ISSN (Print) 0473-453X
ISSN (Online) 2435-0982

Discussion Paper No. 1309

OVERREACTION IN INFLATION EXPECTATIONS: DO BEHAVIORAL ATTRIBUTES OF INDIVIDUALS MATTER?

Takayuki Tsuruga
Zichong Yue

March 2026

The Institute of Social and Economic Research
The University of Osaka
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

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

Takayuki Tsuruga† and Zichong Yue‡

March 29, 2026

Abstract

Previous studies on inflation expectations suggest that agents often overreact to public signals, pointing to a role for behavioral explanations and motivating extensions of standard macroeconomic models beyond full-information rational expectations. Using the survey data that include a variety of questions related to behavioral economics, we examine how behavioral attributes of individuals interact with the overreaction to public signals. We find that even individuals with rational behavioral attributes overreact to public signals. These findings suggest that overreaction is a distinct phenomenon that cannot be fully attributed to well-documented behavioral biases such as present bias or myopia.

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 Institute of Social and Economic Research and the Kyoto Institute of Economic Research, and from Grants-in-Aid for Scientific Research (JP20H05631, JP20H05633, JP21H04397, JP25K00622, and JP25H00544).

†Institute of Social and Economic Research, The University of Osaka; The Centre for Applied Macroeconomic Analysis; Correspondence to: Institute of Social and Economic Research, The University of Osaka, 6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan. E-mail: tsuruga@iser.osaka-u.ac.jp

‡Graduate School of Economics, The University of Osaka.

1 Introduction

While inflation expectations are central to macroeconomic analysis, reaching a consensus on their interpretation and modeling remains a challenging task. Empirical studies on inflation expectations have reached a broad consensus that the full-information rational expectations (FIRE) hypothesis receives little support from the data.[1] Recent studies further show that economic agents often exhibit overreaction to the information observable at the time of forming their expectations. They document that forecast errors of inflation expectations are statistically correlated with information available at the time of forming expectations, which is against 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) also find individuals' forecast errors of 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 the diagnostic expectations, in which past news is overly emphasized.

In this paper, we examine the overreaction of inflation expectations from a behavioral economics perspective. In particular, we study how overreaction to observable information interacts with individuals' behavioral attributes (hereafter, BA). Following Kohlhas and Walther (2021) and Broer and Kohlhas (2024), we measure the extent to which individuals' one-year-ahead inflation expectations overreact to public signals observable at the time expectations are formed. 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*, 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. In each January since 2003, the JHPS-CPS has surveyed respondents' inflation expectations and behavioral attributes, such as present bias, forward-looking planning ability, self-control, and risk attitudes. Although the survey was

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

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

not conducted in some years, these rich data on BA allow us to examine the relationship between overreaction in 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 remains poorly understood. 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 behave more consistently with the FIRE hypothesis to make more accurate forecasts of the real interest rate? 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 pervasive across individuals, regardless of BA? If overreaction is primarily concentrated among individuals with irrational BA, then models that focus exclusively on overreaction may be incomplete, and alternative behavioral assumptions may be more informative. By contrast, if overreaction is pervasive even among respondents with otherwise rational BA, then parsimonious models that allow for overreaction while maintaining other rational features may provide a reasonable approximation.

We find that Japanese households pervasively overreact to public signals. More specifically, although BA partly account for heterogeneity in overreaction, their effects are quantitatively minor. For example, even respondents who report no present bias tend to overreact. Although overreaction to public signals is weaker among households without present bias than among those with present bias, it remains statistically and economically significant. Similarly, households that report being forward-looking and having self-control also exhibit overreaction. More risk-averse respondents tend to respond more weakly to public signals, but the effect is not strong enough to eliminate overreaction.

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, and 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* (hereafter, the 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 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 who were selected through stratified random sampling.[3] 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, namely 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 the analysis below, we use data from 2005 to 2024 because of data availability.[4]

In our regression analysis, we regress forecast errors of inflation expectations on public signals and their interactions with BA constructed from the survey data. 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 expected inflation in the corresponding survey year. Specifically, the question is: "*By what percentage do you*

[3]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.

[4]Although the data are available from 2003 onward, the questionnaire format changed in some years, which may affect respondents' answers. To ensure consistency in the questionnaire format, we use data from 2005 onward.

expect consumer prices will change in [survey year]*, compared with the previous year?*" The answer choices 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%*; *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 a measure of inflation expectations, we use the midpoint of each response category.

To construct forecast errors, we require measures of both actual inflation and inflation expectations. Let π_t denote inflation measured by the change in prices from January to December of year t. We use historical inflation data to measure π_t. The Statistics Bureau of the Ministry of Internal Affairs and Communications (MIAC) reports preliminary year-on-year inflation rates in January for the previous year's Consumer Price Index (CPI) for all items and the core CPI (i.e., the CPI: All items, less fresh food). By contrast, in January the JHPS-CPS asks respondents about their expectations for changes in consumer prices through December of the same year. We denote respondent i's inflation expectation formed in January of year t by π^e_{it}.

In the baseline specification, we use core CPI inflation as the measure of π_t. Thus, forecast errors, $\pi_t - \pi^e_{it}$, are defined using core CPI inflation. Of course, respondents may not necessarily interpret consumer prices in the survey as referring to the core CPI. To address this issue, we also use alternative measures of consumer prices.

Figure 1 plots the cross-sectional mean of inflation expectations from the JHPS-CPS (i.e., the consensus forecast), shown by the solid line with squares, over the period 2004–2024. The time-series mean of the consensus forecast over this period is 1.31%. However, when we divide the sample into 2004–2021 (the low-inflation period) and 2022–2024 (the high-inflation period), the average consensus forecast is 1.05% 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, which is one of the main sources of data on inflation expectations in Japan. Inflation expectations from the CCS also increase markedly after 2021. A comparison of the two series shows that inflation expectations in 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.[5] The second is the rise in inflation following the COVID-19 pandemic. The time-series mean of actual inflation over 2004–2024 is 0.68%. When we divide the sample into the low- and high-inflation periods, the time-series mean of the inflation rate is 0.25% in the former period and 3.1% in the latter.

Comparing the time-series means of the consensus forecast from the JHPS-CPS with actual inflation, we confirm that expected inflation is, on average, higher than actual inflation over the full sample period (1.31% versus 0.68%). This pattern is consistent with D'Acunto et al. (2023), who document an upward bias in inflation expectations. However, for the consensus forecast from the JHPS-CPS, we do not observe a pronounced upward bias during the high-inflation period. The time-series mean of the consensus forecast in the JHPS-CPS is 2.53%, which is lower than the corresponding mean of actual inflation, 3.10%. By contrast, we observe a significant upward bias in the consensus forecast from the CCS 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 imperfect because of 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 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, defined as $\pi_t - \pi^e_{it}$, together with the interval between the 2.5th and 97.5th percentiles. The mean forecast error is negative except in 2022 and 2024, implying that inflation expectations exceed actual inflation in all other years. The interval is relatively stable over time.

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

2.3 Behavioral attributes of individuals

2.3.1 Present bias in time preferences

It is well known that preferences with present bias 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 to take an action that provides immediate utility when deciding in the present, but revise their choice when they reach a future period and reoptimize.[6] 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 discount factor, whereas 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 $t + \tau$ (i.e., 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.

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 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 every year. 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 the BA of forward-looking planning ability and self-control. To examine the effects of these two factors, we follow the approach of Kubota and Fukushige (2016).[7]

We begin with forward-looking planning ability. Using the same survey dataset, Kubota

[6] In other words, agents are likely to regret their current choices in the future.

[7] Unfortunately, these two questions were not included in four years: 2004, 2005, 2006, and 2011.

and Fukushige (2016) find that respondents with forward-looking planning ability tend to predict future income changes more accurately, whereas those lacking such ability systematically overestimate future income changes. Following their approach, we use the same measure. However, rather than studying income forecasts, we examine whether forward-looking planning ability affects how individuals respond to public signals when forming inflation expectations.

The corresponding survey 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 the 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. Over the full sample, 29.18% of respondents choose 1 or 2.

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. Once again, because the survey asks this question every year, z_{it}^{FL} varies both across individuals and over time.

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 responses again consist of five categories. We define respondents who choose 4 or 5 as individuals with self-control in decision-making. Over the full sample, 20.94% of respondents are classified as having self-control.[8]

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.

We emphasize that the qualitative questions described above come with important caveats. First, the responses are self-reported. For example, even if respondents state that they always plan before acting, such responses do not necessarily imply true forward-looking planning ability. Second, these 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.

[8] Respondents with self-control may overlap with dynamically consistent respondents, since decisions with present bias imply a lack of self-control. However, self-control is conceptually distinct from present bias.

2.3.3 Risk attitude

The JHPS-CPS also measures respondents' risk attitudes as 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 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.

Overall, respondents in the survey are risk-averse. Over the sample period 2011–2024, only 1.58% of respondents are risk-neutral. As with time preferences and present bias, there is substantial heterogeneity in the degree of risk aversion. Over the period 2011–2024, the mean of z_{it}^{RA} is 0.785 and its standard deviation is 0.218.

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 category.[9] 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. The grand mean and median of taxable income are 4.47 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.[10] Following D'Acunto et al. (2023), we define young respondents as those under age 40 and middle-aged respondents as those aged 40 to 59. The young, middle-aged, and old groups account for 27.81%, 39.05%, and 33.14% of the sample, respectively.

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 accurate 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.[11] In the JHPS-CPS,

[9]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 survey, the first five categories account for 85.33% of respondents.

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

[11]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."

87.78% of respondents report not being liquidity-constrained in the full sample.

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 is the public signal included in respondents' information set at the time of forming their expectations. In the baseline specification, the public signal y_t is core CPI inflation in the previous year (i.e., $y_t = \pi_t$). Also, z_{it} is the vector of 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 v_{it+1}. Here, 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 for the control variables.

This regression extends the specification used by Kohlhas and Walther (2021), in which forecast errors are regressed on public signals and their interaction terms. Here, our goal is to understand how individuals' overreaction or underreaction to public signals varies with BA. This is captured by the coefficients on the interaction terms in (3). The parameter vector b_z measures how the response of forecast errors to the public signal varies with BA.

Although the dependent variable is the forecast error of inflation expectations in year $t+1$, we use one-year-lagged explanatory variables (i.e., z_{it} and X_{it} rather than z_{it+1} and X_{it+1}) to alleviate concerns about reverse causality. Using one-year-lagged variables allows us to reduce the possibility that respondents' measured BA are affected by their inflation expectations.

Under the FIRE hypothesis, under which π^e_{it+1} is specified as a rational expectation, a_i should be zero for all i because rational expectations imply unbiased forecasts of inflation. We impose the assumption that the sum of the individual fixed effects equals zero. In particular, a_i is given by

$$a_i = a + \sum_{j=1}^{N} \tilde{a}_j \mathbb{I}(j = i), \text{ and } \sum_{j=1}^{N} \tilde{a}_j = 0, \tag{4}$$

where $\mathbb{I}(j = i)$ is the indicator function that takes the value one when $j = i$. Here, a can be interpreted as the mean constant term because the individual fixed effects average to zero. In the subsequent sections, we report the estimated value of a for reference.

Similarly, under the FIRE hypothesis, b and b_z should also be zero because the public signal y_t is included in respondents' information set at the time of forming their expectations, and thus the forecast error $\pi_{t+1} - \pi^e_{it+1}$ is orthogonal to the public signal y_t.[12] Testing $b = 0$ and $b_z = 0$ corresponds to testing the joint hypothesis that respondents' inflation expectations are rational and that their information set includes the public signal (i.e., $y_t \in \Omega_{it}$).[13]

The literature has often found negative estimates of b in specifications without interaction terms with BA. When π_{t+1} and y_t are positively correlated, a negative response of forecast errors to the public signal (i.e., $b < 0$) implies that π^e_{it+1} responds more strongly to y_t than π_{t+1} does. In this case, 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.

In the regressions, we standardize the measure of risk aversion, z^{RA}_{it}.[14] 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, while the coefficient on the interaction term between y_t and $\tilde{z}^{RA}_{it}$ captures the additional change in the response of forecast errors to the public signal associated with a one-standard-deviation increase in risk-aversion measure relative to its grand mean.

One advantage of this regression is that the test is more practical than the test in Coibion and Gorodnichenko (2015a) and Bordalo et al. (2020), in which forecast errors are regressed on forecast revisions that are defined as the changes in forecast made in different periods. The JHPS-CPS collects only one-year-ahead forecasts and does not elicit two forecasts for the same target year in consecutive waves, making it impossible to construct forecast revisions from the data.

[12] Let Ω_{it} be respondent i's information set at the end of year t. Given that $\pi_t \in \Omega_{it}$, $\mathrm{Cov}(\pi_{t+1} - \mathbb{E}(\pi_{t+1}|\Omega_{it}), \pi_t) = \mathbb{E}\{[\pi_{t+1} - \mathbb{E}(\pi_{t+1}|\Omega_{it})](\pi_t - \mathbb{E}\pi_t)\} = \mathbb{E}(\{\mathbb{E}[\pi_{t+1} - \mathbb{E}(\pi_{t+1}|\Omega_{it})](\pi_t - \mathbb{E}\pi_t)\}|\Omega_{it}) = \mathbb{E}\{[\mathbb{E}(\pi_{t+1}|\Omega_{it}) - \mathbb{E}(\pi_{t+1}|\Omega_{it})](\pi_t - \mathbb{E}\pi_t)\} = 0$.

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

[14] We also standardize log taxable income in X_{it}.

3.2 Overreaction to the public signal

3.2.1 Regression results without interaction terms

Table 1 reports the regression results estimated by ordinary least squares. The table presents results from seven specifications. The first two columns report regression results without interaction terms (i.e., $b_z = 0$). In specification (1), we regress the forecast error on the public signal and individual fixed effects. The estimated coefficient, $\hat{b}$, is negative and statistically significant at the 1% significance level.[15] The point estimate is -0.021, implying that inflation expectations slightly overreact to the public signal.

Specification (2) allows for a structural break in the bias of inflation expectations. As Figure 1 suggests, inflation rose so rapidly during 2022–2024 that the bias in individuals' inflation expectations, if any, may differ in this period from that in the period before 2022. The bias in inflation expectations during 2022–2024 may also differ from that before 2022 because of the constraint imposed by the survey's response categories, as discussed in the previous section. To address these possibilities, specification (2) includes a structural break in the mean constant term, a, during 2022–2024. When we estimate the regression with this structural break in bias, the estimated coefficient $\hat{b}$ becomes -0.695, which is strongly negative. This overreaction is substantially stronger than the estimates 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. When only data after the 1990s are used, some coefficients become slightly larger: -0.21 in the U.S. Survey of Professional Forecasters and -0.25 in the Michigan Survey of Consumers.[16]

3.2.2 Pervasive overreaction to the public signal

We next discuss the regression results with interaction terms in the estimation equation. In what follows, we denote the elements of the estimated coefficient vector on the interaction terms by $\hat{b}_z^{DC}$, $\hat{b}_z^{FL}$, $\hat{b}_z^{SC}$, and $\hat{b}_z^{RA}$, based on z_{it}^{DC}, z_{it}^{FL}, z_{it}^{SC}, and $\tilde{z}_{it}^{RA}$, respectively.

Overreaction among dynamically consistent respondents Specification (3) in Table 1 regresses forecast errors on y_t and its interaction with z_{it}^{DC}. The estimates indicate that

[15] We use standard errors clustered at the respondent level. We also estimate standard errors clustered at the prefecture level across Japan's 47 prefectures. The results are essentially unchanged.

[16] See Tables C.7 and C.9 in the appendix to Kohlhas and Walther (2021).

$\hat{b} = -0.723$ and $\hat{b}_z^{DC} = 0.107$, and both coefficients are statistically significant at the 1% significance level.[17]

The positive coefficient on the interaction term implies that the response of forecast errors to the public signal becomes weaker when $z_{it}^{DC} = 1$. This suggests that respondents without present bias overreact less to the public signal than those with present bias. 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 estimated response of forecast errors to the public signal remains $-0.616\,(= -0.723 + 0.107)$, indicating statistically significant overreaction of dynamically consistent respondents' inflation expectations.

Overreaction among respondents with forward-looking planning ability Turning to forward-looking planning ability, specification (4) indicates that $\hat{b} = -0.618$ and $\hat{b}_z^{FL} = -0.099$, and both coefficients are statistically significant at the 1% significance level. For respondents with forward-looking planning ability, the implied coefficient is $-0.618 - 0.099 = -0.717$, which is again significantly negative.[18] Thus, even respondents with forward-looking planning ability exhibit statistically significant overreaction to the public signal. Moreover, the negative coefficient on the interaction term implies that respondents with forward-looking planning ability respond more strongly to the public signal than those without such ability.

This finding stands in sharp contrast to the rationality test in Kubota and Fukushige (2016). Focusing on respondents with forward-looking planning ability, they regress income changes on a constant term and expected income changes, and fail to reject the joint hypothesis that the constant term is zero and the coefficient on expected income changes is one. By contrast, our analysis regresses forecast errors on a constant term and the public signal constructed from lagged inflation. The results suggest that expectations are not consistent with the FIRE hypothesis even among respondents who report having forward-looking planning ability.

Overreaction among respondents with self-control For self-control, the results are similar to those for respondents with forward-looking planning ability. In specification (5),

[17] We also estimate overreaction using the elicited present-bias parameter β directly. The results are qualitatively similar: $\hat{b}$ is -0.805, and the coefficient on the interaction term is 0.186. Both are statistically significant.

[18] We also confirm the robustness of this result using an alternative definition of z_{it}^{FL}, under which $z_{it}^{FL} = 1$ only when respondents choose option 1.

the estimated coefficient on y_t is -0.633, while the estimated coefficient on the interaction term is -0.029. Both are statistically significant at least at the 5% significance level. Again, the negative coefficient on the interaction term implies that self-control does not eliminate overreaction to the public signal.[19]

Overreaction under heterogeneity in risk attitude Specification (6) examines how overreaction varies with the degree of risk aversion. The estimates $\hat{b}$ and $\hat{b}_z^{RA}$ are -0.833 and 0.033, respectively. The positive coefficient on the interaction term, $\hat{b}_z^{RA}$, suggests that more risk-averse respondents exhibit slightly less overreaction. However, this effect is small relative to the negative coefficient on y_t, namely $\hat{b}$. Consequently, even when a respondent's risk-aversion measure deviates from the mean by two standard deviations, the implied response of forecast errors to the public signal is $-0.833 + 2 \times 0.033 = -0.767$. Thus, the effect of risk-aversion measure on overreaction is quantitatively minor.[20]

Overreaction among respondents with rational BA Specification (7) reports the results from a regression that includes all of the above BA. The signs of $\hat{b}$ and $\hat{b}_z$ in this specification remain unchanged from those in specifications (3)–(6). Based on the estimates in this column, we can hypothetically consider respondents with rational BA. The response of forecast errors for respondents with rational BA (i.e., $z_{it}^{DC} = z_{it}^{FL} = z_{it}^{SC} = 1$) is given by the sum of the estimated coefficients, $\hat{b} + \hat{b}_z^{DC} + \hat{b}_z^{FL} + \hat{b}_z^{SC}$.[21] This sum is $-0.826\,(= -0.870 + 0.122 - 0.064 - 0.014)$, suggesting that even respondents with rational BA overreact to the public signal. To evaluate the response of forecast errors for respondents with rational BA conservatively, let us consider respondents with extreme risk aversion. Even when rational respondents' risk-aversion measure deviates from the mean by two standard deviations, so that their overreaction is mitigated, the implied response of forecast errors to the public signal is $-0.826 + 2 \times 0.030 = -0.766$. Since we still observe overreaction even in this extreme case, we conclude that respondents with rational BA also pervasively overreact to the public signal.

[19] We also confirm the robustness of this result under a slightly different definition of z_{it}^{SC}, under which $z_{it}^{SC} = 1$ only when respondent i chooses option 5.

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

[21] More precisely, the sum of the estimated coefficients implies the response of forecast errors to the public signal for a rational respondent whose degree of risk aversion is equal to the mean. This is because we standardize z_{it}^{RA}.

In Table 2, we re-estimate (3) with additional control variables: respondents' taxable income, age-group dummies, and a dummy for liquidity-constrained respondents. Even when we include these control variables in the regressions, the estimated coefficients remain similar to those in Table 1 in both sign and magnitude.

3.2.3 Robustness

Table 3 reports a range of robustness checks. In all specifications, the regressions include control variables for respondents' taxable income, age-group dummies, and a dummy for liquidity-constrained respondents.

Structural break Overreaction to the public signal remains pervasively observed even when we assume no structural break. In specifications (1) and (2) of Table 3, we compare the effects of including and excluding a structural break on our results. For comparison, specification (1) in this table reproduces the results of specification (5) in Table 2. The comparison reveals that the structural break strongly influences the estimated value of b. Whereas $\hat{b}$ is -0.894 in specification (1), it is much smaller in specification (2), where $\hat{b} = -0.185$. Nevertheless, aside from the magnitude of $\hat{b}$, each element of $\hat{b}_z$ is not substantially different between specifications (1) and (2). In particular, the estimates $\hat{b}_z^{DC}$ are both positive and similar in magnitude (i.e., 0.120 and 0.147) in the two specifications. Likewise, $\hat{b}_z^{FL}$ are both negative and statistically different from zero (-0.054 and -0.090) in the two specifications. The estimates $\hat{b}_z^{SC}$ differ in sign across the two specifications, but neither is statistically different from zero.

We then ask whether respondents with rational BA continue to overreact to the public signal in specification (2). Their implied response of forecast errors is $-0.185+0.147-0.090+0.003 = -0.125$, with a standard error of 0.019. The corresponding t-statistic indicates that this response is statistically significant. Thus, again, we find pervasive overreaction: even respondents with rational BA exhibit overreaction to the public signal.

Alternative measures of the price index for constructing forecast errors In our benchmark regression, forecast errors are calculated based on core CPI inflation (i.e., the CPI excluding fresh food). However, respondents do not necessarily forecast core CPI inflation when answering the JHPS-CPS, because the survey simply asks by how much they expect consumer prices to change.

To address this issue, we examine robustness using headline CPI (the CPI: All items) inflation and fresh-food inflation.[22] The latter assumes that respondents associate consumer prices with the prices of fresh food rather than with the core CPI. It also reflects the fact that consumers often refer to the prices of frequently purchased goods when forming their inflation expectations.[23]

Specification (3) in Table 3 reports the results when we use headline CPI inflation to construct forecast errors. Our main finding that overreaction is pervasively observed even among respondents with rational BA remains unchanged. In particular, the implied response of forecast errors for respondents with rational BA is $-0.833\,(= -0.904+0.135-0.062-0.002)$ with its standard error of 0.021, and remains significantly negative.

In specification (4), we again confirm overreaction under the assumption that respondents associate consumer prices with fresh-food prices. Under this assumption, the overreaction implied by $\hat{b}$ is substantially smaller than in the baseline case. Nevertheless, the implied response of forecast errors for respondents with rational BA is $-0.125\,(= -0.159 + 0.047 - 0.022 + 0.009)$, with a standard error of 0.008, and remains significantly negative.

Alternative measures of the public signal So far, the public signal we use is inflation in the previous year. However, any variable can serve as a public signal as long as it is included in an individual's information set. Thus, we next conduct robustness checks using an alternative measure of the public signal. Broer and Kohlhas (2024) regress forecast errors on the consensus forecasts of professional forecasters, since such consensus forecasts are readily observable to them. These consensus forecasts are not directly observable to all respondents in our survey, unlike in the case of professional forecasters. Nevertheless, we examine the robustness of our results by regressing forecast errors of inflation on consensus forecasts.

As specification (5) indicates, we continue to find overreaction to this alternative public signal. The estimated coefficient $\hat{b}$ is -0.506. Regarding the coefficients on the interaction terms, we continue to obtain $\hat{b}_z^{DC} > 0$ and $\hat{b}_z^{FL} < 0$. We also find that $\hat{b}_z^{SC}$ remains insignificant. Thus, we continue to observe pervasive overreaction among respondents with rational BA. The response of forecast errors to the public signal is given by $-0.385\,(= -0.506 + 0.221 - 0.099 - 0.001)$, and its standard error is 0.033, again indicating statistical significance.

In specification (6), we observe pervasive overreaction to fresh-food inflation. Unlike

[22] We take the price index for fresh food from *the Annual Report on the Consumer Price Index*.

[23] 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.

specification (4), in which fresh-food inflation is assumed to represent respondents' notion of consumer prices, here we examine how core CPI inflation expectations respond to fresh-food inflation as the public signal. This specification reflects the empirical finding that inflation expectations in Japan are sensitive to changes in the prices of frequently purchased goods (e.g., Coibion and Gorodnichenko (2015b) and Kikuchi and Nakazono (2023)). The estimation results are not substantially different from those in the baseline case, specification (1). The response of forecast errors to fresh-food inflation is somewhat smaller than that in specification (1). The response of forecast errors to the public signal for respondents with rational BA is $-0.181\,(= -0.214 + 0.050 - 0.026 + 0.009)$, with a standard error of 0.008.

Specification (7) uses month-on-month price changes from November to December (i.e., month-on-month inflation), rather than year-on-year price changes from January to December (i.e., year-on-year inflation), as the public signal. This month-on-month inflation is also based on the preliminary CPI published by the Statistics Bureau of the MIAC. Again, we observe overreaction to this alternative public signal, suggesting pervasive overreaction in inflation expectations.

Finally, we use experienced inflation as a quasi-public signal and check robustness. 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:

$$\pi_{it}^{EI} = \sum_{k=1}^{\mathrm{age}_{it}-1} w_{it}(k,\rho)\,\pi_{t-k}, \tag{5}$$

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

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

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, the respondent assigns greater weight to more recent inflation; if ρ is negative, less weight is assigned 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 the weights assigned to public information differ across respondents. As long as experienced inflation, as defined above, is included in respondent i's information set, it is orthogonal to the forecast error of that respondent's inflation expectations. Thus, as before, the coefficients on this signal and its interaction terms with BA should be zero.

Specification (8) reports the results using experienced inflation as the public signal. The estimate of $\hat{b}$ is -0.630. The estimated coefficients on the interaction terms have the same sign as those in specification (1). The response of forecast errors to the public signal for respondents with rational BA is $-0.498\,(= -0.630 + 0.499 - 0.253 - 0.114)$, with a standard error of 0.094, again suggesting pervasive overreaction.

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 full-information rational expectations (FIRE) hypothesis is inconsistent with observed behavior: individuals often overreact to recent events. To account for this pattern, recent behavioral models emphasize that people place excessive weight on observable signals.

Using unique survey data with rich information on behavioral traits, this paper examines how overreactions of inflation expectations to public signals are related to individuals' behavioral attributes. We show that even individuals who appear relatively rational along these dimensions continue to overreact to public signals. The evidence therefore suggests that overreaction is pervasive and cannot be easily explained by observable behavioral heterogeneity. More broadly, our findings support recent macroeconomic models that allow for behavioral overreaction while retaining rationality in other aspects of expectation formation.

A limitation of this paper is that it does not fully identify the mechanism behind overreaction. Our evidence shows that overreaction cannot be explained solely by observable behavioral attributes. Clearly, the possibility of other behavioral explanations for heterogeneous overreaction remains. 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.

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

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

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.

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.

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.

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.

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.

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.*

Notice that the time preference $\beta\delta^\tau$ is obtained from (1), whereas δ^τ is obtained from (2), where τ corresponds to 7 days. Time preference is identified by the choice at which the respondent switches from option A to option B. For example, the list of choices for (1) in the 2024 survey is shown in Table 4. 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 (2), we derive the range of δ^τ from the nine choices and use its midpoint as the measure of time preference obtained from (2). In an extreme case, the implied time preference may exceed 1, which implies a negative discount rate.[24] In this case, we recode any time preference greater than unity as unity. In another extreme case, inconsistent choices across options may prevent us from uniquely identifying a range for time preference. In such cases, we drop those respondents from the sample.

[24] In the example in Table 4, 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 $0 < \beta \leq 1$.[25] Overall, 77.7% 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 $\beta = 1$.

There are some remarks on changes in the 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 2 days from today, or 9 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 eight 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 4 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 83.87% to 86.60%, the corresponding fraction from 2011 onward ranges from 70.14% to 74.78%. We therefore take this shift into account.

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

A.2 Measuring risk attitudes

Examples of the choices are shown in Table 5. 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 attitude is identified by the choice at which the respondent switches from Option "A" to Option "B." For example, when 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 risk-neutral, since the amount paid equals the expected return of the risky lottery.

There are some remarks on risk attitudes. First, there are changes in the questions. 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. Notice that when the expected return is low, the expected loss from buying the lottery is also low. In this case, respondents may be more willing to take the risk of buying the lottery.[26]

Second, for robustness, Cramer et al. (2002) and Hanaoka et al. (2018) also use the Arrow-Pratt measure of absolute risk aversion.[27] This measure is a monotonically increasing but nonlinear function of the reservation price.

[26] In fact, the mean of z_{it}^{RA} during 2004–2009 is 0.699, which is lower than the full-sample mean of 0.785 and implies a lower degree of risk aversion.

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

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

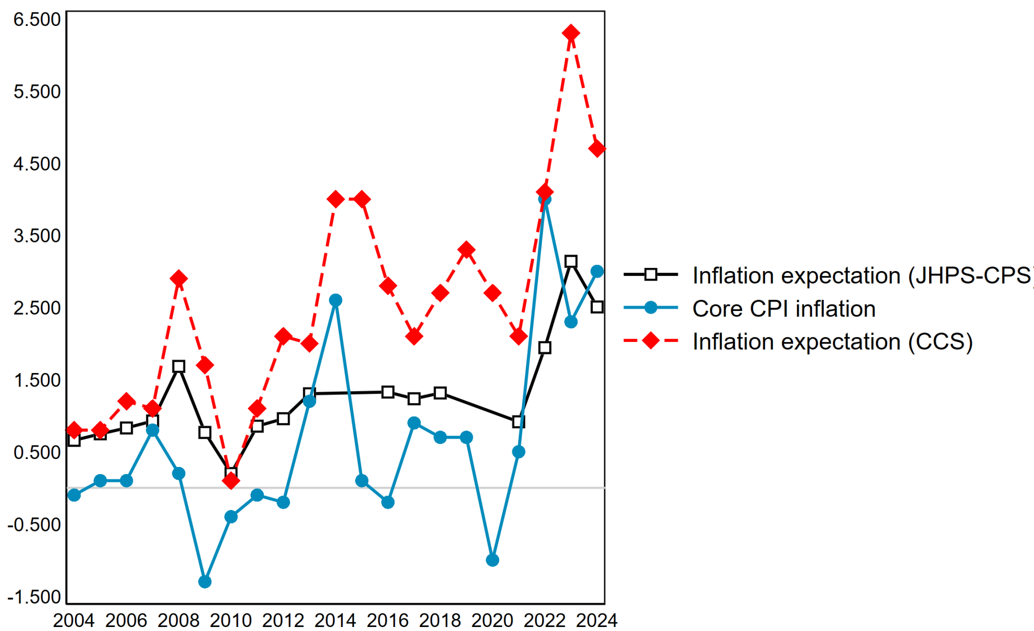

Notes: The figure displays the mean of respondents' inflation expectations and historical inflation data from 2004 to 2024. The line with squares represents inflation expectations, and the line with circles represents inflation. For comparison, the figure also plots inflation expectations calculated from the Consumer Confidence Survey. The x-axis shows years, and the y-axis is measured in percent.

Figure 2: Histogram of inflation expectations in 2023

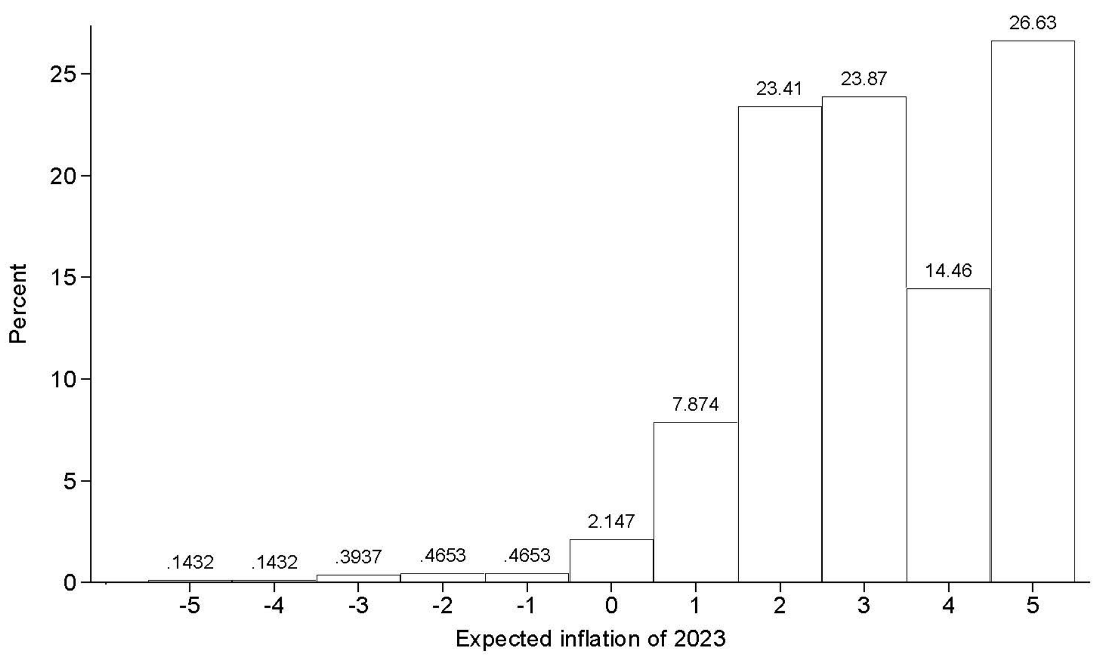

Notes: The figure displays the distribution of inflation expectations in 2023. The x-axis shows inflation expectations, and the y-axis is measured in percent.

Figure 3: Forecast errors of mean inflation expectations

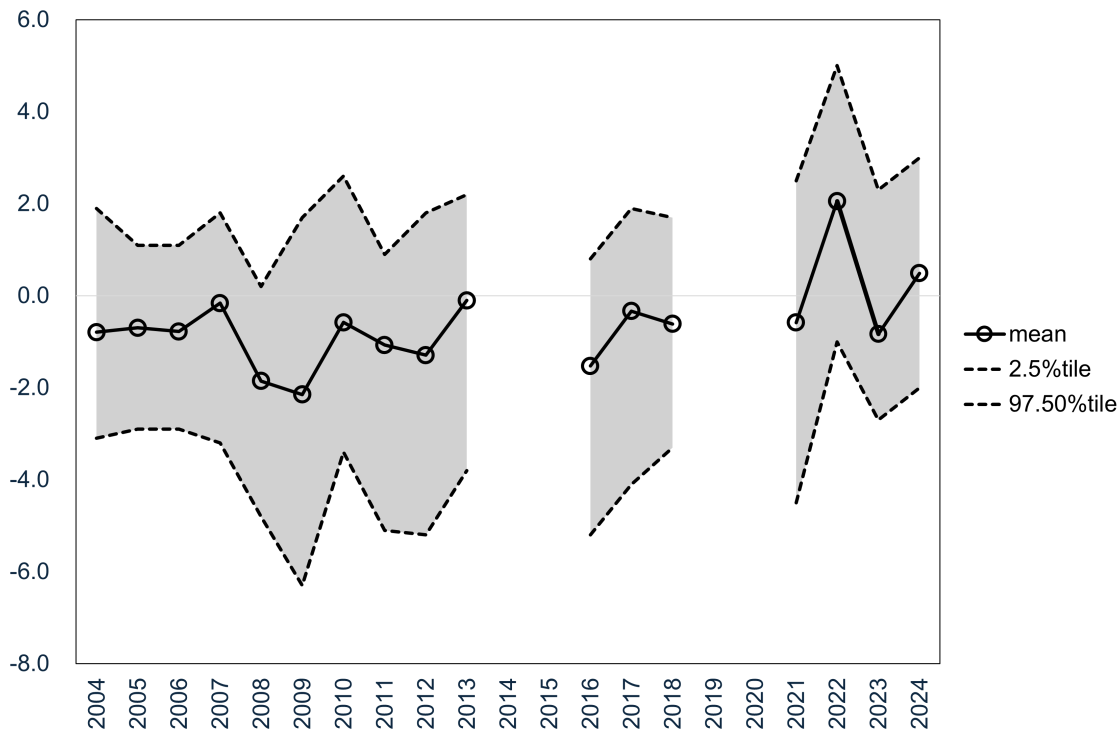

Notes: The figure displays the forecast error of mean inflation expectations, $\pi_t - \sum_i \pi^e_{it}$, from 2004 to 2024. The shaded area represents the cross-sectional variation in forecast errors, based on the interval between the 2.5th and 97.5th percentiles. The x-axis shows years, and the y-axis is measured in percent.

Table 1: Estimated responses to past inflation and bias in forecast errors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
y_t	-0.021***	-0.695***	-0.723***	-0.618***	-0.633***	-0.833***	-0.870***
	(0.007)	(0.008)	(0.014)	(0.010)	(0.011)	(0.010)	(0.019)
$z_{it}^{DC} \times y_t$			0.107***				0.122***
			(0.015)				(0.016)
$z_{it}^{FL} \times y_t$				-0.099***			-0.064***
				(0.014)			(0.016)
$z_{it}^{SC} \times y_t$					-0.029**		-0.014
					(0.014)		(0.015)
$\tilde{z}_{it}^{RA} \times y_t$						0.033***	0.030***
						(0.007)	(0.008)
z_{it}^{DC}			-0.158***				-0.268***
			(0.022)				(0.033)
z_{it}^{FL}				0.267***			0.217***
				(0.028)			(0.035)
z_{it}^{SC}					0.166***		0.163***
					(0.026)		(0.034)
$\tilde{z}_{it}^{RA}$						0.071***	0.082***
						(0.014)	(0.017)
Intercept	-0.710***	-1.079***	-0.848***	-1.094***	-1.102***	-0.957***	-0.800***
	(0.001)	(0.004)	(0.018)	(0.010)	(0.014)	(0.007)	(0.034)
Dummy for		3.220***	2.947***	2.973***	2.979***	3.379***	3.208***
high-inflation period		(0.027)	(0.028)	(0.030)	(0.030)	(0.033)	(0.035)
Observations	53,943	53,943	37,083	28,845	28,847	24,195	17,827
R-squared	0.259	0.428	0.442	0.487	0.485	0.543	0.593

Notes: The dependent variable is the forecast error, defined as the difference between core CPI inflation and respondents' inflation expectations. Standard errors are reported in parentheses below the coefficients. Coefficients are statistically significant at the *10%, **5%, and ***1% significance levels. Variables with a tilde are standardized to have mean zero and unit variance. Individual fixed effects are included in all regressions.

Table 2: Robustness analysis: Estimation with control variables

	(1)	(2)	(3)	(4)	(5)
y_t	-0.768***	-0.703***	-0.716***	-0.833***	-0.894***
	(0.016)	(0.012)	(0.013)	(0.012)	(0.021)
$z_{it}^{DC} \times y_t$	0.100***				0.120***
	(0.016)				(0.018)
$z_{it}^{FL} \times y_t$		-0.074***			-0.054***
		(0.016)			(0.018)
$z_{it}^{SC} \times y_t$			-0.021		-0.007
			(0.016)		(0.017)
$\tilde{z}_{it}^{RA} \times y_t$				0.034***	0.037***
				(0.008)	(0.009)
z_{it}^{DC}	-0.108***				-0.205***
	(0.024)				(0.037)
z_{it}^{FL}		0.159***			0.174***
		(0.030)			(0.037)
z_{it}^{SC}			0.126***		0.135***
			(0.028)		(0.038)
$\tilde{z}_{it}^{RA}$				0.052***	0.027
				(0.016)	(0.019)
Intercept	-0.389***	-0.397***	-0.408***	-0.576***	-0.442***
	(0.033)	(0.036)	(0.037)	(0.044)	(0.065)
Dummy for	2.902***	2.978***	2.980***	3.151***	3.162***
high-inflation period	(0.036)	(0.038)	(0.038)	(0.042)	(0.040)
Observations	29,243	22,444	22,451	18,887	13,836
R-squared	0.480	0.543	0.542	0.581	0.630

Notes: The regressions include respondents' log taxable income, age-group dummies, and a dummy for liquidity-constrained respondents as control variables. See the notes to Table 1 for further details.

Table 3: Robustness analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
y_t	-0.894***	-0.185***	-0.904***	-0.159***	-0.506***	-0.214***	-0.384***	-0.630***
	(0.021)	(0.022)	(0.023)	(0.009)	(0.039)	(0.009)	(0.020)	(0.105)
$z_{it}^{DC} \times y_t$	0.120***	0.147***	0.135***	0.047***	0.221***	0.050***	0.165***	0.499***
	(0.018)	(0.021)	(0.020)	(0.008)	(0.037)	(0.009)	(0.018)	(0.086)
$z_{it}^{FL} \times y_t$	-0.054***	-0.090***	-0.062***	-0.022***	-0.099***	-0.026***	-0.071***	-0.253***
	(0.018)	(0.020)	(0.019)	(0.008)	(0.034)	(0.008)	(0.016)	(0.078)
$z_{it}^{SC} \times y_t$	-0.007	0.003	-0.002	0.009	-0.001	0.009	-0.018	-0.114
	(0.017)	(0.019)	(0.019)	(0.008)	(0.031)	(0.008)	(0.015)	(0.071)
$\tilde{z}_{it}^{RA} \times y_t$	0.037***	0.047***	0.037***	0.037***	-0.006	0.039***	0.167***	0.150***
	(0.009)	(0.009)	(0.009)	(0.004)	(0.015)	(0.004)	(0.007)	(0.035)
z_{it}^{DC}	-0.205***	-0.360***	-0.285***	-0.231***	-0.448***	-0.258***	-0.263***	-0.418***
	(0.037)	(0.047)	(0.041)	(0.048)	(0.063)	(0.049)	(0.041)	(0.049)
z_{it}^{FL}	0.174***	0.348***	0.221***	0.163***	0.283***	0.192***	0.178***	0.260***
	(0.037)	(0.049)	(0.042)	(0.047)	(0.059)	(0.048)	(0.042)	(0.049)
z_{it}^{SC}	0.135***	0.234***	0.159***	0.083*	0.194***	0.114**	0.163***	0.208***
	(0.038)	(0.047)	(0.042)	(0.048)	(0.054)	(0.049)	(0.041)	(0.047)
$\tilde{z}_{it}^{RA}$	0.027	-0.103***	-0.003	-0.252***	0.024	-0.295***	-0.187***	-0.060**
	(0.019)	(0.023)	(0.021)	(0.024)	(0.028)	(0.024)	(0.020)	(0.024)
Intercept	-0.442***	0.358***	0.074	-1.309***	0.209**	0.012	-0.570***	-0.195**
	(0.065)	(0.074)	(0.070)	(0.077)	(0.086)	(0.079)	(0.074)	(0.077)
Dummy for	3.162***		3.235***	6.835***	1.536***	1.410***	2.193***	1.581***
high-inflation period	(0.040)		(0.042)	(0.031)	(0.039)	(0.032)	(0.043)	(0.073)
Observations	13,836	13,836	13,836	13,836	13,836	13,836	13,836	11,653
R-squared	0.630	0.405	0.625	0.905	0.475	0.557	0.531	0.461

Notes: Specification (1) replicates specification (5) in Table 2 to highlight the effect of the structural break. In specification (2), we remove the assumption of a structural break during 2022–2024. In specifications (3) and (4), we use alternative measures of forecast errors and public signals. Specifically, we use headline CPI inflation in specification (3) and fresh-food inflation in specification (4), instead of core CPI inflation in the baseline estimation. In specifications (5)–(8), we use alternative measures of the public signal y_t: specifications (5), (6), (7), and (8) use consensus forecasts, fresh-food inflation, month-on-month inflation, and experienced inflation, respectively. In all regressions, we include respondents' log taxable income, age-group dummies, and a dummy for liquidity-constrained respondents as control variables. See the notes to Table 1 for further details.

Table 4: Examples of choices for time preferences in the JHPS-CPS

Choices	Option "A" Receive today	Option "B" Receive 7 days from today	Which one do you prefer? 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 5: Examples of choices for risk aversion in the JHPS-CPS

		Which one do you prefer?	
Choices	Price of the "speed lottery" ticket	Option "A" (buy the "speed lottery" ticket)	Option "B" (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 ☐