Tohoku University Research Center for Policy Design Discussion Paper

TUPD-2026-013

People Prefer Zero Inflation: Evidence from Conjoint Analysis on Inflation–Unemployment Trade-offs

Ritsu Yano
Graduate School of International Management, Yokohama City University

Yoshiyuki Nakazono
Graduate School of International Management, Yokohama City University
Graduate School of Economics and Management, Tohoku University

Jun Takahashi
Graduate School of International Management, Yokohama City University

Sep 2026

TUPD Discussion Papers can be downloaded from:

https://www2.econ.tohoku.ac.jp/~PDesign/dp.html

People Prefer Zero Inflation: Evidence from Conjoint Analysis on Inflation–Unemployment Trade-offs*

Ritsu Yano† Yoshiyuki Nakazono‡ Jun Takahashi§

September 11, 2026

Abstract

Many central banks, including the Bank of Japan, define price stability as 2% inflation. Do households agree? We answer this question with a conjoint experiment in which Japanese respondents chose between hypothetical economies that differed in their inflation and unemployment rates. We find that households prefer zero inflation. An economy with 2% inflation is chosen significantly less often than one with 0% inflation. On average, respondents are indifferent between 0% and −2% inflation, although men and younger respondents prefer 0% to deflation. We also find that households weigh unemployment more heavily than inflation: for a one-percentage-point fall in inflation, they accept a rise in unemployment of only about 0.65 percentage points. The results point to a gap between the inflation rate households prefer and the 2% that the Bank of Japan targets. This gap raises the possibility that the weak anchoring of Japanese households' inflation expectations at 2% partly reflects their preference for inflation closer to zero.

JEL Classification: E31; E52; E58

Keywords: inflation; unemployment; monetary policy; inflation target; conjoint experiment

*Nakazono acknowledges financial support from JSPS KAKENHI (Grant Number: JP25H00544; JP26K00321; and JP26K22091). Takahashi acknowledges financial support from JST SPRING, Japan (Grant Number: JPMJSP2179).

†Yokohama City University

‡Yokohama City University; Tohoku University

§Yokohama City University

1 Introduction

Two percent has become the global benchmark for price stability. The Federal Reserve made a 2% objective explicit in 2012, the European Central Bank adopted a symmetric 2% target in 2021, and the Bank of Japan introduced a 2% "price stability target" in 2013. The case for a small positive target rests on room to cut interest rates before they hit the zero lower bound (Summers, 1991; Blanchard et al., 2010; Kiley and Roberts, 2017), upward bias in measured inflation (Boskin et al., 1998; Shiratsuka, 1999), and downward nominal wage rigidity (Akerlof et al., 1996). Theory alone does not deliver the number two: models of money holding and price dispersion imply an optimal rate at or below zero (Schmitt-Grohé and Uribe, 2010), New Keynesian models with a zero lower bound imply a positive rate below 2% (Coibion et al., 2012), and the lower-bound argument has also been used to motivate 4% (Blanchard et al., 2010; Ball, 2014). The 2% target is thus a policy judgment rather than a number uniquely pinned down by theory. Whether the public shares that judgment matters: most households do not know the target and their expectations are far from anchored at it (Coibion et al., 2020). This raises the possibility that expectations anchor at a target only if households accept it. Surveys in the United States and Europe suggest that the public dislikes inflation more than the consensus presumes (Shiller, 1997; Di Tella et al., 2001; Dräger et al., 2022; Stantcheva, 2024; Afrouzi et al., 2024). Yet in recent decades, aggregate inflation in these economies has fallen below zero only briefly, so for most current respondents prolonged zero inflation is not a lived macroeconomic experience.[1]

This paper examines how Japanese households evaluate inflation, and how they trade it off against unemployment, with a conjoint experiment. Conjoint experiments, developed in marketing research and now a standard tool in political science (Hainmueller et al., 2014; Abramson et al., 2022), present respondents with pairs of profiles whose attributes are randomly assigned and ask them to choose one. Following Binetti et al. (2024), our profiles are hypothetical economies described by two attributes, the inflation rate and the unemployment rate. We targeted 1,000 respondents from an online panel, sampled with quotas on gender and age, and analyze the 942 who completed the survey and

[1] U.S. consumer price inflation was negative from March to October 2009, and euro area inflation was negative from December 2014 to March 2015 and from August to December 2020.

answered every question we use; the rest left the survey partway through or declined to answer. Each completed ten choices, giving 18,840 profile evaluations. Because the attribute levels are randomly and independently assigned, the design identifies the causal effect of each level on the probability that an economy is chosen without imposing a model of how respondents decide, and because both attributes vary at once, respondents must trade one off against the other rather than simply state that they dislike inflation.

We do three things with these data. First, we ask which inflation rate households prefer. We estimate the population average marginal component effect (pAMCE) of de la Cuesta et al. (2022), which averages the effect of each inflation level over the historical distributions of inflation and unemployment in Japan rather than over the uniform distribution of the experiment, and we rank inflation levels with the average component preference (ACP) of Ganter (2023), which measures how often a level is chosen head-to-head against the other levels. Second, we ask whether that preference is shared across the population. Because the ACP is not defined relative to a baseline level, differences in the ACP across subgroups reflect differences in preferences, which Leeper et al. (2020) show differences in AMCEs need not, and we re-estimate it by gender and by age. Third, we ask how households trade inflation off against unemployment. We estimate the marginal rate of substitution (MRS) between the two, following Binetti et al. (2024), and we repeat every step of the analysis on their U.S. data to see which of our results are specific to Japan.

There are three findings. First, Japanese households prefer zero inflation. The ACP falls monotonically with inflation across the positive range, and respondents prefer 0% inflation to 2% inflation at any conventional significance level. The pAMCE says the same in causal terms: moving an economy from 0% to 2% inflation lowers the probability that it is chosen, and the effect is more than three times as large when the hypothetical economies are weighted by their historical frequencies as when they are weighted uniformly. The reason is that under uniform weights the competing economy often has extreme inflation, and against an alternative with 20% inflation it hardly matters whether one's own economy has 0% or 2%; under historical weights the competing economy usually has inflation between 0% and 4%, so the difference between 0% and 2% is pivotal. Deflation is a different matter. On average, respondents are indifferent between 0% and -2% inflation: the ACP of 0% is slightly

higher than that of -2%, but not significantly so. Households therefore do not value positive inflation for its own sake. The level they favor is at or just below zero, not the 2% that the Bank of Japan targets.

Second, the average masks systematic heterogeneity, and the heterogeneity concerns deflation rather than the preference for zero over 2%. Men prefer 0% to -2% inflation, whereas women are indifferent between the two. Older respondents dislike high inflation much more than younger respondents do, whereas younger respondents dislike deflation and prefer 0% to -2%. Every group, however, prefers 0% to 2%. The age pattern is consistent with the view that lifetime inflation experiences shape how people think about inflation, which Malmendier and Nagel (2016) document for inflation expectations, and with the life-cycle interpretation that Afrouzi et al. (2024) offer for the finding that older people prefer lower inflation. Our survey does not ask about the mechanism behind the gender difference, and we treat its explanation as a question for future research.

Third, households weigh unemployment more heavily than inflation. The estimated marginal rate of substitution is 0.65: to lower inflation by one percentage point, respondents accept a rise in unemployment of only about 0.65 percentage points. The estimate lies below one in every demographic subgroup. This contrasts with the U.S. data of Binetti et al. (2024), in which the preferences over inflation are far stronger than the preferences over unemployment. The reweighting, by contrast, works the same way in both countries: weighting the U.S. profiles by the U.S. historical distributions roughly doubles the estimated aversion to 2% and 4% inflation relative to 0%, so the conventional uniform design understates the public's aversion to moderate inflation in the United States as well as in Japan.

This paper contributes to three strands of the literature. The first strand asks why people dislike inflation, how much they dislike it, and how that dislike varies across groups. Shiller (1997) showed that ordinary people dislike inflation for reasons that economists rarely emphasize, Di Tella et al. (2001) use happiness surveys from Europe and the United States to show that inflation lowers reported well-being, although less than unemployment does, and Stantcheva (2024) shows that Americans dislike inflation intensely and mostly because they believe it erodes their real income. Afrouzi et al. (2024) find that the inflation rate U.S. households prefer lies well below the Federal

Reserve's 2% target, Pfajfar and Winkler (2025) elicit households' preferences over inflation and the trade-offs monetary policy faces, Dräger et al. (2022) show that households with the same inflation expectations differ in whether they regard that level as appropriate and that these preferences vary with income, education, risk attitudes, and trust in the central bank, and D'Acunto et al. (2021a) show that the prices households see shape their inflation expectations. Binetti et al. (2024) use a conjoint experiment to show that U.S. respondents place a larger weight on inflation than on unemployment. We add evidence from Japan and identify preferences over inflation from repeated choices between economies that require respondents to trade inflation off against unemployment. Most closely related, Niizeki (2026) directly elicits Japanese households' preferred inflation rate in isolation and finds a mean of -0.55% and a median of 0%. Rather than asking respondents to report an ideal inflation rate directly, we infer preferences from repeated choices between economies that force respondents to trade inflation off against unemployment.

The second strand is the debate over the level of the inflation target, which has not settled on two. Models of the opportunity cost of holding money and price dispersion place the optimal rate at or below zero (Lucas, 2000; Schmitt-Grohé and Uribe, 2010), while models with an occasionally binding zero lower bound generally imply a positive target (Billi, 2011; Coibion et al., 2012). Other work points to higher targets when the natural rate of interest is low or the costs of trend inflation are limited (Kiley and Roberts, 2017; Andrade et al., 2019; Blanco, 2021; Nakamura et al., 2018), and estimates based on micro price data and calibrated models span a range around the conventional 2% target (Adam and Weber, 2019; Adam et al., 2022; Cho et al., 2025). The 2% target thus lies within a range of defensible values rather than at a point uniquely pinned down by theory. Existing work also studies whether information about the inflation target anchors household expectations and whether households regard a given inflation rate as desirable. What remains unclear is whether acceptance of the target itself conditions the extent to which expectations anchor at that target. Our evidence speaks to the first step in this link by showing whether households themselves regard 2% inflation as preferable to alternative inflation rates.

The third strand concerns the design and interpretation of conjoint experiments. On design, the central issue is the distribution of profiles over which effects are averaged. The standard estimand,

the average marginal component effect (AMCE) of Hainmueller et al. (2014), averages the effect of changing one attribute over the distribution of all other attributes in both profiles, and de la Cuesta et al. (2022) show that its value, and even its sign, can change with that distribution. Most applications, including Binetti et al. (2024), average over the uniform distribution used to randomize the profiles, which in our setting treats an economy with 24% inflation as exactly as likely as one with 0% inflation. We instead estimate the pAMCE of de la Cuesta et al. (2022), which averages over a target profile distribution chosen on substantive grounds, in our case the historical distributions of Japanese inflation and unemployment,[2] and we show that the choice of distribution changes the answer in Japan and in the United States alike. On interpretation, the AMCE answers a causal question but does not say which level respondents prefer or how strongly: Ganter (2023) shows that in a forced-choice design it mixes the preference for one level over another with the distribution of that attribute in the competing profile, and Abramson et al. (2022) show that it can favor a level most respondents rank below the alternative. We therefore use Ganter's ACP, which compares each level head-to-head with every other level of the same attribute. Because the ACP is not measured relative to a baseline, differences in the ACP across subgroups describe differences in preferences, which Leeper et al. (2020) show differences in AMCEs need not, and because every ACP is a choice probability minus one half, it allows a descriptive comparison of preferences over inflation with preferences over unemployment within the ranges of levels used. Applications of the pAMCE and the ACP to macroeconomic attitudes remain limited, and we use both to distinguish the effects of profile distributions from underlying preferences.

The rest of the paper proceeds as follows. Section 2 describes the experiment and the historical distributions used to weight the hypothetical economies. Section 3 explains how we estimate the pAMCE, the ACP, and the MRS, and why we use each. Section 4 reports the results and compares them with the United States. Section 5 concludes.

[2]Bansak et al. (2023) make a related point from the other direction: when the goal is the effect of an attribute on a candidate's expected vote share in elections that mirror the experimental design, the AMCE, with the expectation taken over the target distribution of voters and candidate attributes, is the quantity of interest.

2 Data

2.1 Conjoint Experiment in Japan

We ran an online conjoint experiment in Japan. The survey was administered by GMO Research & AI, Inc. to members of its online panel on July 23, 2026. We set a target of 1,000 respondents aged 20 to 79, with quotas on gender and age group so that the sample matches the Japanese population on those two dimensions. Not all of the respondents who started the survey completed it. Some left the survey before finishing the conjoint tasks, and some chose "prefer not to answer" for a demographic question that we use in the analysis. We exclude these respondents. After these exclusions, 942 respondents remain, and they constitute our sample. Table 1 compares the 942 respondents with the population. The sample matches the population closely on gender and age, which were the quota targets. It over-represents college graduates and slightly under-represents the lowest income group, as online panels typically do. About 14% of respondents hold a mortgage, and about 35% report having been unemployed at some point.

The survey had four parts. First, respondents reported their gender, age, prefecture of residence, education, income, and whether they held a mortgage. Second, we defined the inflation rate and the unemployment rate so that all respondents interpreted the two terms the same way. The inflation rate was defined as the annual change in the price of a basket of goods and services typically consumed by Japanese households, with a numerical example, and the unemployment rate as the share of the labor force that is not working but actively looking for work. We then told respondents that both rates had been about 3% in Japan in the previous year, and on the next page we asked them to recall the unemployment rate. Third, respondents completed the conjoint choice tasks. Finally, they reported whether they had ever been unemployed.[3]

In each choice task, respondents saw two hypothetical economies, labeled Option A and Option B, each defined by an inflation rate and an unemployment rate, and chose the one they would rather live in for the next year. The choice was forced: respondents could not declare indifference. Each respondent completed ten tasks. We drew the inflation rate uniformly from 14 levels between -2%

[3] Appendix B reproduces the full questionnaire.

and 24% in steps of 2 percentage points, and the unemployment rate uniformly from 10 levels between 1.0% and 5.5% in steps of 0.5 percentage points. The range of inflation levels covers Japan's postwar experience, from the deflation of the 2000s to the peak of the first oil shock, and it includes both the Bank of Japan's 2% target and the zero inflation that Japan experienced for much of the past three decades. The range of unemployment levels covers Japan's experience since 1971, in which the unemployment rate has never exceeded 5.5%. Because we drew the two rates independently, the two attributes are orthogonal, and the levels of one attribute carry no information about the other. To avoid framing effects, we randomized across respondents the order in which the two attributes appeared. Because the levels were randomly assigned, the design identifies how each level changes the choice probability without assuming any particular model of how respondents decide.

We drew the levels uniformly rather than from the historical distribution for two reasons that de la Cuesta et al. (2022) discuss. A uniform design delivers observations at every level, including rare ones such as -2% or 20% inflation, and so allows us to estimate the effect of every level with reasonable precision. It also allows us to compute the AMCE under any target distribution after the fact, so that the uniform AMCE and the population AMCE can be compared on the same data. The historical distribution therefore enters only at the estimation stage, and only for the pAMCE. The uniform AMCE and the ACP, which we report alongside the pAMCE, are computed under the uniform distribution of the experiment, and the ACP is designed not to depend on how often each competing level appears in the experiment, although it remains defined relative to the set of levels included and the weights used. Section 3 makes the role of each distribution explicit.

2.2 Historical Distributions of Inflation and Unemployment

The central methodological choice in this paper is the distribution over which we average the effect of each attribute. The uniform distribution used to randomize the profiles is convenient for identification, but it describes no economy that has ever existed. Under uniform weights, a respondent evaluating 2% inflation is, in effect, evaluating it against competing economies whose inflation rate is drawn uniformly from -2% to 24%, and whose average inflation rate is 11%. Japanese households do not face that menu. The pAMCE of de la Cuesta et al. (2022) instead averages the effect over a

target profile distribution that the researcher chooses on substantive grounds; it may be a real-world distribution or a counterfactual distribution of theoretical interest. We choose the distribution of inflation and unemployment that Japanese households actually face, and we take that distribution from history.

We compute the marginal distribution of each variable from monthly Japanese data between January 1971 and May 2025: the year-on-year change in the consumer price index (CPI) for inflation and the unemployment rate from the Labour Force Survey. We round each monthly observation to the nearest level used in the experiment and compute the share of months at each level. We start the sample in 1971 so that it includes the inflation of the 1970s, which is the only period in which Japanese inflation reached the upper levels of our design, as well as the deflation of 1999–2012 and the inflation after 2022. Table 2 reports the two distributions.

The historical distributions are far from uniform. Inflation was at roughly 0% in 45% of the months and at roughly 2% in a further 24%, so the two levels at the center of the policy debate account for about 70% of Japan's experience. Inflation of 12% or more occurred in about 4% of the months, and 20% inflation did not occur at all after rounding. Unemployment is concentrated between 2% and 3%, which together account for 57% of the months, whereas the 1% level at the bottom of our design occurred in only 2% of the months. Weighting by these distributions therefore shifts the comparison from a world in which most competing economies have high inflation to one in which most competing economies have inflation between 0% and 4%. Section 4.1 shows that this shift changes the estimated aversion to moderate inflation by a factor of more than three, and Appendix A shows that the same reweighting changes the U.S. estimates of Binetti et al. (2024) in the same direction. We weight by the marginal distribution of each variable; because the experiment assigned the two attributes independently, this amounts to weighting each combination of inflation and unemployment by the product of the two marginals, which approximates the historical joint distribution by treating the two variables as independent.[4]

[4]Using the historical joint distribution instead would additionally reflect the correlation between inflation and unemployment in Japanese data. We leave this extension for future work.

3 Estimation

We estimate three quantities. The average marginal component effect (AMCE; Hainmueller et al., 2014) measures how each level of an attribute changes the probability that a profile is chosen, averaged over a distribution of the other attributes. Because the standard AMCE averages over the uniform distribution used in the experiment, we estimate the population AMCE (pAMCE; de la Cuesta et al., 2022), which averages over the historical distribution described in Section 2.2. The average component preference (ACP; Ganter, 2023) measures how strongly respondents prefer each level relative to the other levels of the same attribute, and unlike the AMCE it does not depend on the experimental frequency of ties or of the competing levels of the focal attribute, although it is still defined relative to the included level set and the chosen weights.[5] The marginal rate of substitution measures the relative weight that respondents place on inflation and unemployment. To be explicit about the role of the two distributions: the historical distribution of Table 2 is used only to compute the pAMCE. The uniform AMCE, the ACP, and the MRS are all estimated on the experimental data as randomized, that is, under the uniform distribution of levels.

Throughout, i indexes respondents, k indexes choice tasks, $j \in \{1, 2\}$ indexes the two profiles in a task, and l indexes attributes. The outcome Y_{ijk} equals one if respondent i chooses profile j in task k and zero otherwise.

3.1 Population Average Marginal Component Effect (pAMCE)

The AMCE of moving attribute l from level t_0 to level t_1 is (Hainmueller et al., 2014; de la Cuesta et al., 2022)

$$\tau_l(t_1, t_0) = \sum_{(t_{ijk,-l}, t_{i,-j,k}) \in \mathcal{T}} \mathbb{E}\left[Y_{ijk}(t_1, t_{ijk,-l}, t_{i,-j,k}) - Y_{ijk}(t_0, t_{ijk,-l}, t_{i,-j,k})\right] \mathrm{P}(t_{ijk,-l}, t_{i,-j,k}), \quad (1)$$

where $Y_{ijk}(\cdot)$ is the potential outcome, $t_{ijk,-l}$ collects the levels of the other attributes of profile j, $t_{i,-j,k}$ collects the levels of all attributes of the competing profile in the same task, P is the joint

[5] We compute the pAMCE and the ACP with the replication materials of de la Cuesta et al. (2020) and Ganter (2021).

distribution of those levels, and $\mathcal{T}$ is its support. The expectation runs over respondents and tasks.

The AMCE thus averages the effect of the focal attribute over the distribution P of everything else: the other attribute of the same profile and both attributes of the competing profile. The choice of P is part of the definition of the estimand, not a detail of estimation. Most conjoint studies set P to the uniform distribution that assigned the levels in the experiment; we call the result the uniform AMCE (uAMCE). The uAMCE gives every combination of inflation and unemployment the same weight, however rare that combination is in practice. de la Cuesta et al. (2022) show that the AMCE under one distribution need not carry over to another, and they propose the pAMCE, which sets P to the distribution in the target population. In our application the target profile distribution is the distribution of inflation and unemployment that Japanese households have faced historically. That choice is ours: de la Cuesta et al. (2022) ask that the target distribution be chosen according to the substantive interests of the study, and they allow either a real-world or a counterfactual target. We set P to the product of the historical marginal distributions of inflation and unemployment in Japan reported in Table 2.

Following the model-based approach of de la Cuesta et al. (2022), we fit a linear probability model to the experimental data that includes the main effects of both attributes and their interactions, including the interactions of each attribute across the two profiles in a pair (Egami and Imai, 2019), and we average the fitted effects using the historical distributions as weights. The cross-profile interactions are what allow the effect of an economy's own inflation rate to depend on the inflation rate of the economy it is compared with, which is the channel through which the profile distribution matters. Uniform randomization gives every profile pair a positive probability of being drawn, and with 18,840 profile evaluations every level and every pair of levels that enters the model is observed many times. Not every full profile pair is observed, so the pAMCE rests on the model's assumption that three-way and higher-order interactions are zero. The resulting pAMCE tells us how each level of inflation or unemployment changes the choice probability in economies that Japanese households could plausibly face. The cost of reweighting is precision: because the historical distribution puts most of its mass on a few levels, the pAMCE is estimated from fewer effective observations than the uAMCE, and its confidence intervals are wider.

3.2 Average Component Preference (ACP)

The AMCE, under either weighting, is a causal effect on a choice probability. It is not a measure of preference. Ganter (2023) shows that in a forced-choice conjoint experiment the AMCE of level t_l relative to t_0 depends on the distribution of the same attribute in the competing profile: it aggregates the preference for t_l over t_0 with the preferences for t_l and t_0 over every other level, weighted by how often those other levels appear. Two groups with identical preferences can therefore have different AMCEs, and a change in the design changes the AMCE even when preferences do not change. Abramson et al. (2022) make a related point: because the AMCE aggregates over respondents, it reflects the intensity of preferences as well as their direction, and a level can have a positive AMCE even when most respondents rank it below the baseline. Ganter (2023) proposes the ACP, which isolates the preference between two levels by comparing profiles that differ only in the level of the focal attribute. We therefore use the ACP to measure how strongly respondents prefer each level of inflation and unemployment.

Let T_{ijkl} be the level of attribute l in profile j and $T_{i[-j]kl}$ the level of the same attribute in the competing profile. The preference for level t_l over level t'_l is

$$p_{t_l,t'_l} \equiv \mathbb{E}\left[Y_{ijk}(T_{ijkl} = t_l, T_{i[-j]kl} = t'_l)\right] - 0.5, \tag{2}$$

where the two profiles share the same levels of all other attributes, and the expectation runs over respondents and profiles. Thus p_{t_l,t'_l} measures how far the probability of choosing t_l over t'_l departs from indifference, 0.5. Because the two profiles differ only in the focal attribute, p_{t_l,t'_l} is a head-to-head comparison of the two levels and does not depend on the distribution of the other levels.

The ACP of level t_l averages these pairwise preferences over all other levels (Ganter, 2023):

$$\begin{aligned}\pi_l(t_l; w_{t_l}) &\equiv \mathbb{E}\left[Y_{ijk}(T_{ijkl} = t_l, T_{i[-j]kl} \neq t_l) \,\middle|\, w_{t_l}\right] - 0.5 \\ &= \sum_{t'_l \in \mathcal{T}_l,\, t'_l \neq t_l} w_{t_l}(t'_l)\, p_{t_l,t'_l},\end{aligned} \tag{3}$$

where $\mathcal{T}_l$ is the set of levels of attribute l and w_{t_l} is a vector of weights that sum to one. Following

Ganter (2023), we weight all other levels equally: $w_{t_l}(t'_l) = 1/(|\mathcal{T}_l| - 1)$ for every $t'_l \neq t_l$. The uniform randomization of the experiment identifies every pairwise preference p_{t_l,t'_l}, and the ACP combines them with these fixed weights; the historical distribution plays no role here, which is the point of the ACP as a measure of preference rather than of the effect in a particular environment. The ACP therefore does not depend on how often each competing level appears in the experiment, but it remains defined relative to the set of levels included in the design and to the weights w_{t_l}. The ACP is positive when respondents, on average, prefer level t_l to the other levels and negative when they dislike it. Because every ACP is a choice probability minus one half, it lies between -0.5 and 0.5 for every level of every attribute, and Ganter (2023) presents it as comparable within and across attributes; cross-attribute comparisons are nevertheless descriptive and relative to the range of levels included for each attribute. Because it is not defined relative to a reference category, it also lets us compare preferences across subgroups without the baseline problem that Leeper et al. (2020) identify for the AMCE. We exploit both properties in Section 4.2, where we rank inflation levels, compare the preferences of men and women and of different age groups, and compare the strength of preferences over inflation with that over unemployment. To test whether two ACPs differ, we compute the difference and its standard error.

3.3 Marginal Rate of Substitution between Inflation and Unemployment

The pAMCE and the ACP describe preferences level by level. To summarize the trade-off between the two attributes in a single number, we follow Binetti et al. (2024) and estimate the marginal rate of substitution from the regression

$$Y_{ijk} = \psi_i + \beta_\pi \Delta\pi_{ijk} + \beta_u \Delta u_{ijk} + \epsilon_{ijk}, \tag{4}$$

where ψ_i is a respondent fixed effect, $\Delta\pi_{ijk}$ is the inflation rate of profile j minus that of the competing profile in the same task, Δu_{ijk} is the corresponding difference in unemployment rates, and ϵ_{ijk} is an error term. The regression treats the two attributes as continuous and imposes that a one-percentage-point difference has the same effect at every level.

If preferences are linear, the marginal rate of substitution is the ratio of the two coefficients:

$$\lambda = \beta_\pi / \beta_u. \tag{5}$$

The parameter λ is the rise in unemployment, in percentage points, that respondents accept in exchange for a one-percentage-point fall in inflation. If $\lambda > 1$, respondents dislike inflation more than unemployment; if $\lambda < 1$, they dislike unemployment more. We compute the confidence interval for λ by the delta method.

4 Results

This section reports the pAMCE, the ACP, and the marginal rate of substitution between inflation and unemployment, together with a comparison with the United States. The three estimands answer different questions. The pAMCE tells us how each level of inflation or unemployment changes the probability that an economy is chosen, in the economies Japanese households actually face. The ACP tells us which levels respondents prefer and how strongly. The MRS summarizes the trade-off in a single number. One rule governs how we read the estimates. An AMCE, whether uniform or population-weighted, is a weighted average of effects over a particular profile distribution and is measured relative to a particular baseline level (Hainmueller et al., 2014; de la Cuesta et al., 2022). Its magnitude is therefore meaningful only within an attribute and within a design: we compare the uAMCE and the pAMCE of a given inflation level with each other, but we do not compare the AMCE of an inflation level with the AMCE of an unemployment level, or an AMCE from one figure with an AMCE from another. Comparisons across attributes and across groups are reserved for the ACP, which is on a common scale, and for the MRS.

4.1 Population Average Marginal Component Effect (pAMCE)

This subsection shows that weighting by the historical distributions sharpens the aversion to positive inflation but leaves the evaluation of deflation unchanged. Figure 1 plots the uAMCE and the pAMCE

of each inflation level relative to 0% inflation. Under uniform weights (uAMCE), an economy with positive inflation is chosen less often than one with 0% inflation, and the effect grows almost linearly with the inflation rate. Under historical weights (pAMCE), the picture changes in two ways. First, the gap between 0% and moderate inflation widens sharply: the pAMCE of 2% inflation is more than three times as large in absolute value as the uAMCE of the same level, and for every level from 2% to 16% the pAMCE is more negative than the uAMCE. Second, the pAMCE flattens at high inflation: above about 10%, additional inflation lowers the choice probability little further. This is what one expects when the competing economies are drawn from history. Against a competitor that usually has inflation between 0% and 4%, an economy with 2% inflation is already at a disadvantage, and an economy with 12% inflation is rejected almost regardless of its unemployment rate. Under either weighting, an economy with -2% inflation is chosen about as often as one with 0% inflation.

Figure 2 plots the same estimates for unemployment relative to 1% unemployment. Under uniform weights, every unemployment rate above 1% lowers the probability that an economy is chosen. Under historical weights, unemployment rates of 2.5% or below no longer differ significantly from 1%, whereas for rates above 2.5% the pAMCE and the uAMCE of each level are close to each other. For unemployment, then, reweighting changes the estimates little, which is what one expects given that the historical distribution of unemployment is much closer to the uniform distribution than the historical distribution of inflation is (Table 2). We do not compare the size of the unemployment effects with the size of the inflation effects here; that comparison requires a common scale and is made with the ACP and the MRS below.

4.2 Average Component Preference (ACP) for Inflation and Unemployment

This subsection shows that households prefer zero inflation. Figure 3 plots the ACP of each level of inflation and unemployment, and Table 3 tests the differences between the ACPs of the lowest inflation levels. Within the positive range, the ACP falls monotonically as inflation rises: each 2-percentage-point step from 0% to 8% lowers the ACP by about 0.05, and the ACP turns negative at about 12%, the level at which an economy is chosen less often than not against the alternatives. The ranking at the bottom of the range is the one that matters for the policy debate. Respondents prefer

0% inflation to 2% ($p < 0.001$), and the difference in ACPs, 0.067, is larger than the difference between 2% and 4% (0.051). Respondents also prefer -2% to 2% ($p < 0.001$). The ACP of 0% exceeds that of -2%, but the difference, 0.013, is not significant. On average, then, respondents rank zero inflation and mild deflation together at the top, and they rank the central bank's 2% target clearly below both. Unemployment shows the same monotone pattern: the ACP falls as the unemployment rate rises, and it turns negative at about 3.5%.

Figure 3 also allows a descriptive comparison across attributes, because the ACP is on the same bounded scale for both. The ACPs of inflation range from about 0.33 at 0% to about -0.38 at 24%, whereas the ACPs of unemployment range from about 0.10 at 1% to about -0.10 at 5.5%. The wider range for inflation partly reflects the wider range of levels in the design, and we return to the relative weight on the two attributes when we estimate the MRS.

Some groups, however, prefer 0% inflation to -2%. To see this, we re-estimate the ACPs by gender and by age group. Because the ACP is not defined relative to a baseline, the subgroup estimates can be compared directly (Leeper et al., 2020). Figure 4 and Table 4 report the estimates by gender. Men prefer 0% inflation to -2% ($p < 0.05$), whereas women are indifferent between the two, and the preference for -2% differs significantly between men and women. Both groups prefer 0% to 2% and 2% to 4%.

Figure 5 and Table 5 report the estimates by age group. Older respondents dislike very high inflation much more than younger respondents do, and they show a clearer preference for low inflation: for respondents in the oldest group, the ACP of -2% is slightly above that of 0%, and both are well above that of 2%. Younger respondents, by contrast, dislike deflation: their ACP of -2% is significantly below their ACP of 0% ($p < 0.05$), and they do not distinguish significantly between -2% and 2%. Younger respondents nevertheless prefer 0% to 2% ($p < 0.05$), so the preference for zero inflation holds in every age group; what differs across groups is the evaluation of deflation.

4.3 Marginal Rate of Substitution between Inflation and Unemployment

This subsection shows that households weigh unemployment more heavily than inflation. Table 6 reports the regression. Both coefficients are negative: an economy is chosen less often as either

its inflation rate or its unemployment rate rises relative to the competing economy. The estimated marginal rate of substitution is 0.65, with a 95% confidence interval of $[0.578, 0.738]$, so it lies well below one. To lower inflation by one percentage point, respondents accept a rise in unemployment of only about 0.65 percentage points. Under linear preferences, then, Japanese households place more weight on unemployment than on inflation. Figure 6 plots the estimate for the full sample and for demographic subgroups defined by gender, age, income, education, region, mortgage status, and unemployment experience. In every subgroup the estimate lies below one, and its confidence interval excludes one.

4.4 Comparison with the United States

Japan's history of low inflation and low unemployment makes it natural to ask whether the preferences we document are specific to Japan. Appendix A repeats our analysis on the U.S. conjoint data of Binetti et al. (2024), in which 2,264 respondents each evaluated five pairs of hypothetical economies with inflation rates from 0% to 16% and unemployment rates from 2% to 16%. Binetti et al. (2024) chose those levels from the range of monthly U.S. outcomes since the 1970s, but they average over the levels uniformly. We reweight their profiles by the historical U.S. distributions of inflation and unemployment and compute the ACP. Two results stand out.

First, the effect of reweighting is the same in both countries. Figure A.1 shows that the pAMCE of every positive inflation level from 2% to 14% is more negative than the uAMCE, and that historical weighting roughly doubles the estimated aversion to 2% and 4% inflation relative to 0%. As in Japan, the reweighting matters most at the moderate inflation rates that are common in history and that central banks target, because under historical weights the competing economy usually has low inflation and the difference between 0% and 2% becomes pivotal. Figure A.2 shows that for unemployment, as in Japan, the pAMCE and the uAMCE of each level are similar. Binetti et al. (2024) report the uniform AMCE and conclude that Americans dislike inflation; the pAMCE says that they dislike moderate inflation relative to zero considerably more than the uniform estimate suggests. In both countries, then, the conventional uniform design understates the public's aversion to inflation at the levels that central banks target.

Second, Americans also prefer the lowest inflation rate offered. In the U.S. data the ACP of inflation falls monotonically from 0% upward (Figure A.3), just as in Japan, and the ACP of 0% exceeds that of 2% by a wide margin. The U.S. design does not include deflation, so the U.S. data cannot tell us whether Americans, like Japanese respondents, would treat -2% and 0% alike.

5 Conclusion

Many central banks define price stability as 2% inflation. This paper examines whether households in Japan, the economy that has tried hardest to raise inflation to 2%, share that definition. In an online conjoint experiment, 942 respondents chose between hypothetical economies that differed in their inflation and unemployment rates. We estimate three quantities. The population average marginal component effect weights hypothetical economies by the historical distributions of inflation and unemployment in Japan rather than by the uniform distribution of the experiment. The average component preference measures how strongly respondents prefer each level, while the marginal rate of substitution captures the trade-off between inflation and unemployment, following Binetti et al. (2024).

Our main finding is that households prefer zero inflation. An economy with 2% inflation is chosen significantly less often than one with 0% inflation, and the gap is several times larger once the hypothetical economies are weighted by their historical frequencies. On average, respondents are indifferent between 0% and -2% inflation, so they do not value positive inflation for its own sake. The preference for zero over 2% holds for men and women and in every age group; what differs across groups is the evaluation of deflation, which men and younger respondents dislike and women and older respondents do not. We also find that households weigh unemployment more heavily than inflation: the marginal rate of substitution is about 0.65 and lies below one in every subgroup.

The methodological lesson is that the answer depends on the estimand. The uniform AMCE that conjoint studies conventionally report averages over a set of economies most of which have never existed. de la Cuesta et al. (2022) recommend that the averaging distribution reflect substantive knowledge of the target profiles rather than the uniform distribution used for randomization. Weighting by

history is our implementation of that advice, and it more than triples the estimated aversion to 2% inflation relative to zero in Japan and roughly doubles it in the U.S. data of Binetti et al. (2024). Measuring preferences with the ACP of Ganter (2023) rather than with the AMCE lets us rank inflation levels and compare groups without the dependence on the design and on the baseline category that the AMCE carries. Researchers who use conjoint experiments to study macroeconomic attitudes should report the estimand that matches the question, and for questions about the economies that households actually face, that estimand is the pAMCE evaluated at the historical distribution. Other target distributions, including counterfactual ones motivated by theory, answer different questions (de la Cuesta et al., 2022).

What do our findings imply for the 2% target? They do not settle what the target should be, nor do they imply that it should be zero. A positive target can give central banks more room to cut nominal interest rates, accommodate measurement bias in inflation, and ease adjustment when nominal wages resist downward cuts. Our experiment does not tell respondents about these benefits and ask whether they justify higher inflation. It does, however, reveal a gap between the inflation rate households prefer and the 2% rate the Bank of Japan targets. That gap matters for two reasons. First, households may anchor their expectations less firmly to a target they regard as too high. Japanese households' weak anchoring at 2% may therefore partly reflect their preference for inflation closer to zero.[6] Second, raising inflation from 0% to 2% moves inflation away from the level households prefer. The case for doing so therefore depends on what households gain in return, such as stronger employment or more room for monetary policy. Our results do not show whether these benefits justify the 2% target, but they identify a trade-off that the broader debate should confront.

[6] Kikuchi and Nakazono (2023) show that Japanese households' long-term inflation expectations are not anchored at 2%.

References

Abramson, Scott F., Korhan Koçak, and Asya Magazinnik (2022). "What Do We Learn about Voter Preferences from Conjoint Experiments?" *American Journal of Political Science* 66(4), 1008–1020.

Adam, Klaus, and Henning Weber (2019). "Optimal Trend Inflation." *American Economic Review* 109(2), 702–737.

Adam, Klaus, Erwan Gautier, Sergio Santoro, and Henning Weber (2022). "The Case for a Positive Euro Area Inflation Target: Evidence from France, Germany and Italy." *Journal of Monetary Economics* 132, 140–153.

Afrouzi, Hassan, Alexander Dietrich, Kristian Myrseth, Romanos Priftis, and Raphael Schoenle (2024). "Inflation Preferences." *NBER Working Paper* 32379, National Bureau of Economic Research.

Akerlof, George A., William T. Dickens, and George L. Perry (1996). "The Macroeconomics of Low Inflation." *Brookings Papers on Economic Activity* 1996(1), 1–76.

Andrade, Philippe, Jordi Galí, Hervé Le Bihan, and Julien Matheron (2019). "The Optimal Inflation Target and the Natural Rate of Interest." *Brookings Papers on Economic Activity* 2019(2), 173–230.

Ball, Laurence (2014). "The Case for a Long-Run Inflation Target of Four Percent." *IMF Working Paper* 14/92, International Monetary Fund.

Boskin, Michael J., Ellen L. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson (1998). "Consumer Prices, the Consumer Price Index, and the Cost of Living." *Journal of Economic Perspectives* 12(1), 3–26.

Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins, and Teppei Yamamoto (2023). "Using Conjoint

Experiments to Analyze Election Outcomes: The Essential Role of the Average Marginal Component Effect." *Political Analysis* 31(4), 500–518.

Billi, Roberto M. (2011). "Optimal Inflation for the US Economy." *American Economic Journal: Macroeconomics* 3(3), 29–52.

Binetti, Alberto, Francesco Nuzzi, and Stefanie Stantcheva (2024). "People's Understanding of Inflation." *Journal of Monetary Economics* 148, 103652.

Blanchard, Olivier, Giovanni Dell'Ariccia, and Paolo Mauro (2010). "Rethinking Macroeconomic Policy." *Journal of Money, Credit and Banking* 42(s1), 199–215.

Blanco, Andrés (2021). "Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound." *American Economic Journal: Macroeconomics* 13(3), 108–141.

Cho, Daeha, Jung Hyun Kim, Kwang Hwan Kim, and Suk Joon Kim (2025). "Optimal Trend Inflation in an Open Economy." *Journal of Monetary Economics* 149, 103675.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland (2012). "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?" *Review of Economic Studies* 79(4), 1371–1406.

Coibion, Olivier, Yuriy Gorodnichenko, Saten Kumar, and Mathieu Pedemonte (2020). "Inflation Expectations as a Policy Tool?" *Journal of International Economics* 124, 103297.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber (2021a). "Exposure to Grocery Prices and Inflation Expectations." *Journal of Political Economy* 129(5), 1615–1639.

de la Cuesta, Brandon, Naoki Egami, and Kosuke Imai (2020). "Replication Data for: Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." *Harvard Dataverse*, V1. doi: 10.7910/DVN/HVY5GR.

de la Cuesta, Brandon, Naoki Egami, and Kosuke Imai (2022). "Improving the External Validity of Conjoint Analysis: The Essential Role of Profile Distribution." *Political Analysis* 30(1), 19–45.

Di Tella, Rafael, Robert J. MacCulloch, and Andrew J. Oswald (2001). "Preferences over Inflation and Unemployment: Evidence from Surveys of Happiness." *American Economic Review* 91(1), 335–341.

Dräger, Lena, Michael J. Lamla, and Damjan Pfajfar (2022). "The Hidden Heterogeneity of Inflation and Interest Rate Expectations: The Role of Preferences." *CESifo Working Paper* 9637, CESifo.

Egami, Naoki, and Kosuke Imai (2019). "Causal Interaction in Factorial Experiments: Application to Conjoint Analysis." *Journal of the American Statistical Association* 114(526), 529–540.

Ganter, Flavien (2021). "Replication Data for: Identification of Preferences in Forced-Choice Conjoint Experiments: Reassessing the Quantity of Interest." *Harvard Dataverse*, V1. doi: 10.7910/DVN/3RNU1Q.

Ganter, Flavien (2023). "Identification of Preferences in Forced-Choice Conjoint Experiments: Reassessing the Quantity of Interest." *Political Analysis* 31(1), 98–112.

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto (2014). "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1), 1–30.

Kiley, Michael T., and John M. Roberts (2017). "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* 2017(1), 317–396.

Kikuchi, Junichi, and Yoshiyuki Nakazono (2023). "The Formation of Inflation Expectations: Microdata Evidence from Japan." *Journal of Money, Credit and Banking* 55(6), 1609–1632.

Leeper, Thomas J., Sara B. Hobolt, and James Tilley (2020). "Measuring Subgroup Preferences in Conjoint Experiments." *Political Analysis* 28(2), 207–221.

Lucas, Robert E., Jr. (2000). "Inflation and Welfare." *Econometrica* 68(2), 247–274.

Malmendier, Ulrike, and Stefan Nagel (2016). "Learning from Inflation Experiences." *Quarterly Journal of Economics* 131(1), 53–87.

Nakamura, Emi, Jón Steinsson, Patrick Sun, and Daniel Villar (2018). "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation." *Quarterly Journal of Economics* 133(4), 1933–1980.

Niizeki, Takeshi (2026). "Why Do People Dislike Inflation? The Role of Missing Debt Erosion Channel." *Applied Economics Letters* 33(9), 1437–1442.

Pfajfar, Damjan, and Fabian Winkler (2025). "Households' Preferences over Inflation and Monetary Policy Tradeoffs." *Working Paper* 25-12, Federal Reserve Bank of Cleveland.

Schmitt-Grohé, Stephanie, and Martín Uribe (2010). "The Optimal Rate of Inflation." In *Handbook of Monetary Economics*, Vol. 3, edited by Benjamin M. Friedman and Michael Woodford, 653–722. Amsterdam: Elsevier.

Shiratsuka, Shigenori (1999). "Measurement Errors in the Japanese Consumer Price Index." *Monetary and Economic Studies* 17(3), 69–102.

Shiller, Robert J. (1997). "Why Do People Dislike Inflation?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer, 13–70. Chicago: University of Chicago Press.

Stantcheva, Stefanie (2024). "Why Do We Dislike Inflation?" *Brookings Papers on Economic Activity* Spring 2024, 1–46.

Summers, Lawrence (1991). "How Should Long-Term Monetary Policy Be Determined?" *Journal of Money, Credit and Banking* 23(3), 625–631.

Table 1: Sample Characteristics

	Survey	Japan population
Targeted quotas		
Male	0.51	0.49
Female	0.49	0.51
20-29 years old	0.12	0.14
30-39 years old	0.13	0.16
40-49 years old	0.17	0.17
50-59 years old	0.20	0.20
60-69 years old	0.22	0.18
70-79 years old	0.16	0.16
Non-targeted characteristics		
0-4,999,999 yen	0.46	0.54
5,000,000-9,999,999 yen	0.31	0.32
10,000,000 yen +	0.10	0.14
I don't know or prefer not to answer	0.13	
College graduate	0.46	0.24

Notes: The table compares the survey sample with the Japanese population. The sample consists of the 942 respondents who remain, out of a target of 1,000, after excluding respondents who left the survey partway through or declined to answer a question used in the analysis, as described in Section 2. Population statistics come from the 2020 Population Census (except education) and the 2025 Comprehensive Survey of Living Conditions.

Table 2: Historical Marginal Distributions of Inflation and Unemployment in Japan

Inflation Rate	**Probability**	**Unemployment Rate**	**Probability**
-2%	0.033	1.0%	0.023
0%	0.450	1.5%	0.051
2%	0.239	2.0%	0.182
4%	0.132	2.5%	0.259
6%	0.041	3.0%	0.125
8%	0.039	3.5%	0.083
10%	0.027	4.0%	0.077
12%	0.006	4.5%	0.096
14%	0.011	5.0%	0.072
16%	0.003	5.5%	0.033
18%	0.002		
20%	0.000		
22%	0.009		
24%	0.009		

Notes: The table reports the marginal distributions of inflation and unemployment in Japan used to weight the hypothetical economies in the pAMCE. We compute the probabilities from monthly data on the year-on-year change in the CPI and the unemployment rate from January 1971 to May 2025, rounding each observation to the nearest level used in the experiment. The uniform distribution used to randomize the profiles assigns probability 0.071 to each inflation level and 0.100 to each unemployment level.

Table 3: Differences in Average Component Preferences (ACP) for Inflation Rates

	Comparison Level			
Inflation Rate	-2%	0%	2%	4%
-2%	—	-0.013	0.054^{***}	0.105^{***}
0%		—	0.067^{***}	0.118^{***}
2%			—	0.051^{***}
4%				—

Notes: Each cell reports the row ACP minus the column ACP. A positive entry means that respondents prefer the row level to the column level. Sample: 942 respondents, 18,840 profiles.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 4: Differences in Average Component Preferences (ACP) for Inflation Rates by Gender

	Comparison Level (Column)			
Inflation Rate (Row)	-2%	0%	2%	4%
Panel A: Male ($N = 480$)				
-2%	—	-0.055^{*}	0.023	0.084^{***}
0%		—	0.078^{***}	0.139^{***}
2%			—	0.060^{**}
4%				—
Panel B: Female ($N = 462$)				
-2%	—	0.032	0.079^{***}	0.120^{***}
0%		—	0.047^{*}	0.088^{***}
2%			—	0.041^{*}
4%				—

Notes: Each cell reports the row ACP minus the column ACP. A positive entry means that respondents prefer the row level to the column level. N is the number of respondents.

$^{***}\ p < 0.001$, $^{**}\ p < 0.01$, $^{*}\ p < 0.05$

Table 5: Differences in Average Component Preferences (ACP) for Inflation Rates by Age Group

	Comparison Level (Column)			
Inflation Rate (Row)	-2%	0%	2%	4%
Panel A: Elder, 60–79 years old ($N = 353$)				
-2%	—	0.012	0.074^{**}	0.130^{***}
0%		—	0.061^{**}	0.118^{***}
2%			—	0.056^{*}
4%				—
Panel B: Middle, 40–59 years old ($N = 354$)				
-2%	—	0.008	0.059^{*}	0.139^{***}
0%		—	0.051^{*}	0.130^{***}
2%			—	0.079^{**}
4%				—
Panel C: Young, 20–39 years old ($N = 235$)				
-2%	—	-0.071^{*}	0.001	0.039
0%		—	0.072^{*}	0.110^{**}
2%			—	0.039
4%				—

Notes: Each cell reports the row ACP minus the column ACP. A positive entry means that respondents prefer the row level to the column level. N is the number of respondents.
$^{***}\ p < 0.001$, $^{**}\ p < 0.01$, $^{*}\ p < 0.05$

Table 6: Estimation Results for the Marginal Rate of Substitution

	Y_{ijk}
$\Delta\pi_{ijk}$	-0.0266^{***}
	(0.0003)
Δu_{ijk}	-0.0409^{***}
	(0.0014)
MRS (λ)	0.650
95% CI	$[0.578, 0.738]$
Observations	18,840

Notes: Estimates of equation (4). The regression includes respondent fixed effects. Standard errors in parentheses. The 95% confidence interval for λ is computed by the delta method. Sample: 942 respondents, 9,420 tasks, 18,840 profiles.

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

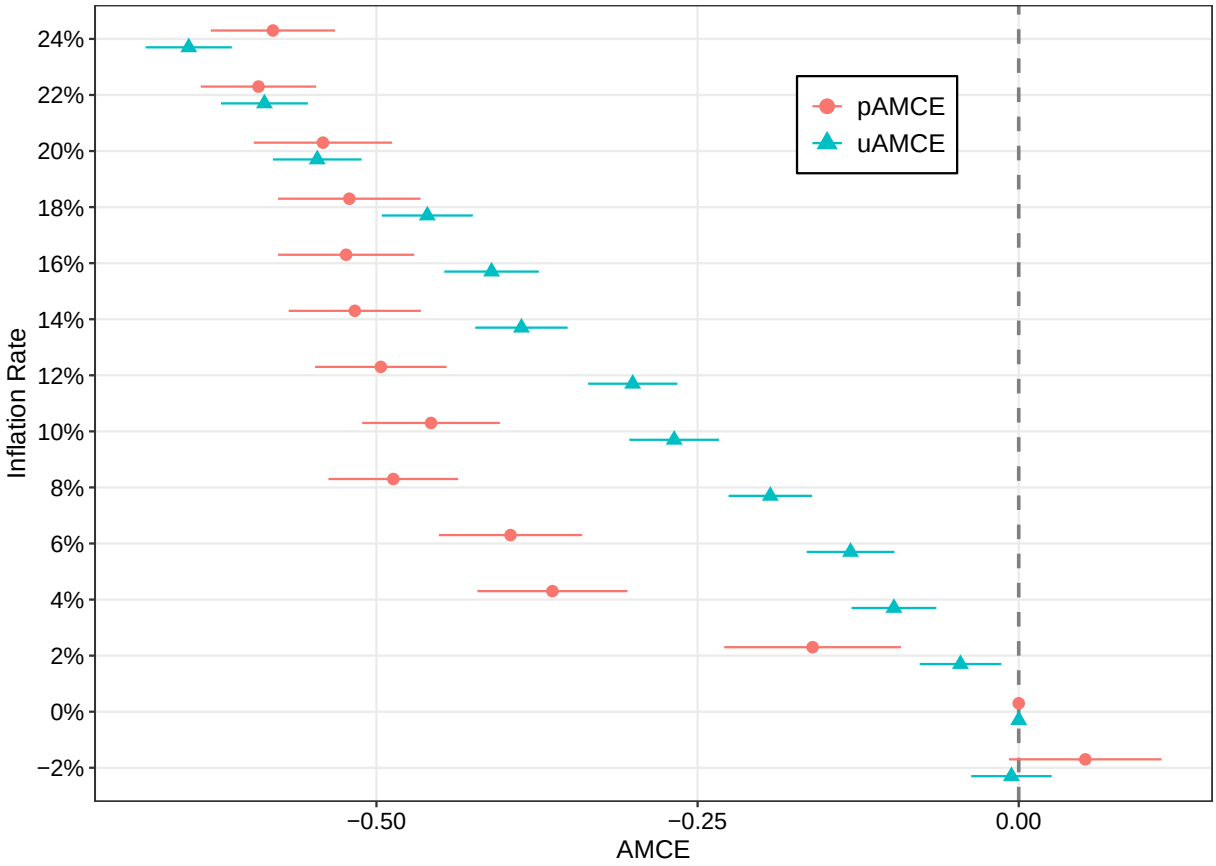

Figure 1: uAMCE and pAMCE of inflation rates in Japan. Each dot is the uAMCE or pAMCE of the corresponding inflation level relative to 0% inflation, estimated with the model-based method of de la Cuesta et al. (2022). The uAMCE averages over the uniform distribution used in the experiment; the pAMCE averages over the historical distributions of Table 2 (January 1971 to May 2025). Horizontal bars are 95% confidence intervals. Sample: 942 respondents, 9,420 tasks, 18,840 profiles.

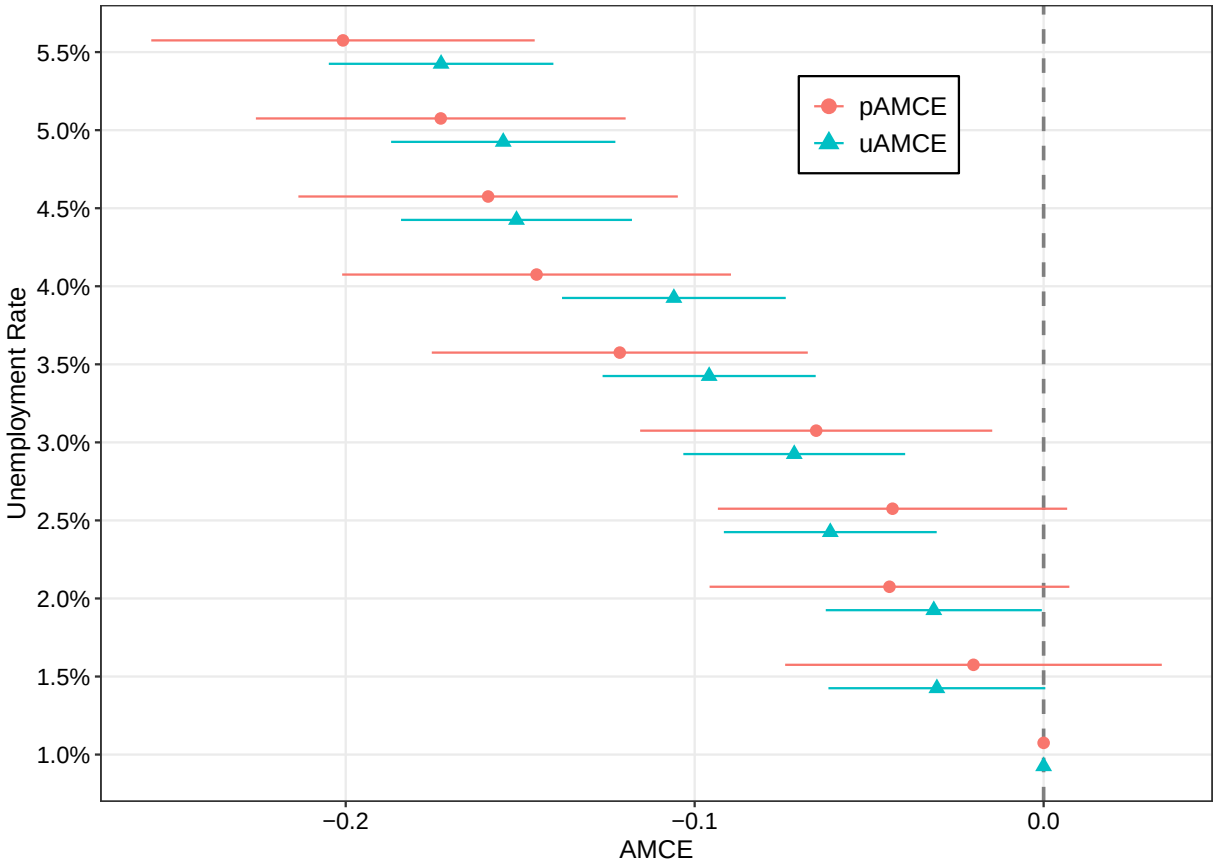

Figure 2: uAMCE and pAMCE of unemployment rates in Japan. Each dot is the uAMCE or pAMCE of the corresponding unemployment level relative to 1% unemployment, estimated with the model-based method of de la Cuesta et al. (2022). The uAMCE averages over the uniform distribution used in the experiment; the pAMCE averages over the historical distributions of Table 2 (January 1971 to May 2025). Horizontal bars are 95% confidence intervals. Sample: 942 respondents, 9,420 tasks, 18,840 profiles.

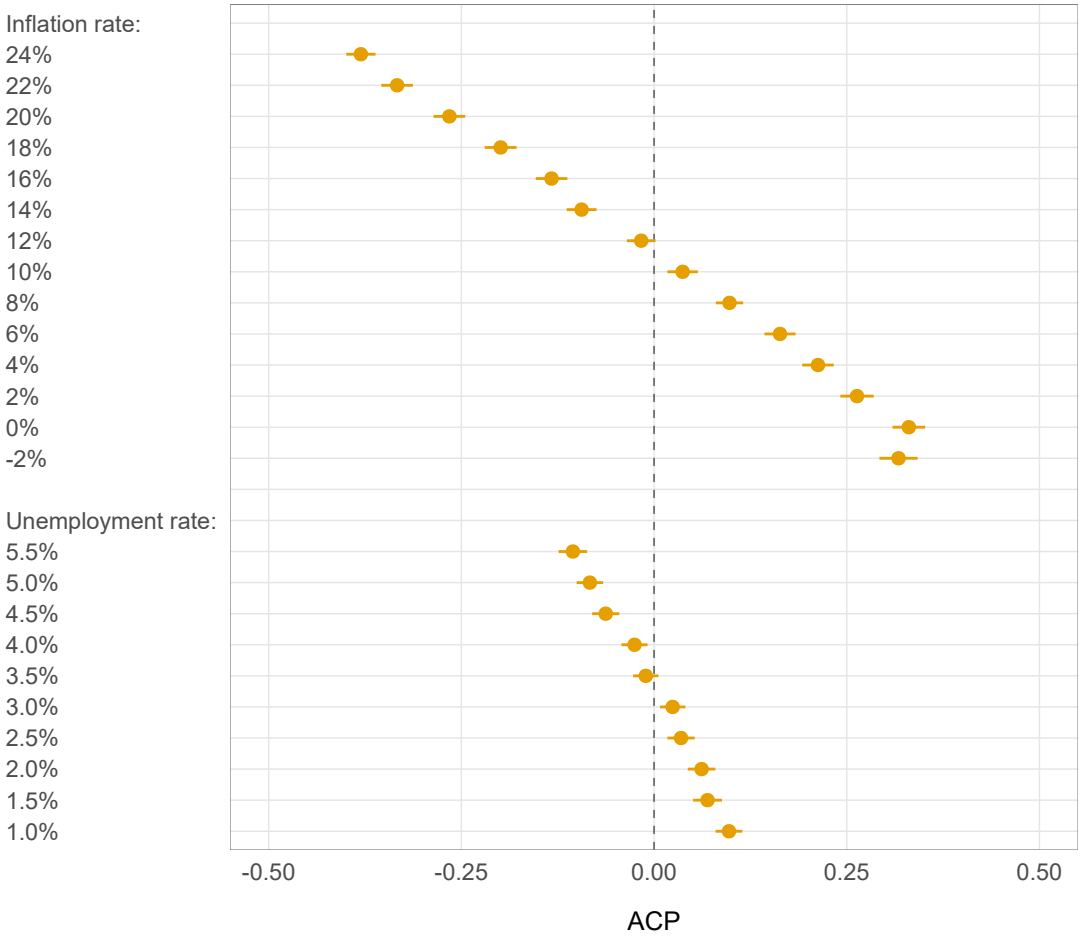

Figure 3: ACPs of inflation rates and unemployment rates in Japan. Each dot is the ACP of the corresponding inflation or unemployment level, estimated with the method of Ganter (2023) with equal weights on the other levels of the same attribute. A positive ACP means that the level is chosen more often than not against the other levels of the same attribute. Horizontal bars are 95% confidence intervals. Sample: 942 respondents, 9,420 tasks, 18,840 profiles.

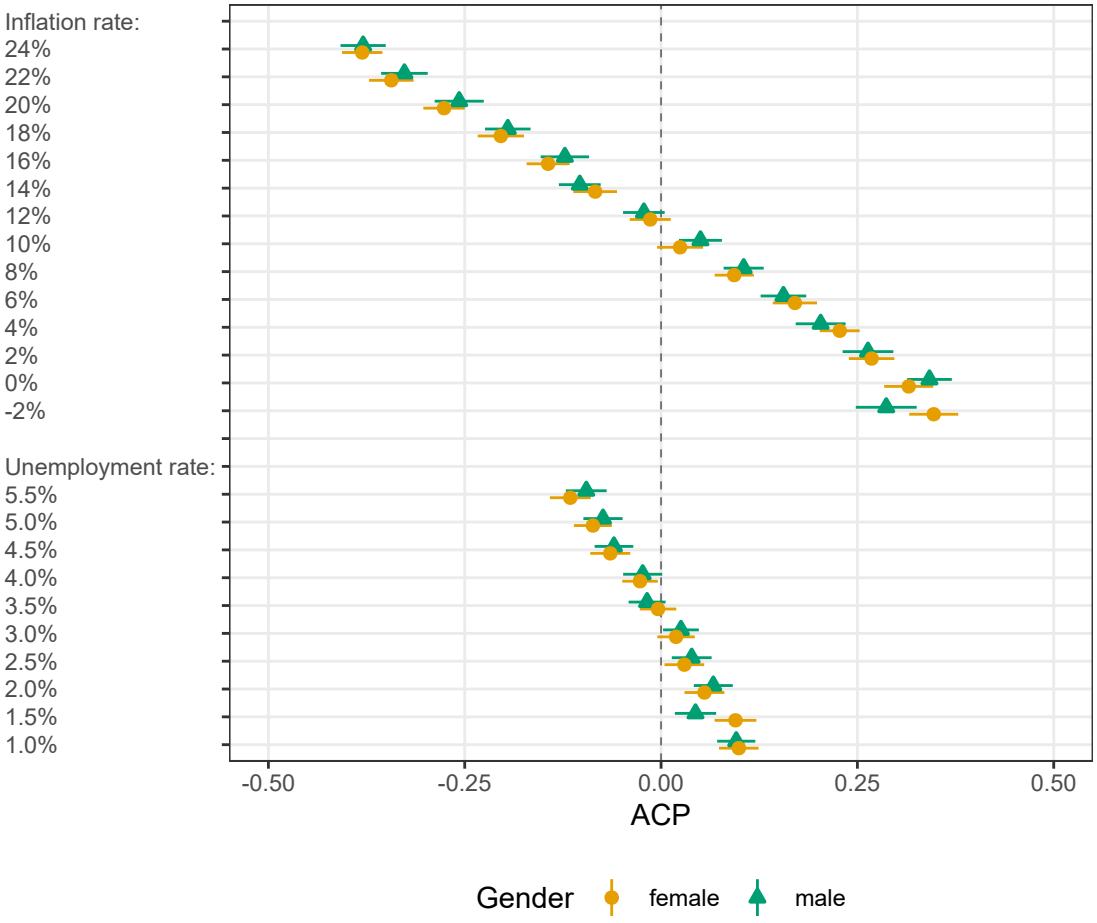

Figure 4: ACPs of inflation rates and unemployment rates by gender. Each dot is the ACP of the corresponding inflation or unemployment level, estimated separately for men ($N = 480$ respondents) and women ($N = 462$) with the method of Ganter (2023). Horizontal bars are 95% confidence intervals.

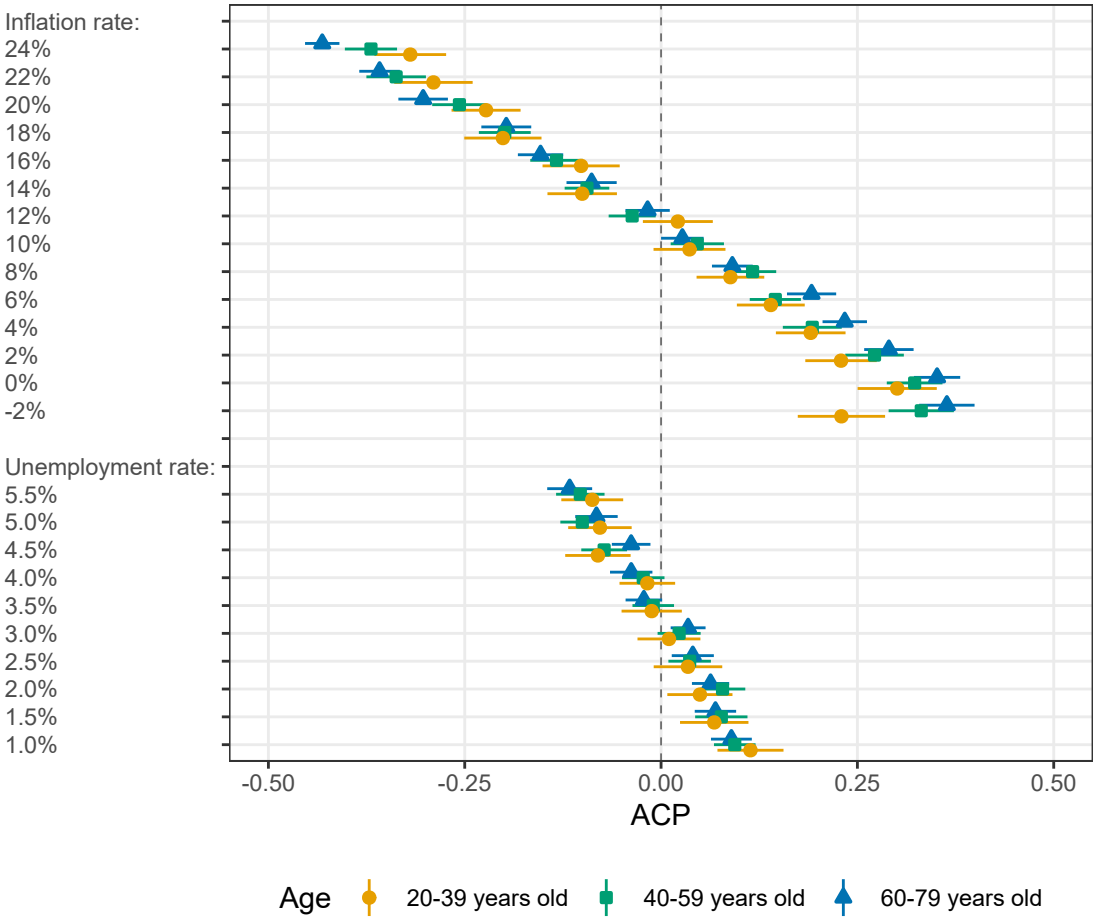

Figure 5: ACPs of inflation rates and unemployment rates by age group. Each dot is the ACP of the corresponding inflation or unemployment level, estimated separately for respondents aged 20–39 (Young, $N = 235$), 40–59 (Middle, $N = 354$), and 60–79 (Elder, $N = 353$) with the method of Ganter (2023). Horizontal bars are 95% confidence intervals.

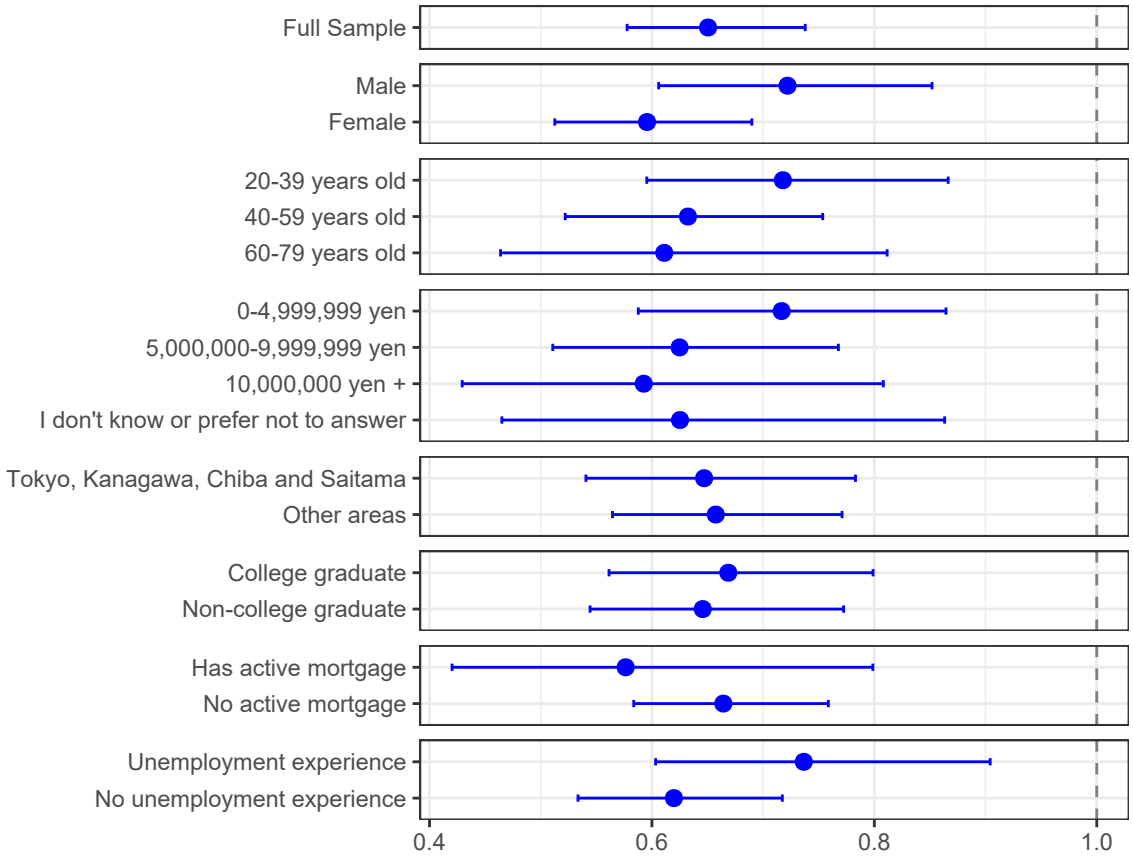

Figure 6: Marginal rate of substitution between inflation and unemployment, for the full sample and by subgroup. Each dot is the estimate of λ from equation (4), estimated with the method of Binetti et al. (2024) on the indicated subsample, and the horizontal bar is its 95% confidence interval computed by the delta method. The vertical line at one marks equal weight on inflation and unemployment. Numbers of respondents: full sample 942; men 480, women 462; aged 20–39 235, 40–59 354, 60–79 353; income below 5 million yen 437, 5–10 million yen 287, 10 million yen or more 95, income not reported 123; college graduates 433, others 509; with a mortgage 133, without 809; with unemployment experience 328, without 614.

Appendix

A pAMCE and ACP for Inflation and Unemployment in the United States

This appendix shows that the difference between the uniform AMCE and the population AMCE documented in Section 4.1 is not specific to Japan. We repeat the analysis with the U.S. conjoint data of Binetti et al. (2024), in which 2,264 respondents each completed five choice tasks, giving 22,640 profile evaluations, with inflation rates from 0% to 16% and unemployment rates from 2% to 16% in steps of 2 percentage points. We weight the profiles by the historical U.S. distributions of inflation and unemployment, computed from monthly data in the same way as for Japan. Section 4.4 discusses the comparison with Japan.

Figures A.1 through A.3 report the results. Figure A.1 plots the uAMCE and the pAMCE of inflation relative to 0%. Weighting by the historical U.S. distributions, rather than uniformly as in Binetti et al. (2024), makes the aversion to positive inflation more pronounced: for every level from 2% to 14%, the pAMCE is more negative than the uAMCE, and at 2% and 4% the pAMCE is roughly twice the uAMCE in absolute value. Figure A.2 plots the two estimates for unemployment relative to 2%. They are similar, except that 6% unemployment lowers the choice probability significantly relative to the 2% baseline under uniform weights but not under historical weights. Figure A.3 plots the ACPs of inflation and unemployment. The ACPs of inflation are much larger in magnitude than those of unemployment, which is consistent with the conclusion of Binetti et al. (2024) that Americans weigh inflation more heavily than unemployment and contrasts with the Japanese pattern in Figure 3, in which the relative weight on unemployment is much larger. We base this cross-attribute comparison on the ACP rather than on the AMCEs, because the AMCE of an inflation level and the AMCE of an unemployment level are each defined relative to their own baseline and are not on a common scale (de la Cuesta et al., 2022).

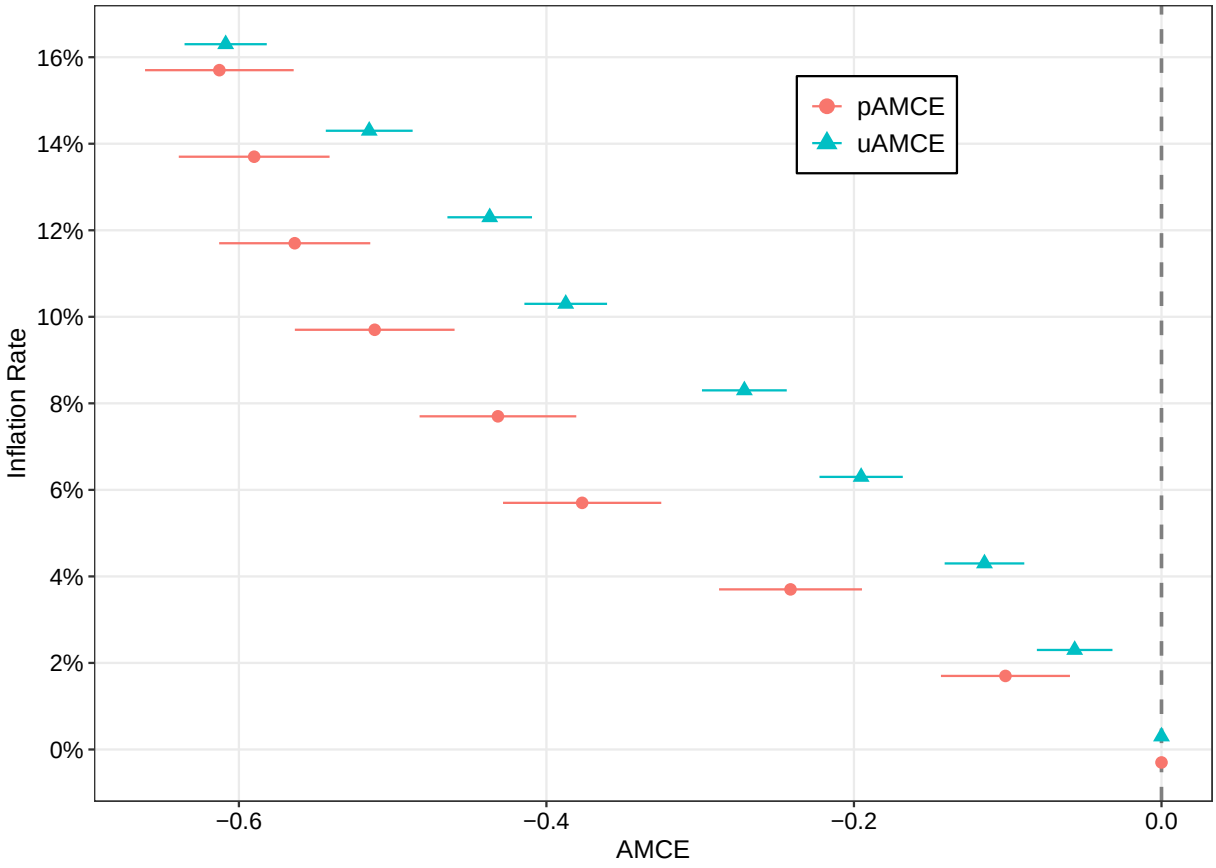

Figure A.1: uAMCE and pAMCE of inflation rates in the United States. Each dot is the uAMCE or pAMCE of the corresponding inflation level relative to 0% inflation, estimated with the data of Binetti et al. (2024) and the model-based method of de la Cuesta et al. (2022). The pAMCE averages over the historical U.S. distributions of inflation and unemployment. Horizontal bars are 95% confidence intervals. Sample: 2,264 respondents, 11,320 tasks, 22,640 profiles.

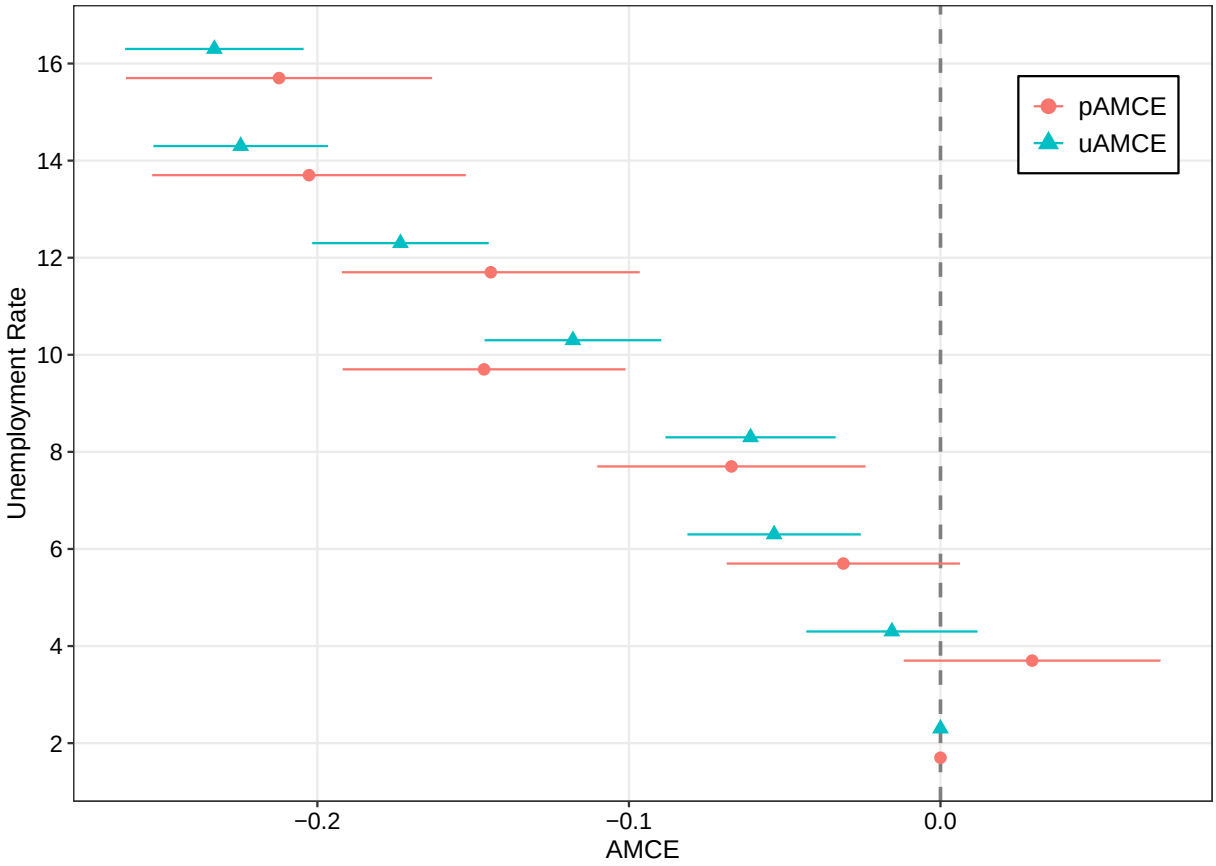

Figure A.2: uAMCE and pAMCE of unemployment rates in the United States. Each dot is the uAMCE or pAMCE of the corresponding unemployment level relative to 2% unemployment, estimated with the data of Binetti et al. (2024) and the model-based method of de la Cuesta et al. (2022). The pAMCE averages over the historical U.S. distributions of inflation and unemployment. Horizontal bars are 95% confidence intervals. Sample: 2,264 respondents, 11,320 tasks, 22,640 profiles.

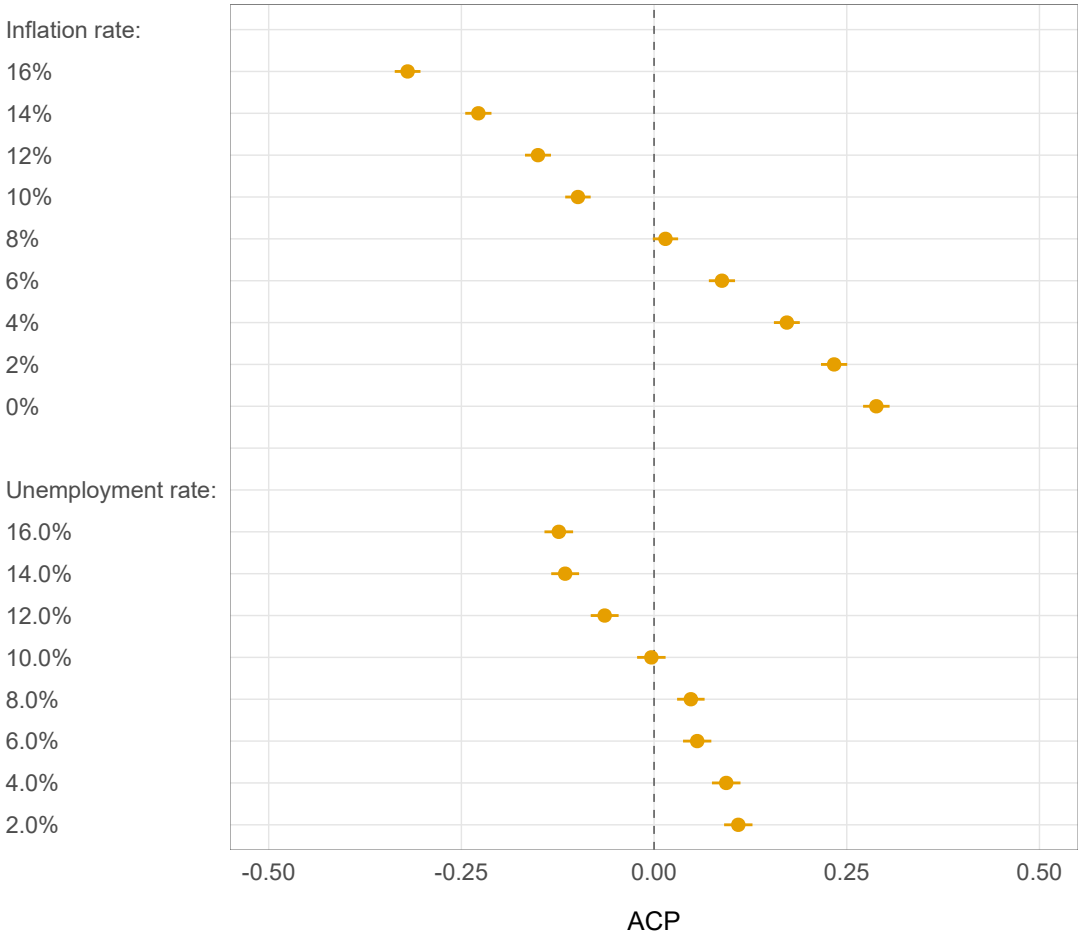

Figure A.3: ACPs of inflation rates and unemployment rates in the United States. Each dot is the ACP of the corresponding inflation or unemployment level, estimated with the data of Binetti et al. (2024) and the method of Ganter (2023). Horizontal bars are 95% confidence intervals. Sample: 2,264 respondents, 11,320 tasks, 22,640 profiles.

B Survey Questionnaire

B.1 Respondent's Demographic Questions

Q.1. Please select your gender.

- Male
- Female
- Prefer not to answer

Q.2. Please provide your age.

- 20s
- 30s
- 40s
- 50s
- 60s
- 70s

Q.3. Please select the prefecture where you live.

- (Prefecture selection: Hokkaido to Okinawa, Other, 48)

Q.4. Please select your highest level of education.

- Elementary / Junior high school graduate
- High school graduate
- Technical college graduate
- Vocational school graduate
- Junior college graduate

- University graduate
- Graduate school graduate

Q.5. Please select your total annual income (before taxes) for the past year.

- No income
- Less than 1 million yen
- 1 million to less than 2 million yen
- 2 million to less than 3 million yen
- 3 million to less than 4 million yen
- 4 million to less than 5 million yen
- 5 million to less than 6 million yen
- 6 million to less than 7 million yen
- 7 million to less than 8 million yen
- 8 million to less than 9 million yen
- 9 million to less than 10 million yen
- 10 million to less than 12 million yen
- 12 million to less than 14 million yen
- 14 million to less than 16 million yen
- 16 million to less than 18 million yen
- 18 million to less than 20 million yen
- 20 million yen or more
- Don't know
- Prefer not to answer

Q.6. Do you currently have an active mortgage you are paying off?

- Yes
- No

B.1.1 Explanation of the Inflation Rate and the Unemployment Rate

Before proceeding to the following questions, we will explain the **inflation rate** and **unemployment rate**.

In this survey, the **inflation rate** represents **how much the price of a "basket of goods and services" typically consumed by Japanese households has changed over a specific period (usually one year)**. For example, an annual inflation rate of 10% means that a basket of goods and services that cost 100 yen at the beginning of the year will cost 110 yen on average by the end of the year.

The **unemployment rate** is the proportion of **unemployed persons** to the labor force (the sum of employed and unemployed persons). Here, **unemployed persons** refers to **"people who are not currently working but are actively looking for a job and are available to start working immediately."** People who are not looking for a job or cannot work, such as students, retirees, and full-time homemakers, **are not included in the unemployed**.

For your reference, Japan's unemployment rate last year was **approximately 3%**, and the inflation rate (year-on-year) was **approximately 3%**.

next page:

Q.7. What was Japan's approximate unemployment rate last year? Please choose the correct answer.

- 1%
- 2%
- 3%
- 4%

”3%” is the correct answer.

B.2 Conjoint Experiment

In the next question, we will present two hypothetical economies with different inflation and unemployment rates. Please choose which economy you would prefer to live in.

next page:

	Option A	Option B
Inflation Rate	XA%	XB%
Unemployment Rate	YA%	YB%

If you had to choose, in which of the two hypothetical economies would you prefer to live for the next year? Please choose one from the following options.

Option A / Option B

- Forced choice
- The combination of attribute levels is randomly drawn for each task.
- There are 10 tasks in total.
- The order of the attributes is randomized across respondents.

B.3 Respondent’s Unemployment Experience

Q.8. Have you ever been unemployed and looking for a job for at least one week up to the present?

- Yes
- No