ECONOMIC RESEARCH CENTER
DISCUSSION PAPER

E-Series

No.E26-4

Overreaction in inflation expectations: Do behavioral attributes of individuals matter?

by

Takayuki Tsuruga
Zichong Yue

July 2026

ECONOMIC RESEARCH CENTER
GRADUATE SCHOOL OF ECONOMICS
NAGOYA UNIVERSITY

Overreaction in inflation expectations: Do behavioral attributes of individuals matter?*

Takayuki Tsuruga† and Zichong Yue‡

July 8, 2026

Abstract

Previous studies on inflation expectations suggest that agents often overreact to public signals, motivating extensions of standard macroeconomic models beyond full-information rational expectations and pointing to a role for behavioral explanations. Using Japanese survey data from 2008 to 2024 that include a variety of questions related to behavioral economics, we examine how individuals' behavioral attributes are associated with responses to public signals. Allowing for a shift in mean forecast errors during Japan's high-inflation period, we find evidence broadly consistent with overreaction in inflation expectations. However, this overreaction is not fully explained by the behavioral attributes examined in this paper. Our findings suggest that overreaction is not concentrated among individuals classified as behaviorally biased.

JEL Classification: D84; D91; E31; E70

Keywords: Inflation forecasts, Over-extrapolation, Behavioral agents, Behavioral bias

*We thank M. Eguchi for helpful comments and E. Kurimune for excellent research assistance. The authors also gratefully acknowledge financial support from the Joint Usage/Research Centers at the Kyoto Institute of Economic Research, and from Grants-in-Aid for Scientific Research (JP20H05631, JP20H05633, JP21H04397, JP25K00622, and JP25H00544). The authors used a generative AI tool, ChatGPT (OpenAI, 5.4), to improve English expression. The authors reviewed and revised the output as necessary and take full responsibility for the final manuscript.

†Graduate School of Economics, Nagoya University; The Centre for Applied Macroeconomic Analysis; Correspondence to: Graduate School of Economics, Nagoya University, Furocho Chikusa-ku, Nagoya, 464-8601, Japan. E-mail: tsuruga.takayuki.a3@f.mail.nagoya-u.ac.jp

‡Graduate School of Economics, The University of Osaka.

1 Introduction

While inflation expectations are central to macroeconomic analysis, their interpretation and modeling remain challenging. Empirical studies on inflation expectations find limited support for the full-information rational expectations (FIRE) hypothesis.[1] Recent studies further show that economic agents often overreact to information that is observable at the time they form their expectations. Specifically, they document that forecast errors in inflation expectations are statistically correlated with information available at the time of expectation formation, contrary to the FIRE hypothesis. For example, Bordalo et al. (2020) find that market participants' forecast errors of macroeconomic expectations are negatively correlated with their past forecast revisions, suggesting overreaction in their expectations.[2] Similarly, Kohlhas and Walther (2021) and Broer and Kohlhas (2024) find that individuals' forecast errors in inflation expectations are negatively correlated with past inflation or past consensus forecasts. This evidence suggests that inflation expectations overreact to commonly observed macroeconomic variables, namely, public signals. To explain these findings, the literature has proposed models that depart from full rationality, such as diagnostic expectations, in which past news is overly emphasized.

In this paper, we examine how inflation expectations respond to public signals from a behavioral economics perspective. In particular, we study how this response interacts with individuals' behavioral attributes (BA). Following Kohlhas and Walther (2021), we measure the extent to which individuals' one-year-ahead inflation expectations overreact to recently realized inflation. Under the assumption that this recent realized inflation is included in individuals' information sets at the time expectations are formed, it can be interpreted as a public signal. Based on this interpretation, we then estimate how the degree of overreaction in inflation expectations varies with BA.

We use data from the *Japan Household Panel Survey on Consumer Preferences and Satisfaction* (JHPS-CPS), conducted by the University of Osaka. A notable feature of the survey is that it includes a wide range of questions related to behavioral economics, along with measures of respondents' inflation expectations. Each January since 2003, the JHPS-CPS has surveyed respondents' inflation expectations and BA, such as present bias, forward-looking planning ability, self-control, and risk attitudes. Although the survey was not conducted in

[1] See, among others, Coibion and Gorodnichenko (2015a) and Malmendier and Nagel (2016).

[2] Forecast errors of mean expectations, or consensus forecasts, are often positively correlated with revisions in mean expectations, which suggests underreaction at the aggregate level. See, for example, Coibion and Gorodnichenko (2012) and Coibion and Gorodnichenko (2015a).

some years, these rich data on BA allow us to examine the relationship between inflation expectations and BA.

Exploring the relationship between overreaction and BA has important implications for macroeconomic modeling. Previous studies have shown that inflation expectations, on average, overreact to observable information. However, heterogeneity in overreaction has not been well understood, especially from the viewpoint of behavioral economics. We explicitly examine which individuals overreact to public signals. For example, do respondents with present bias overreact more strongly to public signals? Do respondents who describe themselves as forward-looking form more precise inflation expectations? Do respondents who report frequently lacking self-control overreact more strongly to public signals? Does overreaction in inflation expectations vary with the degree of risk aversion? More fundamentally, is overreaction concentrated among behaviorally biased individuals or widely observed even among behaviorally disciplined individuals? If overreaction is primarily concentrated among behaviorally biased individuals, then models that focus exclusively on overreaction may be incomplete, and alternative behavioral assumptions may be more informative. By contrast, if overreaction is observed even among behaviorally disciplined individuals, then parsimonious models that allow for overreaction without explicitly incorporating other behavioral biases may provide a reasonable approximation.

We find that overreaction in inflation expectations is not concentrated among behaviorally biased respondents. Forecast errors in inflation expectations are negatively correlated with the public signal, which is consistent with overreaction in inflation expectations. Although BA partly account for heterogeneity in the negative response to the public signal, their effects are not large enough to make expectations consistent with the FIRE hypothesis. Even for respondents classified as behaviorally disciplined, the estimated response to public signals remains consistent with overreaction. For example, respondents classified as having no present bias exhibit a negative response of forecast errors to public signals. Although this negative response is weaker than that among respondents with present bias, it remains statistically significant, suggesting that dynamically consistent respondents' expectations still overreact to public signals under our maintained interpretation. Similarly, more risk-averse respondents also tend to respond more weakly to public signals, but the effect is not strong enough to offset the negative response.

Empirical studies have long documented deviations from the FIRE hypothesis and heterogeneity in how expectations respond to information. Coibion and Gorodnichenko (2015a), Bordalo et al. (2020), and Kohlhas and Walther (2021) study deviations from the FIRE hypothesis through the response of expectations to information or news. A related strand of the literature emphasizes heterogeneity in these responses. For example, Armantier et al. (2016) show that individual characteristics, such as financial literacy, education, and income, are important for understanding how inflation expectations respond to information. Das et al. (2020) suggest that socioeconomic status, measured by income and education, is associated with heterogeneity in macroeconomic expectations, particularly in how individuals perceive or process macroeconomic information and news.[3] Relatedly, Ilan and Mugerman (2025) show that socioeconomic status is associated with how borrowers use inflation-related information in mortgage decisions. Turning to Japan, Ueda (2010) examines the determinants of aggregate household inflation expectations. Using microdata on inflation expectations, Kikuchi and Nakazono (2023) investigate the frequency of information updating and show that expectation revisions are sensitive to changes in food-price inflation. Our paper contributes to this literature by focusing on individuals' BA in order to better understand deviations from the FIRE hypothesis. Using a unique dataset from a survey on behavioral economics, we exploit microdata on inflation expectations to study how expectations respond to public information and how this response differs across BA.

The remainder of the paper is organized as follows. Section 2 describes the survey data used in the regression analysis. Section 3 presents the regression framework, empirical results, and their implications. Section 4 concludes.

2 Data

In this section, we describe the data used in the regression analysis. The data include forecast errors of inflation expectations as the dependent variable with public signals and BA as explanatory variables.

2.1 Overview

Our analysis is mainly based on panel data from *the Japan Household Panel Survey on Consumer Preferences and Satisfaction* (JHPS-CPS), conducted by the University of Osaka. The goal of this survey is to better understand households' preferences and BA from a behavioral economics perspective. Accordingly, the JHPS-CPS includes questions related to behavioral

[3] Cavallo et al. (2017) also provide suggestive evidence that individual characteristics, such as confidence, education, gender, and age, are associated with heterogeneity in learning rates.

economics, such as those on individuals' preferences (e.g., subjective discount rates and risk attitudes), along with standard questions on basic individual characteristics. The survey also collects households' inflation expectations.

The JHPS-CPS is a nationally representative annual panel survey of residents in Japan, covering the period from 2003 to 2024. The target population consists of individuals aged 20 and over, selected through stratified random sampling.[4] The JHPS-CPS distributes paper-based questionnaires in January and collects responses in February. Over this period, the survey was not conducted in four years: 2014, 2015, 2019, and 2020. Between 2003 and 2018, the questionnaires were distributed and collected in person. From 2021 onward, they were distributed and collected by mail. To include younger generations in the sample, new respondents were added in 2004, 2006, 2009, and 2022. In the 2023 survey, the number of valid respondents was 2,921. Respondents receive cash vouchers worth 1,000 to 1,500 yen for completing the survey.

In our regression analysis, we regress forecast errors of inflation expectations for year $t+1$ on public signals measured in year t, their interactions with BA measured in year t, and other control variables measured in year t. We define the *common sample* as the sample for which forecast errors, public signals, BA, and control variables are all available for regressions. This common sample consists of 8,858 observations and forms an unbalanced panel over seven forecast years. These forecast years correspond to $t+1 = 2008$, 2009, 2011, 2017, 2022, 2023, and 2024, or equivalently $t = 2007$, 2008, 2010, 2016, 2021, 2022, and 2023. In what follows, we describe how we construct the measures of inflation expectations and BA from the survey data.

2.2 Inflation expectations and forecast errors

In constructing measures of inflation expectations from the JHPS-CPS, we follow Jinnai et al. (2021) and Mineyama and Tokuoka (2025), who study inflation expectations using this survey. The JHPS-CPS asks respondents to select a numerical range for their expected inflation rate in the survey year. Specifically, the question is: "*By what percentage do you expect consumer prices will change in* [survey year]*, compared with the previous year?*" The response options are identical across survey years and consist of 11 categories: *Increase by at least 4.5%*; *Increase by at least 3.5% but less than 4.5%*; *Increase by at least 2.5% but less than 3.5%*;

[4] The JHPS-CPS uses stratified random sampling by region and city size. Specifically, Japan is divided into ten regions, and each region is further classified into four categories according to city size.

Increase by at least 1.5% but less than 2.5%; *Increase by at least 0.5% but less than 1.5%*; *Change by less than 0.5% in either direction*; *Decrease by at least 0.5% but less than 1.5%*; *Decrease by at least 1.5% but less than 2.5%*; *Decrease by at least 2.5% but less than 3.5%*; *Decrease by at least 3.5% but less than 4.5%*; and *Decrease by at least 4.5%*. As our measure of inflation expectations, we use the midpoint of each response category, with the top and bottom open-ended categories assigned values of 5.0% and -5.0%, respectively.

To construct forecast errors, we require measures of both actual inflation and inflation expectations. Let π_{t+1} denote the year-on-year CPI inflation rate for December of year $t+1$, that is, the change in the CPI from December of year t to December of year $t+1$. The JHPS-CPS asks respondents in January of year $t+1$ by what percentage they expect consumer prices to change in year $t+1$, compared to year t, that is, the change in consumer prices from December of year t to December of year $t+1$. We denote respondent i's inflation expectation formed in January of year $t+1$ by π^e_{it+1}. Forecast errors are defined as $\pi_{t+1} - \pi^e_{it+1}$.

We use headline CPI inflation, that is, CPI for all items, as the baseline measure of "consumer prices." Of course, respondents may not necessarily interpret consumer prices in the survey as referring to headline CPI inflation. To address this issue, we also examine alternative measures of consumer prices.

We denote the public signal relevant for inflation expectations by y_t. Here y_t is assumed to be π_t. We take π_t from government publications. The Statistics Bureau of the Ministry of Internal Affairs and Communications (MIAC) releases year-on-year CPI inflation for December in January of the following year.[5] At least in terms of timing, it is reasonable to assume that this public signal can be included in respondents' information sets.

Figure 1 plots the cross-sectional mean of inflation expectations from the JHPS-CPS, shown by the solid line with squares, over the period 2007–2024. The time-series mean of JHPS-CPS inflation expectations over this period is 1.43%. When we divide the sample into 2007–2021, the low-inflation period, and 2022–2024, the high-inflation period, the mean expectation is 1.12% in the former period and 2.53% in the latter. Thus, inflation expectations increased markedly after 2021.

For comparison, the figure also plots inflation expectations from the Consumer Confidence Survey (CCS), shown by the dashed line with diamonds. The CCS is one of the main sources of data on inflation expectations in Japan. Inflation expectations from the CCS also increased markedly after 2021. A comparison of the two series shows that inflation expectations in

[5] We use vintage CPI data because the CPI series are periodically rebased.

the JHPS-CPS move broadly in line with those in the CCS over time, at least in terms of direction. In terms of magnitude, however, inflation expectations are substantially lower in the JHPS-CPS than in the CCS, especially during the high-inflation period.

Finally, Figure 1 also shows actual inflation, indicated by the solid line with circles. We observe two major increases. The first is associated with the consumption tax hike in 2014.[6] The second is the rise in inflation following the COVID-19 pandemic. The time-series mean of actual inflation over 2007–2024 is 0.86%. When we divide the sample into the low- and high-inflation periods, the time-series mean of inflation is 0.35% in the former period and 3.40% in the latter.

Comparing inflation expectations from the JHPS-CPS with actual inflation, we confirm that expected inflation is, on average, higher than actual inflation over 2007–2024, at 1.43% versus 0.86%. This pattern is consistent with D'Acunto et al. (2023), who document an upward bias in inflation expectations. However, for the JHPS-CPS, we do not observe a pronounced upward bias during the high-inflation period. The time-series mean of inflation expectations from the JHPS-CPS is 2.53%, which is lower than the corresponding mean of actual inflation of 3.40%. By contrast, inflation expectations from the CCS remain substantially above actual inflation even during the high-inflation period.

It is important to note that responses in the JHPS-CPS may be constrained by the survey's response categories. Figure 2 suggests that the elicitation of inflation expectations in 2023 may have been affected by the censoring of response categories. At the time of the survey, actual inflation was unusually high. In that year, more than 26% of respondents selected the top-coded category, *Increase by at least 4.5%*. In our coding, we assign this category an expected inflation rate of 5%. However, if the questionnaire had offered more flexible categories in the upper tail, such as *Increase by at least 5.5%* or *Increase by at least 4.5% but less than 5.5%*, some respondents would likely have selected higher numerical ranges, implying higher measured inflation expectations. By contrast, the top-coded category in the CCS is *Increase by at least 10%*. As a result, measured inflation expectations in the JHPS-CPS may be lower than those in the CCS and may even fall below actual inflation during the high-inflation period.

Figure 3 plots the cross-sectional mean of forecast errors, $\pi_t - \pi_{it}^e$, together with the 5th–95th percentile range. The mean forecast error is negative except in 2022 and 2024, implying that mean inflation expectations exceed actual inflation in all other years. The percentile

[6] Unfortunately, inflation expectations were not surveyed in that year, so we cannot examine the effect of the consumption tax hike on inflation expectations.

range is relatively stable over time, although it shifts upward during the high-inflation period.

2.3 Behavioral attributes of individuals

2.3.1 Present bias in time preferences

It is well known that preferences with present bias can lead to dynamically inconsistent choices. According to Ericson and Laibson (2019), present bias in time preferences, or quasi-hyperbolic discounting, refers to a situation in which agents prefer an action that provides immediate utility when making decisions in the present, but revise their choice when they reach a future period and reoptimize.[7] Under quasi-hyperbolic discounting, preferences are specified as

$$U_0 = u_0 + \beta[\delta u_1 + \delta^2 u_2 + \delta^3 u_3 + ...], \tag{1}$$

where δ is the standard discount factor, and β captures present bias. When $\beta < 1$, the agent is said to exhibit present bias. When comparing utility in any period $t > 0$ with utility in a future period $t + \tau$, that is, u_t and $u_{t+\tau}$, the latter is discounted by δ^τ. By contrast, when comparing u_0 and u_τ, utility in period τ is discounted by $\beta\delta^\tau$. When $\beta = 1$, the two discount factors coincide, implying exponential discounting.

From the perspective of behavioral models of expectations, the relationship between present bias and the precision of inflation expectations would be theoretically ambiguous. Present bias can affect inflation expectations through two opposing channels. On the one hand, households with present bias may attach less weight to future utility and may therefore devote less attention to forecasting inflation. In the rational-inattention model of Sims (2003), less attention makes inflation expectations less precise.[8] On the other hand, if expected inflation affects current consumption decisions through the real interest rate, households with present bias may have stronger incentives to pay attention to inflation. Greater attention may then lead to inflation expectations that are closer to rational expectations.

For the regression analysis, we construct z_{it}^{DC} as a dummy variable equal to one if respondent i is dynamically consistent, in the sense that the respondent's survey responses do not

[7] In other words, agents may later regret choices that they make in the present.

[8] See also the comprehensive survey by Maćkowiak et al. (2023) and the references therein.

indicate present bias. Specifically, z_{it}^{DC} is defined as follows:

$$z_{it}^{DC} = \begin{cases} 1 & \text{if } \beta \text{ for respondent } i \text{ in year } t = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Appendix A.1 describes how we elicit respondents' present-bias parameter β from the JHPS-CPS. Note that the JHPS-CPS asks questions on time preferences and present bias in every survey wave. Thus, z_{it}^{DC} varies over time within individuals, beyond individual fixed effects.

2.3.2 Forward-looking planning ability and self-control

We next consider two additional BA: forward-looking planning ability and self-control. We use these measures following the approach of Kubota and Fukushige (2016).[9]

We begin with forward-looking planning ability. The corresponding question is: '*How true for you is each of the following statements? Answer for each on a scale from 1 to 5, where '1' means that it is particularly true for you and '5' means that it doesn't hold true at all for you.*" The statement corresponding to forward-looking planning ability is '*I always plan things before I actually do them.*" Responses to this statement range from 1 to 5. We define respondents who choose 1 or 2 as having forward-looking planning ability, and the remaining respondents as not having such ability. For the regression analysis, we construct z_{it}^{FL} as a dummy variable equal to one if respondent i chooses 1 or 2 in year t. Because the survey asks this question in almost every survey wave, z_{it}^{FL} varies both across individuals and over time.

Using the same survey dataset, Kubota and Fukushige (2016) find that respondents with forward-looking planning ability tend to predict future income changes more precisely, whereas those lacking such ability systematically overestimate future income changes. They test the null hypothesis that expectations of future income growth are unbiased. By contrast, we test the null hypothesis that forecast errors in inflation expectations are uncorrelated with the public signal.

We also examine respondents who report having self-control. The corresponding statement in the survey is: "*When there is something I want, I need to buy it.*" The response options again consist of five categories. We define respondents who choose 4 or 5 as individuals with self-control in decision-making. For the regression analysis, we construct z_{it}^{SC} as a dummy variable equal to one if respondent i chooses 4 or 5 in year t.

[9]Unfortunately, these two questions were not included in the 2011 survey wave.

These BA may capture aspects of respondents' information-processing ability in the rational-inattention model. Respondents with weaker abilities may allocate fewer cognitive resources to forming precise inflation expectations than those with stronger abilities. Such weaker abilities could contribute to larger deviations from FIRE. These abilities may also be related to the diagnosticity parameter in diagnostic expectations (Bordalo et al., 2022), although the direction of this relationship is ultimately an empirical question.

We emphasize that the qualitative questions described above come with important caveats. First, the responses are self-reported and based on Likert-scale questions. As such, responses may vary with respondents' subjective interpretation of the scale. For example, even if respondents state that they always plan before acting, such responses do not necessarily imply true forward-looking planning ability. Second, the survey questions may not perfectly capture the BA that we aim to elicit. For instance, respondents with substantial financial assets may simply feel that they can purchase whatever they want whenever they want, so the self-control measure may partly capture financial circumstances rather than self-control itself. It should be noted that being classified as having forward-looking planning ability or self-control does not directly imply that respondents are rational in the sense of the FIRE hypothesis.

2.3.3 Risk attitude

The JHPS-CPS also measures respondents' risk attitudes as a time-varying BA by asking the following hypothetical question:

> *Suppose that there is a "speed lottery" with a 50% chance of winning ¥100,000. If you win, you get the prize right away. If you lose, you get nothing. How much would you spend to buy a ticket for this lottery? Choose Option "A" if you would buy it at that price, and choose Option "B" if you would not buy the ticket at that price.*

Respondents choose either Option A or B for a given hypothetical price of the "speed lottery." For example, the survey presents prices of 15,000 yen and 25,000 yen and asks respondents to choose A or B for each price. If a respondent switches from A to B between 15,000 yen and 25,000 yen, the respondent's reservation price is defined as the midpoint of these two amounts (i.e., 20,000 yen). By presenting several hypothetical prices, the survey attempts to measure the reservation price at which the respondent is willing to buy the speed lottery. For more details, see Appendix A.2.

To quantify risk attitudes, we follow Cramer et al. (2002) and Hanaoka et al. (2018), who use the following measure of risk aversion:

$$z_{it}^{RA} = 1 - \frac{\text{reservation price of respondent } i \text{ in year } t}{50,000}. \tag{2}$$

Note that respondents are classified as risk-neutral if their reservation price is 50,000 yen, which equals the expected return from the hypothetical speed lottery. In this case, z_{it}^{RA} is zero. By contrast, when the reservation price is strictly lower than 50,000 yen, respondents are risk-averse and z_{it}^{RA} is positive. In the extreme case in which respondents never buy the lottery (i.e., they choose Option B for all prices presented in the survey), the measure of risk aversion is 1.

The relationship between risk attitude and inflation expectations can also be interpreted through the lens of rational inattention. The precision of respondents' inflation expectations may vary with their degree of risk aversion. When the degree of *relative* risk aversion is lower, changes in consumption are more sensitive to the real interest rate. Thus, when the above measure of risk aversion is lower, the utility that respondents maximize may be more sensitive to the real interest rate through current consumption. Under this interpretation, less risk-averse respondents are expected to be more attentive to inflation and to form expectations closer to the FIRE hypothesis. An alternative interpretation is based on respondents' cognitive ability. Dohmen et al. (2010) document that individuals with higher cognitive ability are less risk averse. In the rational-inattention model, respondents with greater cognitive ability or information-processing ability may form inflation expectations that are closer to the FIRE hypothesis. Under this interpretation, if the above measure of risk aversion partly proxies information-processing ability, less risk-averse respondents may be expected to form inflation expectations that are closer to the FIRE hypothesis.

Overall, respondents in the survey are measured as risk-averse. In the common sample, only 0.96% of respondents are risk-neutral. There is substantial heterogeneity in the degree of risk aversion. The mean of z_{it}^{RA} is 0.764 with a standard deviation of 0.225.

2.3.4 Distribution of BA

Table 1 reports the distributions of z_{it}^{DC}, z_{it}^{FL}, and z_{it}^{SC}. This table allows us to examine the extent to which respondents are classified as having these BA, either separately or jointly, and helps us define behaviorally disciplined respondents. The upper panel reports the unconditional distribution of each BA, while the lower panel reports the joint distribution of the three

BA. The first column reports the distributions in the common sample. The second and third columns report the distributions separately for respondents with low and high risk aversion, respectively. Because BA are allowed to vary over time, the table reports percentages based on observations rather than on unique respondents.[10]

The upper panel shows that most respondents are classified as dynamically consistent, as shown in the first row labeled "DC." By contrast, a smaller fraction of respondents report being forward-looking, while about half report having self-control; see the second row labeled "FL" and third row labeled "SC" in the upper panel. We see from the table that the distributions of these BA are not strongly associated with the measured degree of risk aversion.

The lower panel shows the joint distribution of the three BA. We categorize respondents into eight types according to their combination of BA. In the common sample, 13.7% of observations are classified as type (1), for which $(z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC}) = (1, 1, 1)$; see the first row of the lower panel. If we include respondents classified as having at least two of the three BA, 54.2%(=13.7+10.9+25.1+4.5) of observations fall into this group, corresponding to types (1)–(4). We define these respondents as behaviorally disciplined.

Definition *Respondent i is classified as behaviorally disciplined in year t if $z_{it}^{DC} + z_{it}^{FL} + z_{it}^{SC} \geq 2$.*

2.4 Control variables

We include several control variables in our empirical specification because inflation expectations are highly heterogeneous (see Weber et al., 2022 and D'Acunto et al., 2023). Although we include individual fixed effects in all regressions, we additionally control for respondents' taxable income, age groups, and liquidity constraints.

The JHPS-CPS asks respondents about their taxable income in the previous year. The survey question is: "*Approximately how much was the annual earned income of you and your spouse (including common-law marriage), before taxes, including bonuses and business income, in [the previous year]?*" The survey provides 10 taxable income categories, and we assign a representative taxable income value to each respondent based on the selected

[10] In Appendix A.3, we report the evolution of the distribution over time; see Table A.3.

category.[11] The survey question remains unchanged across survey waves, except for the survey year itself. In the regressions, we use the natural logarithm of taxable income in year t as an explanatory variable. In the common sample, the grand mean and median of taxable income are 4.59 million yen and 5 million yen, respectively.

We also include age-group dummies as control variables because age differences are a well-known factor explaining cross-sectional variation in inflation expectations.[12] Following D'Acunto et al. (2023), we define young respondents as those under age 40 and middle-aged respondents as those aged 40 to 60. In the common sample, the young, middle-aged, and old groups account for 12.4%, 46.1%, and 41.5% of the sample, respectively.

Finally, the dummy variable for liquidity-constrained respondents is also included as a control variable. Ichiue et al. (2024) argue that liquidity-constrained households tend to form less precise inflation expectations because they are unable to smooth consumption and therefore have weaker incentives to hold precise inflation expectations for consumption-smoothing purposes. The JHPS-CPS includes a related question each year. The survey question is: *"Have you ever been rejected for a loan application (excluding housing loans)?"* For survey years before 2011, the question offers only two options: Yes or No. Since 2011, the question has offered five options, and we define respondents as liquidity-constrained if they choose option 1, 2, or 3, and as not liquidity-constrained if they choose option 4 or 5.[13] In the common sample, 91.9% of respondents report not being liquidity-constrained.

[11] The options are as follows: "*1. none*; *2. Less than ¥1,000,000*; *3. ¥1,000,000 to less than ¥2,000,000*; *4. ¥2,000,000 to less than ¥4,000,000*; *5. ¥4,000,000 to less than ¥6,000,000*; *6. ¥6,000,000 to less than ¥8,000,000*; *7. ¥8,000,000 to less than ¥10,000,000*; *8. ¥10,000,000 to less than ¥12,000,000*; *9. ¥12,000,000 to less than ¥14,000,000*; *10. ¥14,000,000 or more*. We assign the following income values to the 10 categories: 1. ¥0; 2. ¥500,000; 3. ¥1,500,000; 4. ¥3,000,000; 5. ¥5,000,000; 6. ¥7,000,000; 7. ¥9,000,000; 8. ¥11,000,000; 9. ¥13,000,000; 10. ¥15,000,000." In the common sample, the first five categories account for 71.6% of respondents.

[12] D'Acunto et al. (2023) report that younger respondents tend to have lower mean inflation expectations than older respondents.

[13] The options are as follows: "*1. Yes; 2. No, but I was not approved for the full amount for which I applied and was instead offered a reduced amount; 3. I did not apply because I did not think I would be approved; 4. No, I have always been able to borrow the amount I applied for; 5. I have never attempted to borrow money.*"

3 Regression results

3.1 The estimation equation

Throughout the paper, we estimate regressions of forecast errors on the public signal and its interactions with BA:

$$\pi_{t+1} - \pi^e_{it+1} = a_i + by_t + b_z' z_{it} y_t + c' X_{it} + v_{it+1}. \tag{3}$$

Here, y_t denotes the public signal. In many specifications, we set $y_t = \pi_t$. That is, we interpret lagged inflation as a public signal under the assumption that it is included in respondents' information sets at the time expectations are formed and is relevant to their inflation expectations. The vector z_{it} denotes BA. Typically, $z_{it} = (z^{DC}_{it}, z^{FL}_{it}, z^{SC}_{it}, \tilde{z}^{RA}_{it})'$, where $\tilde{z}^{RA}_{it}$ is the standardized version of z^{RA}_{it}. The vector of control variables is denoted by X_{it} and includes the first-order terms of z_{it}. The error term is denoted by v_{it+1}. The term a_i captures individual fixed effects, b and b_z are parameters governing the response of forecast errors to the public signal, and c is the parameter vector associated with the control variables.

This regression extends the specification used by Kohlhas and Walther (2021). Our goal is to examine how the response of forecast errors to public signals varies with BA. This heterogeneity is captured by the coefficients on the interaction terms in Equation (3). In particular, the parameter vector b_z measures how BA are associated with differential responses of forecast errors to the public signal.

Although the dependent variable is the forecast error of inflation expectations for year $t+1$, we use explanatory variables measured in year t, namely z_{it} and X_{it}, rather than variables measured in year $t+1$. This timing ensures that BA and control variables are measured before the realization of the forecast error. However, this timing choice does not make the measured BA exogenous. BA measured in year t may reflect respondents' past macroeconomic experiences or other factors that are also related to their expectation formation. We therefore interpret the interaction coefficients as describing conditional heterogeneity in response to public signals, rather than as causal effects of BA.

Under the FIRE hypothesis, forecast errors are uncorrelated with y_t since y_t is assumed to be included in respondents' information sets at the time they form their expectations. In other words, the forecast error $\pi_{t+1} - \pi^e_{it+1}$ is orthogonal to the public signal y_t.[14] Testing

[14] Let Ω_{it} be respondent i's information set at the end of year t. If $\pi_t \in \Omega_{it}$ and $\pi^e_{it+1} = \mathbb{E}(\pi_{t+1} \mid \Omega_{it})$, then $\mathrm{Cov}\,(\pi_{t+1} - \mathbb{E}(\pi_{t+1} \mid \Omega_{it}), \pi_t) = 0$ by the law of iterated expectations.

$b = 0$ in a specification without interactions corresponds to testing the joint hypothesis that respondents' inflation expectations are rational and that their information sets include the public signal, i.e., $y_t \in \Omega_{it}$.[15]

More precisely, our estimation equation further allows the response to the public signal to depend on respondents' BA, in order to test whether deviations from the FIRE hypothesis vary across respondent types. Let the elements of the estimated coefficient vector on the interaction terms be denoted by $\hat{b}_z^{DC}$, $\hat{b}_z^{FL}$, $\hat{b}_z^{SC}$, and $\hat{b}_z^{RA}$, corresponding to z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC}, and $\tilde{z}_{it}^{RA}$, respectively. For example, for behaviorally disciplined respondents with $(z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC}) = (1, 1, 1)$, we test the null hypothesis $b + b_z^{DC} + b_z^{FL} + b_z^{SC} = 0$. In general, the relevant null hypothesis depends on the specification and on the combination of BA used to define each respondent type.

The literature has often found negative estimates of b in specifications without interaction terms with BA. Suppose that π_{t+1} and y_t are positively correlated. Then, a negative response of forecast errors to the public signal, i.e., $b < 0$, implies that π_{it+1}^e responds more strongly to y_t than π_{t+1} does. Under this interpretation, inflation expectations overreact to the public signal. By contrast, a positive response of forecast errors, i.e., $b > 0$, implies underreaction to the public signal. More rigorously, in order to make the argument for overreaction or underreaction in the regression framework, the public signal in regressions must correspond to the information that respondents actually use when forming inflation expectations.

In the regressions, we standardize the measure of risk aversion, z_{it}^{RA}.[16] With this adjustment, and holding other factors constant, b can be interpreted as the response of forecast errors when respondents' degree of risk aversion is at its grand mean. The coefficient on the interaction term between y_t and $\tilde{z}_{it}^{RA}$ captures the additional change in the response of forecast errors to the public signal associated with a one-standard-deviation increase in the risk-aversion measure relative to its grand mean.

One advantage of this regression is that the test is more practical in our setting than the tests in Coibion and Gorodnichenko (2015a) and Bordalo et al. (2020), in which forecast errors are regressed on forecast revisions defined as changes in forecasts made at different points in time. The JHPS-CPS collects only one-year-ahead forecasts and does not elicit two forecasts for the same target year in consecutive waves. Therefore, we cannot construct forecast revisions for the same forecast target. Our specification instead uses the response of forecast errors to public signals.

[15] For further details, see Broer and Kohlhas (2024).

[16] We also standardize log taxable income, which is included in X_{it}.

3.2 Overreaction to the public signal

3.2.1 Regression results without heterogeneity in BA

Table 2 reports the regression results without heterogeneity in BA. All specifications are estimated by ordinary least squares (OLS) and include individual fixed effects. In all columns, we use the common sample, so that the results are directly comparable to the regression results with heterogeneity in BA reported in Tables 4–6.

Specification (1) regresses forecast errors on the public signal, individual fixed effects, and control variables. The coefficient on the public signal b is estimated to be negative but statistically insignificant, based on standard errors that are two-way clustered by respondent and year.

The estimated coefficient should be interpreted with caution for two reasons. First, the high-inflation period may raise measurement issues for inflation expectations. In particular, this period may have changed the process of expectation formation, generating a period-specific bias in inflation expectations. In addition, top-coded response categories may directly affect the measurement of inflation expectations during this period. Second, the time-series variation used to identify the main coefficient b is limited because the public signal varies only at the year level. In the common sample, the data cover only seven survey years: $t + 1 = 2008, 2009, 2011, 2017, 2022, 2023$, and 2024. This limited time-series variation makes inference less reliable. Following Cameron and Miller (2015), we conservatively base inference on critical values from a t-distribution with $T - 1$ degrees of freedom, where T is the number of years used in each regression.[17]

To allow for a possible change in inflation expectations during the high-inflation period from 2022 to 2024, we introduce a dummy variable equal to one for observations from this period. Specification (2) allows for a shift in the mean forecast error during the high-inflation period, whereas specification (3) additionally allows the slope with respect to the public signal to differ during this period.

When we allow for a structural break in the mean forecast errors during the high-inflation period, the estimated response of forecast errors to the public signal tends to be consistent with overreaction in inflation expectations. The coefficient on the public signal becomes strongly negative: $\hat{b} = -0.806$ in specification (2) while $\hat{b} = -0.077$, slightly negative, in

[17]The estimated standard errors may still be underestimated when the number of clusters is small, even when inference is based on the t-distribution. The literature has not established an ideal approach for multi-way clustering with few clusters (Cameron and Miller, 2015, p. 350).

specification (3). The coefficient is statistically significant at the 1% significance level in specification (2), while it is imprecisely estimated in specification (3). In specification (2), the estimated response is larger in absolute value than those reported for other countries in previous studies. For example, Kohlhas and Walther (2021) report that $\hat{b}$ is -0.18 in the U.S. Survey of Professional Forecasters, -0.01 in the euro-area Survey of Professional Forecasters, -0.16 in the Livingston Survey, and -0.10 in the Michigan Survey of Consumers. In specification (3), the high-inflation-period dummy implies a statistically significant difference in the intercept, whereas the interaction between the dummy and the public signal does not imply a statistically significant difference in the slope. This makes the interpretation of the main coefficient difficult.

An alternative way to examine the impact of the high-inflation period on the estimation results is to divide the sample into the pre-2022 and 2022–2024 periods. Specification (4) of Table 2 uses the subsample from the pre-2022 period, whereas specification (5) uses the subsample from 2022 to 2024. The coefficient on the public signal is estimated to be positive in the pre-2022 period and negative in the high-inflation period. The estimate for the pre-2022 period is unstable and imprecisely estimated.

Recall also that the common sample contains only seven years. Dividing this common sample into the pre-2022 and 2022–2024 periods leaves too few survey years in each subsample, making the estimates unstable. Thus, we temporarily extend the sample to include all observations for which forecast errors and the public signal are available. This gives 13 usable survey years in total: ten years before 2022 and three years from 2022 to 2024.[18]

Table 3 repeats the estimation in Table 2 using the extended sample. The table shows that, although the magnitude of the coefficients varies across sample periods and specifications, the coefficient on the public signal b remains negative once we allow for the high-inflation-period dummy or estimate the model separately by subperiod. When we allow for more time-series variation in the public signal during the pre-2022 period, the coefficient is estimated more precisely. In particular, specification (4) shows that $\hat{b} = -0.640$ and is statistically significant at the 5% significance level. Appendix A.4 also conducts the subsample analysis after excluding observations in which inflation expectations are classified in the top-coded category. See Table A.5 in the appendix. The estimated coefficients become somewhat smaller in absolute value, but they remain negative and statistically significant in both subsamples.

To summarize, our analysis is broadly consistent with the literature's finding that inflation

[18]More specifically, the usable years for the pre-2022 period are 2008–2013, 2016–2018, and 2021.

expectations overreact to public signals. However, our results should be interpreted with caution. Our results suggest that the recent high-inflation episode in Japan has a non-negligible impact on the mean forecast error and the estimated coefficient on the public signal differs substantially depending on whether the high-inflation-period dummy is included. The subsample analysis dividing the common sample suggests that the coefficient on the public signal is less precisely estimated when we use the sample before the high-inflation period, partly because time-series variation in the public signal is limited. When we use the extended sample with more time-series variation, we find evidence consistent with overreaction even before the high-inflation period, with improved precision. Taken together, the evidence is not inconsistent with overreaction in inflation expectations, but the magnitude of the estimated response is influenced by observations from the high-inflation period.

In what follows, we report regression results that include the high-inflation-period dummy, allowing only for a shift in the mean forecast error. We do not include the interaction between this dummy and the public signal because the coefficient on the interaction term is not statistically significant.

3.2.2 Regression results with heterogeneity in BA

We next discuss the regression results with heterogeneity in BA, based on the interaction terms in the estimation equation.[19]

Dynamically consistent respondents Specification (1) in Table 4 regresses forecast errors on y_t and its interaction with z_{it}^{DC}. The estimates are $\hat{b} = -0.922$ and $\hat{b}_z^{DC} = 0.161$, and both coefficients are statistically significant at the 1% significance level.[20]

The positive coefficient on the interaction term indicates that, when $z_{it}^{DC} = 1$, the negative response of forecast errors to the public signal becomes smaller in absolute value. Under our maintained interpretation, this weaker response suggests that respondents without present bias overreact less to the public signal than those with present bias. More importantly, however, the magnitude of this effect is not large enough to make their expectations consistent with the FIRE hypothesis. Even for dynamically consistent respondents ($z_{it}^{DC} = 1$), the

[19] Table A.6 of Appendix A.4 reports the regression results after excluding observations in which inflation expectations are classified in the top-coded category.

[20] We also estimate the specification using the elicited present-bias parameter β directly. The results are qualitatively similar: $\hat{b}$ is -1.043, and the coefficient on the interaction term is 0.284. Both coefficients are statistically significant.

estimated response of forecast errors to the public signal is -0.761 ($= -0.922 + 0.161$), indicating a statistically significant negative response to the public signal.

Respondents with forward-looking planning ability or self-control Turning to forward-looking planning ability, specification (2) indicates that $\hat{b} = -0.783$ and $\hat{b}_z^{FL} = -0.066$, and both coefficients are statistically significant at the 5% significance level. For these respondents, the implied coefficient is -0.849 ($= -0.783 - 0.066$), which is again significantly negative.[21] Thus, regardless of whether respondents are classified as having forward-looking planning ability, forecast errors respond negatively to the public signal. In fact, the negative coefficient on the interaction term implies that respondents with forward-looking planning ability exhibit a stronger negative response than those without such ability.

This finding contrasts with the rationality test in Kubota and Fukushige (2016), who fail to reject rational expectations for income changes among respondents with forward-looking planning ability. In our setting, by contrast, forecast errors respond negatively to the public signal even among respondents who report having forward-looking planning ability.

For self-control, we do not find a systematic pattern in the response to the public signal. In specification (3), $\hat{b} = -0.813$, which is statistically different from zero at the 1% significance level. By contrast, the estimated coefficient on the interaction term is small, at 0.015. Later, we confirm that the coefficient on the interaction term with z_{it}^{SC} is statistically insignificant in most specifications.[22] This result suggests that self-control may not matter for the response to the public signal. However, it may simply reflect that this self-reported BA is too noisy a proxy to identify its relationship with the response to the public signal.[23]

Heterogeneity in risk attitude Specification (4) examines how the negative response to the public signal varies with the degree of risk aversion. The estimates of $\hat{b}$ and $\hat{b}_z^{RA}$ are -0.788 and 0.067, respectively. The positive coefficient on the interaction term, $\hat{b}_z^{RA}$, suggests that more risk-averse respondents exhibit a weaker negative response to the public signal. However, the effect of risk aversion on the response is small relative to the magnitude of the main coefficient on y_t. For example, even when a respondent's risk-aversion measure is

[21] We also confirm that the result is robust to an alternative definition of z_{it}^{FL}, in which $z_{it}^{FL} = 1$ only when respondent i chooses option 1.

[22] We confirm that the results are quantitatively similar under an alternative definition of z_{it}^{SC}, in which $z_{it}^{SC} = 1$ only when respondent i chooses option 5.

[23] This concern may apply to the case of forward-looking planning ability because both BA are self-reported and based on the Likert-scale questions.

two standard deviations above the mean, the implied response of forecast errors to the public signal is $-0.788 + 2 \times 0.067 = -0.654$.[24]

Behaviorally disciplined respondents Specification (5) reports the results from a regression that includes all of the above BA. This specification is defined as the baseline specification. The signs and magnitudes of $\hat{b}$ and $\hat{b}_z$ in this specification are broadly similar to those in specifications (1)–(4). Based on the estimates in this column, we examine how the response to the public signal differs across types of behaviorally disciplined respondents.

Table 5 summarizes the responses of forecast errors to the public signal for different types of behaviorally disciplined respondents. In the first row, we compute the response of forecast errors for respondents with $z_{it}^{DC} = z_{it}^{FL} = z_{it}^{SC} = 1$ as the sum of the estimated coefficients, $\hat{b}+\hat{b}_z^{DC}+\hat{b}_z^{FL}+\hat{b}_z^{SC}$.[25] This sum is -0.778 ($= -0.887+0.157-0.064+0.016$) and is significantly different from zero at the 1% significance level. As shown in the second to fourth rows, the estimated responses for the other types of behaviorally disciplined respondents are similarly computed and found to be quantitatively similar to that for the first type. The second column of the table shows that the implied response of inflation expectations ranges from 1.221 to 1.442, which exceeds the response of the estimated response of actual inflation to the public signal, 0.507.[26]

We thus argue that forecast errors respond negatively to the public signal, regardless of whether respondents are behaviorally disciplined. Under our maintained interpretation, this result suggests that overreaction in inflation expectations is not concentrated only among respondents who appear more behaviorally biased. This result remains unchanged even when we allow for a large positive deviation in respondents' risk aversion from the mean.[27]

[24] Cramer et al. (2002) and Hanaoka et al. (2018) use the Arrow-Pratt measure of absolute risk aversion. We also construct the Arrow-Pratt measure from the elicited reservation prices as an alternative measure of risk aversion. The results are robust to this alternative definition.

[25] More precisely, this sum gives the response of forecast errors to the public signal for a behaviorally disciplined respondent whose degree of risk aversion is equal to the sample mean, because $\tilde{z}_{it}^{RA}$ has zero mean.

[26] To see this, recall that forecast errors are defined as $\pi_{t+1} - \pi_{it+1}^e$. Thus, the response of π_{it+1}^e is given by the response of π_{t+1} minus the response of forecast errors $\pi_{t+1} - \pi_{it+1}^e$. For example, for respondents with $z_{it}^{DC} = z_{it}^{FL} = z_{it}^{SC} = 1$, the implied response of inflation expectations is $0.507 - (-0.778) = 1.285$. Here the response of the actual inflation is estimated from the regression of π_{t+1} on $y_t (= \pi_t)$ using data for $t + 1 = 2008, \ldots, 2024$ to secure the time-series observation. If the common sample is used, the estimated response is 0.587, close to the value of 0.507.

[27] Even when the risk-aversion measure is two standard deviations above the mean, the implied response of forecast errors to the public signal is -0.650 ($= -0.778 + 2 \times 0.064$).

3.2.3 Robustness

Table 6 reports a variety of robustness checks.

Structural break In the previous section, we argued that overreaction in inflation expectations is not concentrated among behaviorally biased respondents. As noted above, however, the estimated response of forecast errors to the public signal varies depending on whether the high-inflation-period dummy is included. Our main result on heterogeneity in the response may therefore reflect how we control for the structural break associated with the high-inflation period. To address this concern, we replace the main effect of the public signal, $b \times y_t$, in Equation (3) with year fixed effects. This replacement prevents us from identifying the main coefficient on the public signal. Nevertheless, it allows us to avoid relying on the high-inflation-period dummy and to identify heterogeneity in the response of forecast errors to the public signal from the interaction terms.

The first two columns of Table 6 report the results with and without year fixed effects. For comparison, specification (1) reproduces specification (5) in Table 4 as the baseline specification. Specification (2) shows that the coefficients on the interaction terms have the same signs as those in the baseline specification, and that their magnitudes remain quantitatively small. This result supports our finding that the negative response of forecast errors to the public signal is not concentrated among behaviorally biased respondents.

Alternative measures of consumer prices for constructing forecast errors In our benchmark regression, forecast errors are calculated using headline CPI inflation. However, respondents may not necessarily forecast headline CPI inflation when answering the JHPS-CPS, because the survey simply asks how much they expect *consumer prices* to change. To address this issue, we examine the robustness of our results using alternative measures of consumer prices for constructing forecast errors. In this robustness analysis, the public signal is also replaced by the corresponding measure of inflation. That is, as before, the alternative variable with one-year lag is assumed to be included in respondents' information sets and is relevant to their inflation expectation.[28]

Specification (3) in Table 6 reports the results when we use core CPI inflation, namely CPI inflation excluding fresh food.[29] Our main finding remains unchanged: the evidence

[28]Regarding relevance, we confirmed that the forecast target variable and the corresponding public signal are positively correlated.

[29]We use vintage CPI data for core CPI inflation.

consistent with overreaction in inflation expectations is not concentrated among behaviorally biased respondents. In particular, the main coefficient on the public signal is -0.895, and the heterogeneity in the response is small relative to the magnitude of the main coefficient. The implied responses of forecast errors for all four types of behaviorally disciplined respondents are negative and statistically different from zero; see the first column of Table A.7 for the detailed results.

Specification (4) reconfirms the negative response of forecast errors to the public signal when we use fresh-food inflation.[30] The use of fresh-food inflation is also motivated by empirical evidence that consumers often refer to the prices of frequently purchased goods when forming their inflation expectations.[31] Although the response remains negative, the main coefficient is substantially smaller in absolute value than in the baseline specification ($\hat{b} = -0.173$). Nevertheless, the implied responses of forecast errors for all four types of behaviorally disciplined respondents are negative and statistically different from zero at least at the 10% significance level; see the second column of Table A.7 for the detailed results.

Specification (5) reports the results using city-level headline CPI inflation. In this specification, we use headline CPI inflation for the city corresponding to the prefectural capital of respondents' prefectures as a proxy for consumer prices observed by respondents. Here the results are broadly similar to those in the baseline specification. For the response of forecast errors for each respondent type, see the third column of Table A.7. Specification (6) further includes year fixed effects. Because the city-level public signal varies across both years and cities, its main coefficient remains identified even after controlling for year fixed effects. The estimated coefficient on the public signal becomes smaller in magnitude than in specification (5) and is statistically insignificant. However, it remains negative, as in the baseline specification.[32] Moreover, the estimated coefficients on the interaction terms have the same signs as those in the baseline specification.[33]

[30] Unlike the headline and core CPI inflation, we use historical price index data for fresh-food inflation from *the Annual Report on the Consumer Price Index.*

[31] See, for example, Coibion and Gorodnichenko (2015b). Kikuchi and Nakazono (2023) also argue that inflation expectations in Japan are sensitive to changes in the prices of frequently purchased goods.

[32] In this specification, year fixed effects absorb aggregate shocks common to all respondents, so the main coefficient on the city-level public signal is identified from cross-city variation within each year. While we continue to report standard errors two-way clustered by respondent and year, it may also be reasonable to use standard errors clustered by prefecture. With this alternative clustering, we find that the main coefficient is statistically significant at conventional significance levels.

[33] Although the use of city-level inflation in specifications (5) and (6) provides an informative robustness check, it also has an important limitation as the public signal. Namely, respondents in most prefectures except for Tokyo are likely to be less exposed to local-level inflation figures than to national-level inflation figures through government publications and media coverage.

The results from the robustness checks again suggest that the evidence is not inconsistent with overreaction in inflation expectations, although the magnitude of the estimated response varies across specifications.

3.2.4 Alternative measures of the public signal

So far, we have tested the FIRE hypothesis using headline CPI inflation as the public signal. However, alternative measures may also serve as public signals if they are relevant for respondents' expectation formation. Moreover, the literature suggests that individuals often rely on information other than aggregate inflation statistics when forming inflation expectations. Table 7 reports the estimation results when we use alternative measures of the public signal. Specification (1) reproduces the baseline specification reported in Table 4.

In specification (2) of Table 7, fresh-food inflation is used as the public signal. Cavallo et al. (2017) find that individuals use information on supermarket prices rather than relying exclusively on official inflation statistics when forming expectations about general price inflation. If fresh-food inflation serves as a proxy for information on supermarket prices and is included in respondents' information sets, it should be orthogonal to forecast errors under the FIRE hypothesis. We therefore regress forecast errors constructed from headline CPI inflation on lagged fresh-food inflation. Note that, in contrast to the specification in the previous section, the public signal is now fresh-food inflation, while forecast errors are still based on headline CPI inflation.

The response of forecast errors remains negative, but its magnitude is smaller than in the baseline specification reported in specification (1). In the common sample, none of the coefficients on the interaction terms with fresh-food inflation is statistically significant at conventional significance levels. However, the response of forecast errors to this alternative public signal remains negative. Thus, while the magnitude of the main coefficient differs from that in the baseline specification, its sign remains the same.

Specification (3) uses the mean professional forecast of inflation as y_t. Armantier et al. (2016) show that the mean professional forecast of inflation affects inflation expectations more strongly than past-year food-price inflation.[34] If professional forecasts are publicly available through news coverage, policy discussions, and other media reports and are included in respondents' information sets, we can interpret the mean professional forecast as a proxy

[34]More specifically, the mean forecast refers to the mean forecast of next-year overall inflation in the Survey of Professional Forecasters in the United States, while food-price inflation refers to past-year average food-price inflation.

for the public signal based on professional assessments.[35]

The results are ambiguous. In terms of the point estimate, forecast errors exhibit a stronger response to the mean professional forecast than to fresh-food inflation. However, the main coefficient is too imprecisely estimated to assess how different it is from the baseline estimate in specification (1). Likewise, the responses of forecast errors for behaviorally disciplined respondents are estimated to be negative, but none of them is statistically significant.

Specification (4) uses the month-on-month change in headline CPI from November to December of year t, expressed at an annualized rate, as the public signal, rather than the year-on-year price change from December of year $t-1$ to December of year t. Month-on-month inflation is a noisier measure than year-on-year inflation because it may reflect temporary shocks and seasonal movements. In a Bayesian-learning framework, inflation expectations may respond only weakly to such a noisy signal. On the other hand, month-on-month inflation may be salient to respondents because it captures the most recent change in consumer prices, which may make respondents more attentive to this signal.

As shown in specification (4) of Table 7, the main coefficient is estimated to be slightly negative but statistically insignificant. The heterogeneity in the response of forecast errors to this public signal is similar to that in the baseline specification. However, because the main coefficient is imprecisely estimated, the responses of forecast errors for behaviorally disciplined respondents are statistically insignificant, even though their point estimates are negative.

Finally, experienced inflation is used as an alternative measure of the public signal. Malmendier and Nagel (2016) and Diamond et al. (2020) emphasize the importance of households' inflation experiences in shaping their inflation expectations. Following Mineyama and Tokuoka (2025), we construct a measure of experienced inflation and examine the response of forecast errors to this signal. Experienced inflation for respondent i in year t is given by

$$\pi_{it}^{EI} = \sum_{k=0}^{\mathrm{age}_{it}-1} w_{it}(k,\rho)\pi_{t-k}, \tag{4}$$

where $w_{it}(k,\rho)$ denotes the weight assigned to commonly observable past inflation:

$$w_{it}(k,\rho) = \frac{(\mathrm{age}_{it}-k)^{\rho}}{\sum_{k=0}^{\mathrm{age}_{it}-1}(\mathrm{age}_{it}-k)^{\rho}},$$

[35] The mean professional forecast is constructed from professional forecasts for Japan reported in *Consensus Forecasts: G7 and Western Europe*. The forecasts are surveyed monthly by Consensus Economics. For expected inflation in year $t+1$, we use the forecast reported in January of that year.

and ρ determines the shape of $w_{it}(k, \rho)$. If $\rho = 0$, respondent i assigns equal weight to each past year. If ρ is positive (negative), the respondent assigns greater (less) weight to more recent inflation. Mineyama and Tokuoka (2025) find that $\rho = 11.03$ best explains inflation expectations in the JHPS-CPS. We therefore construct π_{it}^{EI} by calibrating ρ to this value in our regressions. Experienced inflation can be regarded as a quasi-public signal because it is based on publicly observable past inflation. It is not a purely common signal, however, because the weights assigned to past inflation differ across respondents.

Specification (5) in Table 7 reports the results. All the estimated coefficients shown in this column have the same signs as those in specification (1). The estimate of b is -0.523, which is closer to the estimate from the baseline specification than the estimates obtained using fresh-food inflation, the mean professional forecasts, or month-on-month inflation. However, again, this coefficient is not statistically significant.

Taken together, the results in Table 7 show that the point estimates of the main coefficients are generally negative and thus not inconsistent with overreaction in inflation expectations. However, the estimates vary substantially across specifications and are not precisely estimated. These imprecise results may reflect the fact that respondents in our survey are not directly provided with information on these public signals, unlike in the information-provision experiments of Armantier et al. (2016) and Cavallo et al. (2017). The limited time-series variation in the data also makes it difficult to draw strong conclusions. Regarding heterogeneity in responses to alternative public signals, the estimated coefficients on the interaction terms with BA are generally small relative to the main coefficient on the public signal. Thus, BA do not appear to fully explain possible overreaction to these alternative public signals.

4 Concluding remarks

Expectations are central to macroeconomics, but how they should be modeled remains an open question. A large empirical literature has shown that the FIRE hypothesis is often inconsistent with observed behavior: individuals' forecast errors are negatively correlated with information available at the time expectations are formed, suggesting overreaction to such information. To account for this pattern, recent behavioral models emphasize that economic agents may place excessive weight on observable signals.

Using unique Japanese survey data with information on BA, this paper examines how overreaction in inflation expectations to public signals is related to individuals' BA. Allowing for a period-specific shift in mean forecast errors during the high-inflation period, we find evidence broadly consistent with overreaction in inflation expectations. However, the overreaction is not fully explained by the BA examined in this paper. Even respondents classified as behaviorally disciplined continue to exhibit a negative response of forecast errors to public signals. The evidence therefore suggests that overreaction is not concentrated among behaviorally biased respondents. More broadly, our findings support parsimonious macroeconomic models that allow for overreaction while otherwise retaining disciplined behavioral assumptions as a reasonable approximation.

A limitation of this paper is that it does not fully identify the mechanism behind overreaction. Our evidence shows that overreaction is not fully explained by the observable BA considered in this paper. However, other behavioral mechanisms may still account for heterogeneity in overreaction. Identifying the underlying mechanism is an important direction for future research.

References

Akesaka, M. (2019). Change in time preferences: Evidence from the Great East Japan Earthquake. *Journal of Economic Behavior & Organization*, 166:239–245.

Armantier, O., Nelson, S., Topa, G., Van der Klaauw, W., and Zafar, B. (2016). The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment. *The Review of Economics and Statistics*, 98(3):503–523.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020). Overreaction in macroeconomic expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Overreaction and diagnostic expectations in macroeconomics. *Journal of Economic Perspectives*, 36(3):223–244.

Broer, T. and Kohlhas, A. N. (2024). Forecaster (mis-)behavior. *The Review of Economics and Statistics*, 106(5):1334–1351.

Cameron, A. C. and Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159.

Coibion, O. and Gorodnichenko, Y. (2015a). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–2678.

Coibion, O. and Gorodnichenko, Y. (2015b). Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.

Cramer, J. S., Hartog, J., Jonker, N., and Van Praag, C. M. (2002). Low risk aversion encourages the choice for entrepreneurship: An empirical test of a truism. *Journal of Economic Behavior & Organization*, 48(1):29–36.

D'Acunto, F., Malmendier, U., and Weber, M. (2023). What do the data tell us about inflation expectations? In *Handbook of Economic Expectations*, pages 133–161. Elsevier.

Das, S., Kuhnen, C. M., and Nagel, S. (2020). Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies*, 33(1):395–432.

Diamond, J., Watanabe, K., and Watanabe, T. (2020). The formation of consumer inflation expectations: New evidence from Japan's deflation experience. *International Economic Review*, 61(1):241–281.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3):1238–1260.

Ericson, K. M. and Laibson, D. (2019). Intertemporal choice. In *Handbook of Behavioral Economics: Applications and Foundations 1*, volume 2, pages 1–67. Elsevier.

Hanaoka, C., Shigeoka, H., and Watanabe, Y. (2018). Do risk preferences change? Evidence from the Great East Japan Earthquake. *American Economic Journal: Applied Economics*, 10(2):298–330.

Ichiue, H., Katagiri, M., Koga, M., Okuda, T., and Ozaki, T. (2024). Households' attention to the central bank, inflation expectations, and spending. Unpublished manuscript.

Ilan, M. and Mugerman, Y. (2025). Misguided mortgage choices: Financial literacy, inflation expectations, and borrowing decisions. *Journal of Behavioral and Experimental Finance*, 47:101077.

Jinnai, R., Mikami, T., Okuda, T., and Nakajima, J. (2021). Household inflation expectation and consumption. *Keizai Kenkyu (The Economic Review)*, 72(3):268–295.

Kikuchi, J. and Nakazono, Y. (2023). The formation of inflation expectations: Microdata evidence from Japan. *Journal of Money, Credit and Banking*, 55(6):1609–1632.

Kohlhas, A. N. and Walther, A. (2021). Asymmetric attention. *American Economic Review*, 111(9):2879–2925.

Kubota, K. and Fukushige, M. (2016). Rational consumers. *International Economic Review*, 57(1):231–254.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273.

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, 131(1):53–87.

Mineyama, T. and Tokuoka, K. (2025). Investigating how inflation expectations affect precautionary wealth. *Japan and the World Economy*, 73:101295.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Ueda, K. (2010). Determinants of households' inflation expectations in Japan and the United States. *Journal of the Japanese and International Economies*, 24(4):503–518.

Weber, M., D'Acunto, F., Gorodnichenko, Y., and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *Journal of Economic Perspectives*, 36(3):157–184.

A Appendix

A.1 Eliciting present bias

To elicit respondents' present-bias parameter β, we follow Akesaka (2019). The JHPS-CPS asks respondents questions about $\beta\delta^\tau$ and δ^τ. Although the wording of the survey questions differs slightly across years, the typical survey questions are as follows:

(1) *Suppose that you are to receive money from someone. You can either choose to receive the money today, or 7 days from today, but the amounts will be different. Compare the amounts and dates below in Option "A" and Option "B", and indicate which option you prefer for each of the nine choices.*

(2) *Now, suppose that you are to receive money from someone. You can choose either to receive the money 90 days from today, or 97 days from today, but the amounts will be different. Compare the amounts and dates below in Option "A" and Option "B", and indicate which option you prefer for each of the nine choices.*

The discount factor $\beta\delta^\tau$ is elicited from question (1), whereas δ^τ is elicited from question (2), where τ corresponds to 7 days. The discount factor is identified by the choice at which the respondent switches from option A to option B. For example, the list of choices for question (1) in the 2024 survey is shown in Table A.1. Focusing on the first and second choices, if option A is strictly preferred to option B in the first choice, then $3,005 > \beta\delta^\tau \times 3,014$ holds, so that $\beta\delta^\tau < 3,005/3,014 = 0.997$. If option B is strictly preferred to option A in the second choice, then $3,003 < \beta\delta^\tau \times 3,297$ holds, so that $\beta\delta^\tau > 3,003/3,297 = 0.911$. Thus, these two choices determine the range of the discount factor, $0.911 < \beta\delta^\tau < 0.997$. Repeating this procedure for the remaining choices allows us to narrow the range of $\beta\delta^\tau$. To pin down a value of $\beta\delta^\tau$, we take the midpoint of the range implied by the nine choices. Similarly, for question (2), we derive the range of δ^τ from the nine choices and use its midpoint as the measure of the discount factor obtained from question (2). In an extreme case, the implied discount factor may exceed unity, which implies a negative discount rate.[36] In this case, we recode any discount factor greater than unity as unity. In another extreme case, inconsistent choices across the nine options may prevent us from uniquely identifying a range for the discount factor. In such cases, we drop those respondents from the sample.

[36] In the example in Table A.1, if a respondent chooses option A only in choice 8, the discount factor exceeds unity.

The present-bias parameter β is calculated as the ratio of the two estimated discount factors, $\beta\delta^{\tau}$ and δ^{τ}. Although the JHPS-CPS is designed to rule out the possibility of $\beta < 0$, the resulting ratio may exceed 1 and thereby violate the assumption that $\beta \leq 1$. Therefore, when β exceeds one, we reset it to unity so that all respondents in the data satisfy $0 < \beta \leq 1$.[37] Under the common sample, 76.3% of respondents exhibit no present bias (i.e., $\beta = 1$), while the remaining respondents exhibit present bias (i.e., $\beta < 1$). Present bias is widely distributed. The minimum value of β is 0.02, which is far from the no-present-bias benchmark of $\beta = 1$.

We make two remarks on changes in the survey questions. First, the dates in question (1) in the original survey before 2010 differed slightly from those in the format described above. In option A, respondents could hypothetically receive the money in 2 days rather than today. Likewise, in option B, respondents could hypothetically receive the money in 9 days rather than 7 days. More specifically, the question was: *Suppose that you are to receive money from someone. You can either choose to receive the money <u>2 days</u> from today, or <u>9 days</u> from today, but the amounts will be different. Compare the amounts and dates below in Option 'A" and Option 'B", and indicate which option you prefer for each of the <u>eight</u> choices.* Here, we underline the differences. Although the length of τ remains the same, this change might affect the distribution of present bias. In fact, although the mean of present bias changes between the periods before and after 2010, we do not observe a substantial change in the fraction of respondents with $z_{it}^{DC} = 1$.

Second, the amount of money that respondents hypothetically receive differs substantially between the surveys before and after 2011. Before 2011, respondents chose between option A, which offered 10,000 yen today, and option B, which offered a specific amount either above or below 10,000 yen in seven days (e.g., 10,019 yen or 9,980 yen). After 2011, however, as Table A.1 shows, the amount hypothetically received in option A is slightly above 3,000 yen, while option B offers a similar but slightly different amount. Although this change does not appear to affect the overall distribution of respondents with present bias very much, the fraction of respondents without present bias declines substantially. In particular, while the fraction of respondents without present bias before 2011 ranges from 83.9% to 86.6%, the corresponding fraction from 2011 onward ranges from 70.1% to 74.8%. We therefore take this shift into account.

[37] Even if we drop respondents with $\beta > 1$, the regression results are essentially unchanged.

A.2 Measuring risk attitudes

Examples of the choices in the survey are shown in Table A.2. The amounts shown in the table can be interpreted as reservation prices in the hypothetical lottery. The question implies that the expected return from the lottery is 50,000 yen. Risk attitudes are identified by the choice at which the respondent switches from Option 'A" to Option 'B." For example, if a respondent switches between 8,000 yen and 15,000 yen, the reservation price is defined as the midpoint of these two amounts (i.e., 11,500 yen). On the other hand, if a respondent never switches to Option "B," the reservation price is set to 50,000 yen. In this case, the respondent is classified as risk-neutral, since the reservation price equals the expected return of the risky lottery.

We make two remarks on the measurement of risk attitudes. First, the survey questions changed over time. In particular, the expected value of the lottery is not necessarily the same across survey years. Specifically, the expected hypothetical return from the lottery is only 1,000 yen in the surveys from 2004 to 2009. Accordingly, the reservation price may differ from the case in which the expected return from the lottery is 50,000 yen if loss aversion is present. In other words, respondents may be less risk-averse when the expected return is 1,000 yen. When the expected return is low, the potential loss from buying the lottery is also low. In this case, respondents may be more willing to take the risk of buying the lottery.[38]

Second, for robustness, Cramer et al. (2002) and Hanaoka et al. (2018) also use the Arrow-Pratt measure of absolute risk aversion.[39] This measure is a monotonically increasing but nonlinear function of the reservation price.

A.3 Changes in distribution of BA over time

One advantage of the JHPS-CPS is that respondents' BA are surveyed in each survey wave, allowing BA to vary over time. However, an important caveat is that the time-series variation available in our analysis is limited. In addition, some survey questions changed slightly over time. The JHPS-CPS also changed its mode of data collection from in-person home-visit interviews to mail surveys from 2021 onward. It is therefore useful to examine whether these changes affected the distribution of BA and whether they matter for the main conclusion of this paper.

[38]In fact, the mean of z_{it}^{RA} during 2007–2008 is 0.960, which is higher than the common-sample mean of 0.76 and implies a lower degree of risk aversion.

[39]See Cramer et al. (2002) for further details.

Table A.3 extends Table 1 by reporting the evolution of the distribution of BA over time. The table is based on the common sample. In the pooled common sample shown in the first column, the fractions of observations classified as dynamically consistent, as having forward-looking planning ability, and as having self-control are 76.9%, 32.6%, and 50.7%, respectively.

The table shows that the distributions exhibit a possible distribution change between 2010 and 2016. This change may be partly due to the revision of the survey questions in 2011 used to elicit discount factors and present bias. In 2011, the JHPS-CPS substantially changed the reference amount used to measure subjective discount factors and present bias, from 3,000 yen to 10,000 yen. Regarding forward-looking planning ability and self-control, we find no changes in the corresponding survey questions between these two years. One possible interpretation is that the Great East Japan Earthquake, which hit Japan in March 2011, affected these BA in a non-negligible way. Hanaoka et al. (2018) and Akesaka (2019) use JHPS-CPS data and provide evidence that respondents' preferences changed after the Great East Japan Earthquake. By contrast, we do not observe a large change in the distribution between 2016 and 2021, suggesting that the change in the mode of data collection did not strongly affect the distributions of BA.

One simple way to examine whether changes in the distribution of BA affect our main results is to split the sample into subperiods before and after the possible change in the distributions. However, as discussed in Section 3.2.1, the limited time-series variation in the data makes the estimates of regression coefficients unstable. We therefore take an alternative approach that preserves the full set of survey years while reducing the influence of year-specific fluctuations in measured BA.

To address possible distribution changes of BA, we construct time-invariant measures of BA for respondent i from z^j_{it}, where $j = DC, FL, SC$. Specifically, we first take the time-series mean $\bar{z}^j_i$ for respondent i and BA j, and then define the time-invariant dummy variable $\mathcal{Z}^j_i$ as

$$\mathcal{Z}^j_i = \begin{cases} 1 & \text{if } \bar{z}^j_i \geq 0.5, \\ 0 & \text{otherwise,} \end{cases}$$

where $\bar{z}^j_i$ is the sample mean of the dummy z^j_{it} over the relevant survey years and $j = DC, FL, SC$.

We construct three versions of the time-series mean $\bar{z}^j_i$: (i) the time-series mean over all available years in the common sample, denoted by $\bar{z}^j_{i,ALL}$; (ii) the time-series mean over the years before the possible distributional change, denoted by $\bar{z}^j_{i,Before}$; and (iii) the time-series mean over the years after the possible distributional change, denoted by $\bar{z}^j_{i,After}$.

We then estimate a regression similar to Equation (3), replacing the time-varying BA measures z_{it} with the time-invariant measures $\mathcal{Z}_i$:

$$\pi_{t+1} - \pi^e_{it+1} = a_i + by_t + b'_z \mathcal{Z}_i y_t + c' X_{it} + v_{it+1}, \tag{5}$$

where $\mathcal{Z}_i = (\mathcal{Z}^{DC}_i, \mathcal{Z}^{FL}_i, \mathcal{Z}^{SC}_i, \tilde{\bar{z}}^{RA}_i)'$, and $\tilde{\bar{z}}^{RA}_i$ is the standardized version of the time-series mean of risk aversion. To preserve the time-series variation in y_t, we use all available years in the common sample. The measure based on $\mathcal{Z}^j_{i,ALL}$ smooths out year-to-year changes in BA because $\bar{z}^j_i$ averages z^j_{it} over all years in the common sample. In addition, comparing the results based on $\mathcal{Z}^j_{i,Before}$ and $\mathcal{Z}^j_{i,After}$ allows us to examine whether the estimated heterogeneity in responses to the public signal is sensitive to the possible distributional change in BA.

The results are not substantially different from the baseline specification. The main coefficient on the public signal remains negative across specifications. Although the estimated coefficients on the interaction terms, especially those for $\mathcal{Z}^{FL}_i$, vary depending on how the time-invariant BA measures are constructed, they are not large enough to offset the negative main response to the public signal. Thus, the evidence remains consistent with overreaction in inflation expectations even when BA are measured by their time-series averages.

These results suggest that our main finding is not driven by year-specific fluctuations or distributional changes in measured BA. Even when BA are smoothed over time and treated as time-invariant respondent-level attributes, overreaction to the public signal is not concentrated among respondents classified as behaviorally biased.

A.4 Excluding observations in the top-coded category of inflation expectations

As discussed in Section 2.2, a number of observations of inflation expectations fall into the top-coded response category during the high-inflation period. This may raise measurement issues for inflation expectations.

To address this concern, Table A.5 reports estimation results after dropping observations in which inflation expectations are classified in the top-coded response category. The table presents the results without heterogeneity in the response to the public signal, using the extended sample that includes all observations for which forecast errors and the public signal are available. Compared with the results in Table 3, the estimated coefficients become somewhat smaller in absolute value, but they remain negative and statistically significant in both subsamples. In addition, the high-inflation-period dummy indicates a statistically significant difference in the intercept, whereas the interaction between the dummy and the public signal is not statistically significant. This result supports our benchmark specification discussed in the main text.

In Table A.6, we estimate Equation (3) using the common sample after excluding observations in which inflation expectations are classified in the top-coded response category. The results are directly comparable to those in Table 4. Overall, the estimation results are qualitatively similar, except for the statistically insignificant coefficient on the interaction term with $\tilde{z}_{it}^{RA}$. The signs of the estimated main coefficient on the public signal are the same as those in Table 4, although their magnitudes are slightly smaller in absolute value. Again, the heterogeneity in the response to the public signal is quantitatively small. Thus, the results remain consistent with our evidence that forecast errors respond negatively to the public signal, regardless of whether respondents are behaviorally disciplined.

Figure 1: Year-on-year inflation and mean one-year-ahead inflation expectations

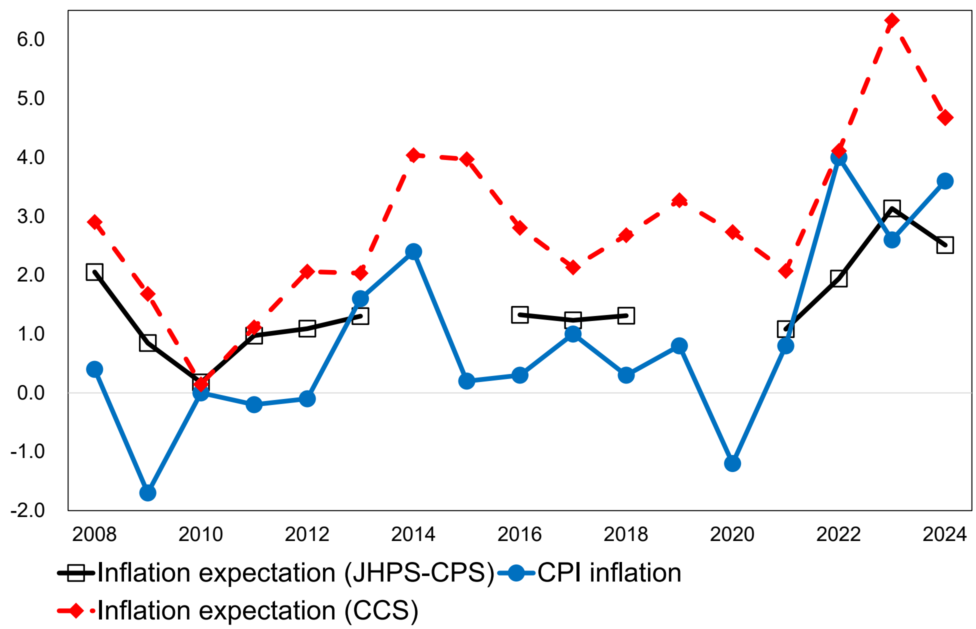

Notes: The figure displays the mean of respondents' inflation expectations and headline inflation from 2007 to 2024. The line with squares represents inflation expectations elicited from the JHPS-CPS, and the line with circles represents actual inflation. For comparison, the figure also plots inflation expectations calculated from the CCS. The x-axis shows years, and the y-axis is measured in percent.

Figure 2: Histogram of inflation expectations in 2023

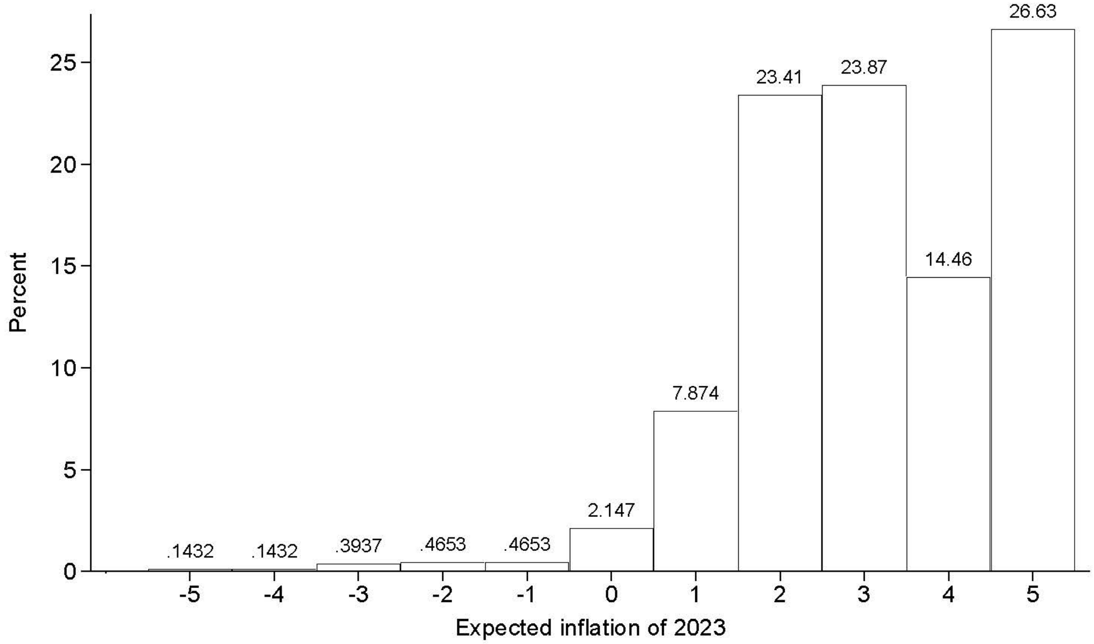

Notes: The figure displays the distribution of inflation expectations in 2023. The x-axis shows inflation expectations, and the y-axis is measured in percent.

Figure 3: Forecast errors of mean inflation expectations

Notes: The figure displays the forecast error of mean inflation expectations from 2007 to 2024. The shaded area represents the cross-sectional variation in forecast errors, based on the 5th – 95th percentile range. The x-axis shows years, and the y-axis is measured in percent.

Table 1: Distributions of behavioral attributes

Behavioral attributes	Common Sample	Low RA	High RA
Single attributes			
DC	76.9	74.8	79.4
FL	32.6	33.8	31.2
SC	50.7	49.3	52.4
Multiple attributes: $(z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC})$			
(1) DC, FL, SC	13.7	13.2	14.4
(2) DC, FL	10.9	11.7	10.1
(3) DC, SC	25.1	23.4	27.1
(4) FL, SC	4.5	5.2	3.8
(5) DC only	27.2	26.6	27.8
(6) FL only	3.4	3.8	3.0
(7) SC only	7.4	7.6	7.1
(8) None	7.8	8.7	6.8
Observations	8,858	4,856	4,002

Notes: The upper panel shows the percentage of observations classified as dynamically consistent, forward-looking, and self-controlled, respectively, regardless of the presence of other BA. The lower panel reports the percentage of observations in each mutually exclusive type defined by combinations of BA; the percentages in each column sum to 100%. Because BA are measured in each survey wave and are allowed to vary over time, the table reports percentages based on observations rather than on unique respondents. Type (1) includes observations with all three BA; types (2)–(4) include observations with exactly two BA; types (5)–(7) include observations with only one BA; and type (8) includes observations with none of the three BA. The columns report the corresponding percentages for the common sample ($N = 8,858$) and for respondents with low and high risk aversion. Low RA and High RA denote respondents below and above the sample median of risk aversion, respectively.

Table 2: Estimation results without heterogeneity in behavioral attributes

	(1)	(2)	(3)	(4)	(5)
y_t	-0.162	-0.806***	-0.077	0.058	-0.822***
	(0.252)	(0.074)	(0.560)	(0.399)	(0.079)
$y_t \times$ Dummy			-0.749		
for high-inflation period			(0.564)		
Dummy for		3.458***	3.632***		
high-inflation period		(0.397)	(0.287)		
Control variables	Y	Y	Y	Y	Y
Observations	8,858	8,858	8,858	3,425	4,057
R-squared	0.463	0.668	0.671	0.577	0.741

Notes: The dependent variable is the forecast error, defined as headline CPI inflation minus respondents' inflation expectations. Standard errors, reported in parentheses, are two-way clustered by respondent and year. Inference is based on the t-distribution with $T-1$ degrees of freedom, where T denotes the number of years used in each regression. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. All regressions include individual fixed effects, respondents' log taxable income, age-group dummies, and a dummy for liquidity-constrained respondents as control variables. Specifications (2) and (3) include a dummy for the high-inflation period to allow for a possible structural break in inflation expectations. Specifications (4) and (5) report results based on subsamples. Specification (4) uses data for $t+1 = 2008, 2009, 2011$, and 2017, whereas specification (5) uses data for $t+1 = 2022, 2023$, and 2024. The numbers of observations in the subsample regressions do not necessarily sum to the number of observations in the full-sample regression because singleton observations are dropped when individual fixed effects are estimated.

Table 3: Estimation results without heterogeneity in behavioral attributes (Extended sample)

	(1)	(2)	(3)	(4)	(5)
y_t	-0.010	-0.703***	-0.640***	-0.640**	-0.791**
	(0.133)	(0.127)	(0.200)	(0.202)	(0.084)
$y_t\times$ Dummy			-0.149		
for high-inflation period			(0.212)		
Dummy for		3.589***	3.784***		
high-inflation period		(0.513)	(0.335)		
Control variables	N	N	N	N	N
Observations	42,516	42,516	42,516	34,088	8,183
R-squared	0.256	0.466	0.467	0.354	0.702

Notes: All regressions are estimated using the extended sample, which includes all observations for which forecast errors and the public signal are available. To increase the number of observations, the regressions include individual fixed effects but exclude the other control variables used in Table 2. Specification (4) uses data for $t + 1 = 2008, 2009, 2010, 2011, 2012, 2013, 2016, 2017, 2018$, and 2021, whereas specification (5) uses data for $t + 1 = 2022, 2023$, and 2024. See the notes to Table 2 for further details.

Table 4: Estimation results with heterogeneity in behavioral attributes

	(1)	(2)	(3)	(4)	(5)
y_t	-0.922***	-0.783***	-0.813***	-0.788***	-0.887***
	(0.078)	(0.076)	(0.069)	(0.074)	(0.073)
$z_{it}^{DC} \times y_t$	0.161***				0.157***
	(0.035)				(0.037)
$z_{it}^{FL} \times y_t$		-0.066**			-0.064**
		(0.021)			(0.021)
$z_{it}^{SC} \times y_t$			0.015		0.016
			(0.026)		(0.025)
$\tilde{z}_{it}^{RA} \times y_t$				0.067**	0.064**
				(0.020)	(0.019)
z_{it}^{DC}	-0.404**				-0.389**
	(0.113)				(0.111)
z_{it}^{FL}		0.272***			0.266***
		(0.058)			(0.057)
z_{it}^{SC}			0.101		0.096
			(0.054)		(0.050)
$\tilde{z}_{it}^{RA}$				-0.106	-0.105
				(0.093)	(0.090)
Dummy for	3.429***	3.440***	3.453***	3.450***	3.399***
high inflation period	(0.387)	(0.391)	(0.395)	(0.394)	(0.376)
Control variables	Y	Y	Y	Y	Y
Observations	8,858	8,858	8,858	8,858	8,858
R-squared	0.671	0.669	0.669	0.670	0.674

Notes: The regressions allow the response to the public signal to vary with BA. Variables with a tilde are standardized to have mean zero and unit variance. See the notes to Table 2 for further details.

Table 5: Responses of behaviorally disciplined respondents

Types	Responses	
	Forecast errors	Inflation expectations
(1) DC, FL, SC	-0.778***	1.285
	(0.077)	
(2) DC, FL	-0.794***	1.301
	(0.065)	
(3) DC, SC	-0.714***	1.221
	(0.085)	
(4) FL, SC	-0.935***	1.442
	(0.080)	

Notes: The estimated responses of forecast errors and inflation expectations for each type are based on specification (5) in Table 4. Types correspond to the categories reported in the lower panel of Table 1. The column labeled "Forecast errors" reports the sum of the estimated coefficient on the public signal and the corresponding interaction coefficients for each type. We test the null hypothesis that this sum equals zero. The column labeled "Inflation expectations" reports the implied response of inflation expectations to the public signal. Because the forecast error is defined as actual inflation minus expected inflation, the implied response of inflation expectations is computed as 0.507 minus the response of forecast errors. The value 0.507 is the estimated response of actual inflation to the public signal, obtained from regressing π_{t+1} on y_t using the sample for $t + 1 = 2008, \ldots, 2024$. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Robustness analysis

	(1)	(2)	(3)	(4)	(5)	(6)
y_t	-0.887***		-0.895***	-0.173***	-0.816***	-0.272
	(0.073)		(0.023)	(0.029)	(0.078)	(0.156)
$z_{it}^{DC} \times y_t$	0.157***	0.090***	0.132***	0.062*	0.164***	0.099***
	(0.037)	(0.022)	(0.026)	(0.031)	(0.038)	(0.022)
$z_{it}^{FL} \times y_t$	-0.064**	-0.032	-0.051**	-0.022	-0.082**	-0.039*
	(0.021)	(0.017)	(0.019)	(0.013)	(0.028)	(0.018)
$z_{it}^{SC} \times y_t$	0.016	0.009	0.002	0.005	0.023	0.008
	(0.025)	(0.023)	(0.018)	(0.009)	(0.029)	(0.025)
$\tilde{z}_{it}^{RA} \times y_t$	0.064**	0.050***	0.052**	0.058	0.054**	0.042***
	(0.019)	(0.010)	(0.013)	(0.037)	(0.017)	(0.009)
z_{it}^{DC}	-0.389**	-0.124**	-0.275***	-0.391*	-0.395**	-0.139**
	(0.111)	(0.049)	(0.070)	(0.200)	(0.109)	(0.044)
z_{it}^{FL}	0.266***	0.112**	0.209***	0.211**	0.292***	0.115**
	(0.057)	(0.045)	(0.041)	(0.077)	(0.054)	(0.045)
z_{it}^{SC}	0.096	0.049	0.089*	0.080	0.099*	0.050
	(0.050)	(0.057)	(0.046)	(0.065)	(0.045)	(0.054)
$\tilde{z}_{it}^{RA}$	-0.105	-0.048	-0.039	-0.416	-0.088	-0.037
	(0.090)	(0.030)	(0.057)	(0.231)	(0.083)	(0.030)
Dummy for	3.399***		3.371***	6.685***	3.262***	
high inflation period	(0.376)		(0.247)	(0.241)	(0.331)	
Control variables	Y	Y	Y	Y	Y	Y
Observations	8,858	8,858	8,858	8,858	8,858	8,858
R-squared	0.674	0.726	0.680	0.913	0.661	0.720

Notes: Specification (1) replicates specification (5) in Table 4 for comparison. Specification (2) includes year fixed effects and identifies differential responses to y_t across behavioral attributes. Specifications (3)–(6) use alternative measures of consumer prices to construct forecast errors and the public signal. Specifically, specification (3) uses core CPI inflation, and specification (4) uses fresh-food inflation. Specifications (5) and (6) use city-level headline CPI inflation to construct both forecast errors and the public signal. Specification (5) includes the high-inflation-period dummy, as in the baseline specification. Specification (6) includes year fixed effects; because the city-level public signal varies across both years and cities, its main coefficient remains identified. See the notes to Table 4 for further details.

Table 7: Alternative measures of the public signal

	(1)	(2)	(3)	(4)	(5)
y_t	-0.887***	-0.187**	-0.235	-0.311	-0.523
	(0.073)	(0.061)	(0.360)	(0.167)	(0.889)
$z_{it}^{DC} \times y_t$	0.157***	0.059	0.194*	0.159***	0.408
	(0.037)	(0.034)	(0.097)	(0.033)	(0.211)
$z_{it}^{FL} \times y_t$	-0.064**	-0.023	-0.057	-0.082***	-0.163
	(0.021)	(0.014)	(0.050)	(0.020)	(0.092)
$z_{it}^{SC} \times y_t$	0.016	0.011	0.018	-0.000	0.061
	(0.025)	(0.010)	(0.034)	(0.014)	(0.079)
$\tilde{z}_{it}^{RA} \times y_t$	0.064**	0.060	0.003	0.092*	0.082
	(0.019)	(0.041)	(0.030)	(0.042)	(0.052)
z_{it}^{DC}	-0.389**	-0.424	-0.288*	-0.215*	-0.402**
	(0.111)	(0.224)	(0.121)	(0.098)	(0.175)
z_{it}^{FL}	0.266***	0.240**	0.154*	0.139**	0.205**
	(0.057)	(0.089)	(0.070)	(0.053)	(0.076)
z_{it}^{SC}	0.096	0.109	0.181**	0.145**	0.160**
	(0.050)	(0.063)	(0.054)	(0.040)	(0.062)
$\tilde{z}_{it}^{RA}$	-0.105	-0.448	-0.078	0.027	-0.111
	(0.090)	(0.257)	(0.079)	(0.094)	(0.086)
Dummy for	3.399***	1.545**	1.723	2.823**	1.782
high inflation period	(0.376)	(0.469)	(0.929)	(0.908)	(1.073)
Control variables	Y	Y	Y	Y	Y
Observations	8,858	8,858	8,858	8,858	8,858
R-squared	0.674	0.604	0.541	0.577	0.542

Notes: Specification (1) replicates specification (5) in Table 4 for comparison. In specifications (2)–(6), we use alternative public signals while continuing to construct forecast errors from headline CPI inflation. Specifically, specifications (2)–(5) use fresh-food inflation, the mean professional forecast, month-on-month headline CPI inflation, and the experienced inflation, respectively, as the public signal. See the notes to Table 4 for further details.

Table A.1: Examples of choices for time preferences in the JHPS-CPS

	Option “A”	Option “B”	Which one do you prefer?	
Choices	Receive today	Receive 7 days from today	X one box for each row	
1	¥3,005	¥3,014	1 ☐	2 ☐
2	¥3,003	¥3,297	1 ☐	2 ☐
3	¥3,008	¥3,037	1 ☐	2 ☐
4	¥3,000	¥3,000	1 ☐	2 ☐
5	¥3,005	¥5,951	1 ☐	2 ☐
6	¥3,009	¥3,068	1 ☐	2 ☐
7	¥3,001	¥3,119	1 ☐	2 ☐
8	¥3,002	¥2,996	1 ☐	2 ☐
9	¥3,008	¥3,011	1 ☐	2 ☐

Table A.2: Examples of choices for risk aversion in the JHPS-CPS

		Which one do you prefer?	
		Option “A”	Option “B”
Choices	Price of the “speed lottery” ticket	(buy the “speed lottery” ticket)	(DO NOT buy the “speed lottery” ticket)
1	¥10	1 ☐	2 ☐
2	¥2,000	1 ☐	2 ☐
3	¥4,000	1 ☐	2 ☐
4	¥8,000	1 ☐	2 ☐
5	¥15,000	1 ☐	2 ☐
6	¥25,000	1 ☐	2 ☐
7	¥35,000	1 ☐	2 ☐
8	¥50,000	1 ☐	2 ☐

Table A.3: Distributions of behavioral attributes over time

Year	Common sample	2007	2008	2010	2016	2021	2022	2023
Behavioral attributes								
Single attributes								
DC	76.9	88.5	90.4	85.2	71.2	70.0	73.1	72.0
FL	32.6	23.1	25.5	23.8	34.1	36.5	38.9	38.8
SC	50.7	46.9	49.7	47.6	51.6	55.4	49.4	52.6
Multiple attributes: $(z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC})$								
(1) DC, FL, SC	13.7	10.7	12.2	10.6	15.0	14.7	14.9	16.4
(2) DC, FL	10.9	9.4	11.6	9.5	9.4	10.5	13.1	12.6
(3) DC, SC	25.1	31.3	32.9	30.1	22.3	24.1	20.7	21.2
(4) FL, SC	4.5	1.3	1.0	1.8	5.1	6.9	6.1	6.1
(5) DC only	27.2	37.2	33.7	35.0	24.6	20.7	24.4	21.9
(6) FL only	3.4	1.7	0.8	1.9	4.6	4.4	4.8	3.7
(7) SC only	7.4	3.6	3.6	5.1	9.2	9.7	7.8	8.9
(8) None	7.8	4.8	4.2	6.0	9.8	9.0	8.3	9.3
Observations	8,858	524	501	2,071	1,213	1,431	1,585	1,533

Notes: This table reports the annual distribution of observations by BA within the common sample ($N = 8,858$). The upper panel reports the percentage of observations classified under each individual BA. The lower panel reports column percentages for mutually exclusive combinations of BA. See the notes to Table 1 for the definitions of the BA categories.

Table A.4: Robustness: Time-invariant measures of behavioral attributes

	(1)		(2)	(3)	(4)
y_t	-0.887***	y_t	-0.869***	-0.811***	-0.868***
	(0.073)		(0.070)	(0.081)	(0.084)
$z_{it}^{DC} \times y_t$	0.157***	$\mathcal{Z}_i^{DC} \times y_t$	0.075*	0.045	0.055*
	(0.037)		(0.034)	(0.033)	(0.025)
$z_{it}^{FL} \times y_t$	-0.064**	$\mathcal{Z}_i^{FL} \times y_t$	0.000	-0.053*	0.033*
	(0.021)		(0.018)	(0.024)	(0.016)
$z_{it}^{SC} \times y_t$	0.016	$\mathcal{Z}_i^{SC} \times y_t$	0.019	-0.016	0.014
	(0.025)		(0.022)	(0.026)	(0.017)
$\tilde{z}_{it}^{RA} \times y_t$	0.064**	$\tilde{\tilde{z}}_i^{RA} \times y_t$	0.085***	0.055*	0.080***
	(0.019)		(0.019)	(0.025)	(0.016)
z_{it}^{DC}	-0.389**				
	(0.111)				
z_{it}^{FL}	0.266***				
	(0.057)				
z_{it}^{SC}	0.096				
	(0.050)				
$\tilde{z}_{it}^{RA}$	-0.105				
	(0.090)				
Dummy for	3.399***		3.440***	3.461***	3.445***
high inflation period	(0.376)		(0.389)	(0.391)	(0.392)
Control variables	Y		Y	Y	Y
Observations	8,858		8,858	8,076	8,405
R-squared	0.674		0.671	0.665	0.664

Notes: Specification (1) replicates specification (5) in Table 4 for comparison. In specifications (2)–(4), we use the time-invariant measures of BA ($\mathcal{Z}_i^{DC}, \mathcal{Z}_i^{FL}, \mathcal{Z}_i^{SC}$ and $\tilde{\tilde{z}}_i^{RA}$) constructed from the time-series means of BA. Specifications (2), (3), and (4) take the mean over all available years in the common sample, the mean over the years before the distribution change, and the mean over the years after the distribution change, respectively. Note that these specifications cannot identify the effect of BA on the mean since individual fixed effects absorb the variation across respondents. See the notes to Table 4 for further details.

Table A.5: Estimation results without heterogeneity in behavioral attributes (Extended sample excluding observations in which inflation expectations are categorized as the top-coded category)

	(1)	(2)	(3)	(4)	(5)
y_t	0.061	-0.628***	-0.605**	-0.606**	-0.686**
	(0.146)	(0.130)	(0.203)	(0.206)	(0.071)
$y_t\times$ Dummy			-0.059		
for high-inflation period			(0.212)		
Dummy for		3.659***	3.738***		
high-inflation period		(0.510)	(0.307)		
Control variables	N	N	N	N	N
Observations	39,735	39,735	39,735	32,680	6,567
R-squared	0.250	0.500	0.501	0.347	0.685

Notes: All regressions are the same as those in Table 3 except that we exclude observations in which inflation expectations are classified in the top-coded category (i.e., inflation expectations are set to 5%). See the notes to Tables 2 and 3 for further details.

Table A.6: Estimation results with heterogeneity in behavioral attributes (Common sample excluding observations in which inflation expectations are categorized as the top-coded category)

	(1)	(2)	(3)	(4)	(5)
y_t	-0.772***	-0.651***	-0.678***	-0.659***	-0.735***
	(0.074)	(0.074)	(0.072)	(0.069)	(0.077)
$z_{it}^{DC} \times y_t$	0.135***				0.131***
	(0.028)				(0.029)
$z_{it}^{FL} \times y_t$		-0.069**			-0.068**
		(0.024)			(0.023)
$z_{it}^{SC} \times y_t$			0.005		0.010
			(0.027)		(0.025)
$z_{it}^{RA} \times y_t$				0.062**	0.060**
				(0.021)	(0.020)
z_{it}^{DC}	-0.394**				-0.384**
	(0.110)				(0.108)
z_{it}^{FL}		0.239***			0.235***
		(0.061)			(0.059)
z_{it}^{SC}			0.057		0.049
			(0.052)		(0.049)
z_{it}^{RA}				-0.110	-0.112
				(0.090)	(0.087)
Dummy for	3.408***	3.418***	3.432***	3.429***	3.384***
high inflation period	(0.350)	(0.357)	(0.362)	(0.360)	(0.342)
Control variables	Y	Y	Y	Y	Y
Observations	7,762	7,762	7,762	7,762	7,762
R-squared	0.710	0.708	0.708	0.709	0.713

Notes: The regressions allow the response to the public signal to vary with BA. In the estimation, we drop forecast errors of inflation expectations from the top-coded response categories. See the notes to Table 2 for further details.

Table A.7: Responses of forecast errors for behaviorally disciplined respondents

Types	Response of forecast errors			
	Core CPI	Fresh foods	City CPI (Dummy)	City CPI (Year FE)
(1) DC, FL, SC	-0.812***	-0.128**	-0.711***	-0.204
	(0.019)	(0.040)	(0.089)	(0.163)
(2) DC, FL	-0.814***	-0.133**	-0.734***	-0.212
	(0.020)	(0.044)	(0.069)	(0.152)
(3) DC, SC	-0.761***	-0.106*	-0.629***	-0.165
	(0.028)	(0.050)	(0.112)	(0.175)
(4) FL, SC	-0.944***	-0.190***	-0.875***	-0.302
	(0.023)	(0.019)	(0.078)	(0.156)

Notes: The responses of forecast errors are estimated using the alternative measures of consumer prices examined in Table 6. More specifically, the columns labeled "Core CPI," "Fresh foods," "City CPI (Dummy)," and "City CPI (Year FE)" correspond to specifications (3)–(6) of that table, respectively.