now turn from general prices to output prices, basic statistics for which are shown in the lower part of Table 1. As mentioned in Section 3,

forecasts for output prices three and five years ahead are for the cumulative change relative to the current level, so that forecasts for the different horizons are not readily comparable.

Looking at the basic statistics for the September 2017 survey, the average of inflation expectations for one year ahead is 0.50 percent. This is lower than that for general prices, suggesting that firms are more pessimistic about their output prices than general prices. In contrast to general prices, the skewness of three- and five-year ahead forecasts is negative, meaning that the left side of the distribution has a longer tail. This implies that only a relatively small share of firms expect their output prices to increase substantially in the medium to long term.

4.2 Rationality

The time series data of average inflation expectations for one year ahead allow us to calculate ex-post forecast errors. Figure 2 reports two types of ex-post forecast errors based on general prices and output prices. The green line in Figure 2 represents the ex-post forecast errors based on general prices (i.e., the CPI). The line indicates that ex-post forecast errors have a negative bias and are persistent—the lag-1 autocorrelation is 0.75, implying that firms' inflation expectations for general prices do not satisfy the rationality assumption.

Unlike general prices, realized output prices are not available, so that it is not possible to calculate ex-post forecast errors of output prices. However, we can compute forecasts for a pseudo-CPI using the output prices of firms belonging to "Retailing," "Services for individuals, and accommodations," "Eating & drinking places," which are included in the CPI. This means that for some sectors, the ex-post forecast errors of output prices can be calculated by comparing the forecasts for their output prices—the pseudo-

CPI—with the actual CPI. The blue line in Figure 2 represents the ex-post forecast errors for the pseudo-CPI. The ex-post forecast errors also have a negative bias and are persistent—the lag-1 autocorrelation is 0.67—which is similar to the result for the CPI, implying that forecasts for output prices also do not satisfy the rationality assumption.

Figure 2: Ex-post forecast errors

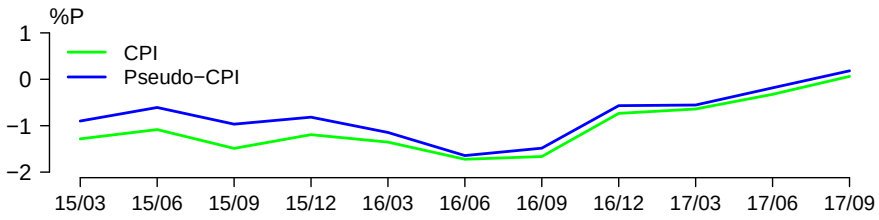

4.3 Differences in inflation expectations across firm sizes and sectors

Firm size Using the *Tankan* data, Inamura et al. (2017) find that firms' inflation expectations depend on their size. In this subsection, we provide more detailed evidence on the difference in inflation expectations across firm sizes. As shown in Table 2, large firms have lower average inflation expectations with regard to both general prices and output prices than small and medium-sized firms. This is also the case for the standard deviations. Given that actual CPI inflation was around zero during our observation period, the lower positive average inflation expectations of large firms imply that their forecasts were relatively accurate. In contrast, Coibion et al. (2015) found that the forecast errors of large firms were larger than those of small and medium-sized firms. The difference between our results and theirs may due to differences in the size of large firms: in the sample used by Coibion et al. (2015), even the largest firm had only 698 employees, while

in our *Tankan* data, the largest firm had 163,406 employees.

Table 2: Basic statistics by firm size and sector

	General prices (1-year)				Output prices (1-year)			
	N	Mean	SD	SKW	N	Mean	SD	SKW
All firms	135600	1.06	1.20	1.16	146272	0.64	3.18	0.30
Large firms	23991	0.82	0.94	1.14	26895	0.36	2.73	0.10
S&M firms	111609	1.11	1.24	1.11	119377	0.71	3.27	0.31
Manu 1	20130	1.08	1.22	1.14	21726	0.83	3.42	0.37
Manu 2	33967	1.00	1.14	1.15	36980	0.14	3.09	−0.16
Non-Manu	81503	1.07	1.22	1.16	87566	0.81	3.13	0.46

Notes. SD and SKW denote the standard deviation and the skewness. 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.

There is also a striking difference in skewness across firm sizes. While there is little difference in the skewness of expectations for general prices, for output prices there is a large difference in the skewness of expectations. Specifically, while the skewness of expectations for output prices is around zero for large firms, it is clearly positive for small and medium-sized firms, implying that a small number of small and medium-sized firms expect their output prices to go up significantly.

Sector Next, we investigate differences in inflation expectations across sectors. The lower three rows of Table 2 report the basic statistics for the manufacturing (basic materials) sector (Manu 1), the manufacturing (processing) sector (Manu 2), and the non-manufacturing sector (Non-manu). Broadly speaking, expectations for general prices appear to be homogeneous across sectors, which contrasts with the differences observed across firm sizes. On the other hand, expectations for output prices exhibit relatively large differences. Specifically, the mean, standard deviation, and

skewness of expectations are substantially lower among firms in the manufacturing (processing) sector than in the manufacturing (basic materials) and non-manufacturing sectors.

5 Simple tests for the sticky information model

The panel structure of our *Tankan* data allows us to examine how firms revise or update their inflation expectations. In this section, we focus on the simple sticky information model proposed by Mankiw and Reis (2002) and Reis (2006), which provides a number of predictions on firms' forecast revisions, which can be directly tested with our *Tankan* data.

5.1 Overview of the model

We start by providing a brief overview of the simple sticky information model presented by Mankiw and Reis (2002) and Reis (2006). The key assumption is that it is costly for firms to acquire, absorb, and process information. Given these costs, firms optimally choose the period of being inattentive to maximize their expected profits conditional on the information they have. Based on the additional assumptions that firms do not have private information and make their choices independently, Reis (2006) derives that the share of firms which revise their forecasts is constant over time and equal to $0 < \lambda < 1$ as follows:

$$F_t = \lambda E[\pi_t \mid \Omega_t] + (1 - \lambda)F_{t-1}, \tag{1}$$

where F_t is the average of inflation expectations in the economy, π_t is the inflation rate, Ω_t represents all information available at time t, and $E[\cdot \mid \cdot]$ denotes the conditional expectations operator. It should be noted that fraction λ of firms have access to all information available at time t and formulate their expectations in line with this model.

5.2 Testable predictions of the model

In this subsection, we discuss a set of empirically testable predictions of the simple sticky information model. Although the predictions are straightforward to obtain, so far it has been impossible to empirically test them due to the lack of firm-level micro-data. This study therefore is the first attempt to test the predictions directly.

Testable prediction 1 There is a fraction $(1 - \lambda)$ of firms which do not revise their forecasts at time t, i.e., $0 < \lambda < 1$. That is, a fraction $(1 - \lambda)$ of firms optimally choose to be inattentive to new information at time t. This is one of the most important predictions of the model, which contrasts with the theoretical prediction of the noisy information model proposed by Sims (2003), which competes with the sticky information model.

The frequency of forecast revisions has been empirically examined using survey data on professional forecasters' and households' expectations by Andrade and Le Bihan (2013), Hori and Kawagoe (2013), and Pfajfar and Santoro (2013). They directly observed whether all professional forecasters or households frequently revised their forecasts over time using micro-data with a panel structure and found that, just as the simple sticky information model predicts, a fraction of professional forecasters or households did not revise their forecasts in every period.

Testable prediction 2 The distribution of inattentiveness across firms tends to the exponential distribution with parameter λ. As a result, at the aggregate level, λ is constant over time. As Reis (2006) has shown, this can be directly derived from the assumption that each firm randomly and independently revises its forecasts.

However, Pfajfar and Santoro (2013) and Andrade and Le Bihan (2013)

find that λ does vary over time in the case of households and professional forecasters, which is inconsistent with this prediction of the model.

Testable prediction 3 There is no disagreement among firms that revise their forecasts with regard to the inflation rate they forecast. This is directly derived from the assumption that a fraction λ of firms rationally revise their forecasts based on the same information set at time t. That is, the simple sticky information model predicts that firms that revise their forecast do not disagree such that at time t one firm revises its inflation forecast for the CPI to, say, 2 percent, while another revises it to 3 percent, since they have access to the same information set and formulate their forecasts based on the same model. Note that in our *Tankan* data, we can only examine disagreement with regard to expectations for general prices, but not output prices, which are specific to each firm.

Professional forecasters who revise their forecast, however, tend to disagree substantially, as shown by Andrade and Le Bihan (2013).

5.3 Two measures of forecast revisions

To empirically test the testable predictions discussed in Section 5.2, we need to specify how to measure forecast revisions in our *Tankan* data. This subsection presents two measures of forecast revisions.

In general, a forecast revision consistent with the theoretical model is defined as a change in the forecast from time $t-1$ to time t with respect to an event at time $t+s$. One important feature is that the forecasting horizon $t+s$ is fixed. As described in Section 3, in the *Tankan* survey, respondents provide rolling horizon forecasts for one, three, and five years ahead. Accordingly, a forecast revision consistent with the model is defined, for example, by comparing the current forecast for one year ahead with the forecast for three years ahead of two years ago. Formally, for any time t and

firm i, the measure of a forecast revision consistent with theory is as follows:

$$r_{i,t,t+s} = E_{i,t}[\pi_{t+s}] - E_{i,t-8}[\pi_{t+s+8}] \tag{2}$$

with $s = 4$, 12, and 20, where $r_{i,t,t+s}$ indicates the forecast change from time $t-8$ to time t for the forecasting horizon $t+s$, $E_{i,t}[\cdot]$ is the expectations operator conditional on firm i's belief on the state of the economy at time t, and π_{t+s} denotes the inflation rate at time $t+s$.

The revision interval is set to eight quarters due to the design of the *Tankan* survey, which may potentially result in downward bias in the estimation of the frequency of forecast revisions. The reason is that we count only one forecast revision during the eight quarters, regardless of how many times firms actually revised their forecasts during the eight quarters. In addition, it should be noted that this measure can be calculated only for general prices, not output prices. As described in Section 3, regarding output prices for three and five years ahead, firms are not asked to answer in terms of year-on-year changes but in terms of the cumulative changes compared to the current level. This implies that we cannot calculate $r_{i,t,t+s}$ for output prices using the forecasts of the different horizons.

The alternative measure for forecast revisions is simply defined as forecast changes from the previous survey:

$$r'_{i,t,t+s} = E_{i,t}[\pi_{t+s}] - E_{i,t-1}[\pi_{t+s-1}]. \tag{3}$$

As shown by Carroll (2003), if certain additional assumptions on agents' views about the inflation process are satisfied, $r'_{i,t,t+s}$ is consistent with theory, i.e., $E_{i,t-1}[\pi_{t+s-1}] = E_{i,t-1}[\pi_{t+s}]$.

Following Andrade and Le Bihan (2013) and others, for these two mea-

sures, the frequency of forecast revisions is defined as

$$\lambda_{t,t+s,8} = \frac{1}{N_t}\sum_{i=1}^{N_t} I(r_{i,t,t+s} \neq 0), \quad \lambda'_{t,t+s} = \frac{1}{N_t}\sum_{i=1}^{N_t} I(r'_{i,t,t+s} \neq 0) \tag{4}$$

where N_t denotes the number of respondents to the survey at time t, $I(x \neq 0)$ is an indicator function equal to 1 if $x \neq 0$ and 0 otherwise. Note that when comparing $\lambda_{t,t+s,8}$ and $\lambda'_{t,t+s}$, units should be adjusted to quarterly values. Here, assuming that the frequency of forecast revisions during the eight quarters is constant, $\lambda_{t,t+s,8}$ is converted to a quarterly rate, i.e., $\lambda_{t,t+s} = 1 - (1 - \lambda_{t,t+s,8})^{1/8}$.

6 Results

This section presents our results of testing the simple predictions discussed in the previous section.

Result for test 1: Is λ less than 1? The results for this test are presented in Table 3. Four results in the table stand out. First, at the aggregate level—represented by the results for "All firms" in the table—the frequency of forecast revisions is less than one regardless of whether this is measured using λ or λ'. Our estimates of λ are 0.139 and 0.115 per quarter, while those of λ' are 0.305 and 0.168 per quarter. Note that λ is lower than λ' because of the downward bias due to the forecast revision interval of λ' being defined as eight quarters as mentioned in Section 5.3.

Second, our estimates of the frequency of forecast revisions for firms are much smaller than those for households and professional forecasters found in previous studies using micro-data. For instance, Andrade and Le Bihan (2013) and Dovern et al. (2014) report that the frequency of forecast revisions of professional forecasters ranges between 0.7 and 0.9 per quarter, while Pfajfar and Santoro (2013) and Hori and Kawagoe (2013) respectively

Table 3: Frequency of forecast revision

	λ		λ'	
	General price		1-year	
	1- vs. 3-year	3- vs. 5-year	General price	Output price
All firms	0.139	0.115	0.305	0.168
	(0.137,0.140)	(0.114,0.116)	(0.303,0.308)	(0.166,0.170)
Large firms	0.112	0.091	0.238	0.119
	(0.109,0.116)	(0.088,0.095)	(0.232,0.244)	(0.114,0.123)
S&M firms	0.145	0.120	0.320	0.180
	(0.144,0.147)	(0.118,0.121)	(0.317,0.323)	(0.177,0.182)
Manu 1	0.141	0.113	0.299	0.171
	(0.138,0.143)	(0.110,0.115)	(0.295,0.303)	(0.168,0.174)
Manu 2	0.139	0.110	0.287	0.150
	(0.136,0.142)	(0.108,0.113)	(0.282,0.292)	(0.147,0.154)
Non-Manu	0.138	0.117	0.309	0.167
	(0.136,0.140)	(0.115,0.119)	(0.306,0.312)	(0.164,0.169)

Notes. Frequencies are quarterly. The table reports the mean frequency with the 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. "1- vs. 3-year" denotes the forecast revision comparing the current forecast for one year ahead with the forecast for three years ahead of two years ago. "3- vs. 5-year" denotes the forecast revision comparing the current forecast for three years ahead with the forecast for five years ahead of two years ago.

found that those of households is 0.98 and 0.86 per quarter.[6] At the same time, we should highlight that our estimates are much closer to the value that Mankiw and Reis (2002) assumed in their simulation exercises—they set $\lambda = 0.25$. However, we should add a caveat to our estimates of the frequency of forecast revisions. As described in Section 3.2, when comparing our estimates to those found in previous studies, the way in which survey participants were asked to respond matters. Since the University of Michigan's *Surveys of Consumers* used by Pfajfar and Santoro (2013) and the *Consumer Confidence Survey* used by Hori and Kawagoe (2013) employ the same response format as our *Tankan* data, the results are comparable. On the other hand, the ECB's *Survey of Professional Forecasters* used by Andrade and Le Bihan (2013) asks respondents to provide their forecasts in words instead of choosing among options, partly explaining why our estimates tend to be lower than those obtained Andrade and Le Bihan (2013).

Third, by firm size, the frequency of forecast revisions is less than one both for large firms and for small and medium-sized firms. Moreover, the estimates for large firms are statistically significantly lower than those for small and medium-sized firms. This suggests that based on the simple sticky information model, the costs for large firms to update their forecasts was relatively large, or that the shocks faced by large firms were relatively small. If it is more costly for large firms to revise their forecast, it is unsurprising that their forecasts, as shown in Section 4.3, are more accurate than those of small and medium-sized firms.

Fourth, by sector, the frequency of forecast revisions is less than one in both the manufacturing and non-manufacturing sectors. Moreover, in contrast with the substantial differences in the frequency of forecast revisions between large firms on the one hand and small and medium-firms on the

[6] Their data are monthly observations, so that we convert their estimates to quarterly values.

other, hardly any difference is observed across sectors. For instance, the estimates of λ based on the 1-year vs. 3-year comparison are not statistically significantly different across sectors.

Combining these results with the findings presented in Section 4.3 suggests that both in terms of the level of inflation forecasts and in terms of the frequency of forecast revisions, differences across firm sizes but not across sectors can be observed. This contrasts with the observation that when it comes to the frequency of prices changes, as highlighted by Nakamura and Steinsson (2008), substantial heterogeneity across sectors can be observed. Taken together, these findings suggest that forecast revisions are not related simply to price changes. Exploring this issue in more detail provides an interesting avenue for future research.

Result for test 2: Is λ constant over time? In order to examine the interval of forecast revisions of individual firms, panel data covering a sufficiently long period is required. Our *Tankan* data, however, as outlined in Section 3, only spans the period from March 2014 to September 2017, i.e., fifteen quarters, which is too short to obtain unbiased estimates of forecast revision intervals.

Given this limitation of our *Tankan* data, we confine our examination to whether λ and λ' are constant over time or not.

Figure 3: λ and λ' over time

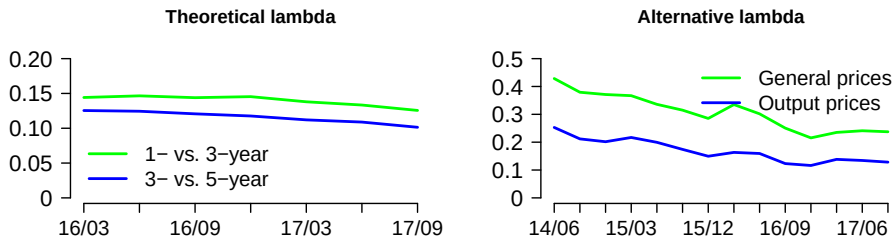

Figure 3 presents the evolution of λ and λ'. The left panel indicates that λ appears to be more or less constant over time, while the right panel suggests that λ' tends to decline over time. Although it is possible that our estimate of λ is biased downward, it seems reasonable to assume that the size of this bias does not vary over time. Therefore, when examining developments over time, λ is likely more reliable than λ'. The fact that λ is more or less constant provides evidence supporting the simple sticky information model.

Result for test 3: Is there disagreement among firms that revise their forecasts? To examine this, we measured, for each period, the disagreement—the cross-sectional standard deviation—of general price expectations among firms that revised their forecasts. As Figure 4 clearly shows, the disagreement among forecast revisers differs from zero for all time horizons in all periods. This finding is similar to the result obtained by Andrade and Le Bihan (2013) for professional forecasters and is not consistent with the prediction of the simple sticky information model.

Figure 4: Disagreement of expectations for general prices

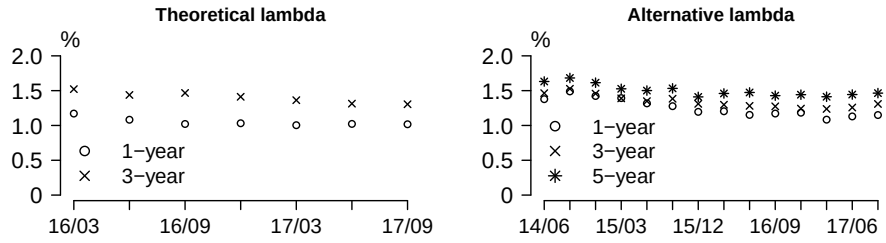

7 Conclusion

In this paper, we presented new evidence on firms' inflation expectations. Using a large dataset based on micro-data from the *Tankan* survey, we first obtained two basic findings on firms' inflation expectations in Japan: (i) firms' inflation expectations are downwardly rigid at zero, and (ii) differences in firms' inflation expectations are larger across firm sizes than across sectors. The latter finding is particularly noteworthy given that there is large heterogeneity across sectors in price changes, meaning that forecast revisions are not simply related to price changes.

We then provided simple tests for the sticky information model proposed by Mankiw and Reis (2002) and Reis (2006). We obtained the following three findings, which are in line with the predictions of the sticky information model. First, in each period, a certain fraction of firms leave their expectations unchanged. Note that this finding contrasts with the theoretical prediction of the simple noisy information model proposed by Sims (2003). Second, the frequency of forecast revisions is constant over time. Third, our estimates of the frequency of forecast revisions based on the *Tankan* survey are much smaller than those for households and professional forecasters reported in previous studies using micro-data and are much closer to the value assumed by Mankiw and Reis (2002) in their simulation exercises.

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] Bryan, Michael F., Brent H. Meyer, and Nicholas B. Parker (2015), "The inflation expectations of firms: What do they look like, are they accurate, and do they matter?" *Working Paper*, Federal Reserve Bank of Atlanta.

[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] Cloyne, James, Lena Boneva, Martin Weale, and Tomasz Wieladek (2016), "Firms' expectations and price-setting: Evidence from micro data." *External MPC Unit Discussion Paper*, No. 48, Bank of England.

[6] Coibion, Olivier and Yuriy Gorodnichenko (2015), "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8), 2644-2678.

[7] Coibion, Olivier, Yuriy Gorodnichenko, and Saten Kumar (2015), "How do firms form their expectations? New survey evidence." *NBER Working Paper Series*, No. 21092, National Bureau of Economic Research.

[8] Dovern, Jonas, Ulrich Fritsche, Prakash Loungani, and Natalia Tamirisa (2014), "Information rigidities: Comparing average and individual forecasts for a large international panel." *International Journal of Forecasting*, 31(1), 144-154.

[9] Hori, Masahiro and Masaaki Kawagoe (2013), "Inflation expectations of Japanese households: Micro evidence from a consumer confidence survey." *Hitotsubashi Journal of Economics*, 54(1), 17-38.

[10] Inamura, Kouki, Kazunori Hiyama, and Kyosuke Shiotani (2017), "Inflation outlook and business conditions of firms: Evidence from the Tankan survey." *IFC Bulletin* 43, https://ideas.repec.org/h/bis/bisifc/43-30.html

[11] Kaihatsu, Sohei and Noriyuki Shiraki (2016), "Firms' inflation expectations and wage-setting behaviors." *Bank of Japan Working Paper Series*, No. 16-E-10, Bank of Japan.

[12] Kamada, Koichiro (2013), "Downward rigidity in households' price expectations: An analysis based on the Bank of Japan's opinion survey on the general public's views and behavior'." *Bank of Japan Working Paper Series*, No. 13-E-15, Bank of Japan.

[13] 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.

[14] Nakamura, Emi and Jon Steinsson (2008), "Five facts about prices: A reevaluation of menu cost models." *Quarterly Journal of Economics*, 123(4), 1415-1464.

[15] Pfajfar, Damjan and Emiliano Santoro (2013), "News on inflation and the epidemiology of inflation expectations." *Journal of Money, Credit and Banking*, 45(6), 1045-1067.

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

[17] Richards, Simon and Matthier Verstraete (2016), "Understanding firms' inflation expectations using the Bank of Canada's Business Outlook Survey." *Staff Working Paper*, 2016-7, Bank of Canada.

[18] Sims, Christopher A. (2003), "Implications of rational inattention." *Journal of Monetary Economics*, 50(3), 665-690.

[19] Uno, Yosuke, Saori Naganuma, and Naoko Hara (2018), "New facts about firms' inflation expectations: Short- versus long-term inflation expectations." *Bank of Japan Working Paper Series*, No. 18-E-15, Bank of Japan.