ISSN 1936-5330

kcFED Research Working Papers

Heterogeneity in Household Inflation Expectations: Policy Implications

Taeyoung Doh, JiHyung Lee, and Woong Yong Park

July 2024
RWP 24-06
http://doi.org/10.18651/RWP2024-06

Federal Reserve Bank of Kansas City

Heterogeneity in Household Inflation Expectations: Policy Implications[*]

Taeyoung Doh[†], JiHyung Lee[‡], Woong Yong Park[§]

Abstract

We empirically characterize the heterogeneity in the conditional distribution of household inflation expectations across demographic groups using the Survey of Consumer Expectations and investigate how monetary policy shocks affect the conditional distribution. We find that across all demographic groups, the peak of the group-specific distribution of household inflation expectations aligns closely with the Federal Reserve's 2 percent target. However, we also find substantial heterogeneity both within and across groups, primarily on the right end of the distribution. Nevertheless, we show that a contractionary monetary policy shock identified by high-frequency financial market response reduces inflation expectations of households more vulnerable to the risk of unanchoring.

Keywords: household inflation expectations, monetary policy, high-frequency identification, quantile regressions

JEL Codes: E31, E52, E58

[*]We thank Michael D. Bauer, Mathieu O. Pedemonte, Kwangyong Park, Jun Il Kim, and Yuriy Gorodnichenko, and seminar participants at Seoul National University, the Federal Reserve Bank of Kansas City, the Conference on "Fixed Income Markets and Inflation" at the Federal Reserve Bank of Chicago, and various places, for their discussions and comments. Views expressed here are solely those of authors do not represent the official view of the Federal Reserve Bank of Kansas City or the Federal Reserve System. This version: June 2024.

[†]*Corresponding author*. Senior Economist, Federal Reserve Bank of Kansas City. 1 Memorial Drive, Kansas City, Missouri 64198. Phone: +1-816-881-2780. Email: Taeyoung.Doh@kc.frb.org.

[‡]Associate Professor. Department of Economics, University of Illinois at Urbana-Champaign

[§]Associate Professor. Department of Economics and Institute of Economic Research, Seoul National University.

1 Introduction

Household inflation expectations are closely watched by the Federal Reserve who tries to manage them at a level close to the inflation target. Since the Federal Reserve announced the 2 percent target inflation rate in 2012, median long-run inflation expectations from consumer survey data have been generally stable. For instance, the median forecast for five-year inflation from the University of Michigan consumer survey (MSC) has been within the range of 2.2 percent to 3.2 percent during the period between January 2012 and December 2023 although yearly inflation measured by the headline consumer price index (CPI) varied a lot, between -0.2 percent and 8.9 percent as shown in Figure 1. While the stability of long-run inflation expectations are encouraging, the inflationary episode in the 1970s suggests that central banks cannot be complacent because the gradual drift in the near-term inflation expectations can signal the risk of losing the inflation anchor (Reis (2021)). Indeed, the median forecast of one-year ahead inflation from the Michigan Survey fluctuated between 2.1 percent and 5.4 percent during the same period, much more than the five-year counterpart. Hence, it is important to evaluate the effectiveness of monetary policy in anchoring near-term inflation expectations close to the central bank's target.

One challenge in assessing the effectiveness of monetary policy for stabilizing household inflation expectations is the substantial heterogeneity in inflation expectation across demographic or socio-economic groups. As well summarized by D'Acunto et al. (2023), women tend to have higher inflation expectations than men. Also, low income and less educated households seem to have higher inflation expectations than other groups. In addition, when we dig into the micro consumer survey data, we find that even households with similar demographic or socio-economic characteristics exhibit significant differences in inflation expectations. This within-group heterogeneity is greater in low income and less educated households. Understanding how monetary policy can influence the heterogeneity in household inflation expectations is important because changes in the cross-sectional distribution of inflation expectations can be informative about future shifts in inflation. For example, Reis (2021) shows that in the early 1970s, a change in right skewness in household inflation expectations in the U.S. was predictive of future inflation but he does not decompose the right tail part of household inflation expectations across different groups.

In this paper, we characterize the distribution of household inflation expectations conditional on demographic and socio-economic characteristics using survey data in the U.S. Our method allows us to analyze changes in the distribution of inflation expectations for different groups after controlling group-specific heterogeneities. By controlling potential confounding factors, we can evaluate the treatment effect of monetary policy on household inflation expectations more precisely. To this end, we estimate a conditional quantile regression, which provides a flexible modeling of the conditional distribution of household inflation expectations. Specifically, we run the conditional quantile regression of one-year ahead inflation expectations from the survey data on demographic and socio-economic characteristics as well as some macroeconomic variables such as a monetary policy shock identified by high-frequency financial market movements from Bauer and Swanson (2023), CPI inflation, the unemployment rate, and gasoline price inflation. The demographic and socio-economic characteristics include income, homeownership, level of education, gender, number of kids and adults in a household, age, region, numeracy score, and survey tenure. The conditional quantile regression is estimated using the Survey of Consumer Expectations (SCE) data published by the Federal Reserve Bank of New York. The sample is monthly and covers the period from June 2013 to December 2019.[1]

Regarding the distribution of household inflation expectations conditional on the demographic and socio-economic characteristics, we find that, across all the groups, the peak of the group-specific distribution of household inflation expectations aligns closely with the 2% target set by the Federal Reserve, but there is substantial heterogeneity in the left and right

[1] MSC provides a longer sample than SCE, which dates back to late 1970s, but it does not have information on numeracy or economic literacy of survey respondents. Since this is an important factor that determines the level of inflation expectations of a household, we chose to work with the SCE data. The cost of working with the SCE data is a shorter sample which starts only in 2013. We do a robustness check using the MSC data.

tails. The between-group heterogeneity primarily arises in the upper quantiles of household inflation expectations.[2] In other words, while most households in each group hold inflation expectations close to the inflation target by the Federal Reserve, some groups have relatively more households predicting high inflation, which generates the between-group heterogeneity. The difference in the length of the right tails across groups is quite substantial while the difference in the length of the left tails across groups is rather limited.

Income, education and gender are estimated to be important characteristics that are associated with the group-specific distribution of household inflation expectations. Households with low income or less education tend to predict higher inflation than households with high income or more education. However, the difference between groups is much bigger in upper quantiles than in lower quantiles. Compared to the households in the high income group, the households in the low income group predict higher inflation by 0.207%p at the 25% quantile but by 1.619%p at the 75% quantile. Households with at most high school diplomas predict higher inflation by 0.077%p than households with some college education or more at the 25% quantile but by 1.537%p at the 75% quantile. The same pattern is also observed when we compare female and male survey respondents. Women tend to predict higher inflation than men but the difference by gender is much bigger in the upper quantiles than in the lower quantiles. The female respondents predict higher inflation by 0.105%p than the male respondents at the 25% quantile but by 1.214%p at the 75% quantile. That is, some female respondents predicting relatively high inflation one year ahead are responsible for most of the difference in inflation expectations between the male and female group. The difference in the 25% quantile between the male and female group is not very large, although it is statistically significant.

Another important characteristic associated with the group-specific distribution of household inflation expectations is economic literacy of the survey respondents. We find that a survey

[2] Here and henceforth, unless specified otherwise, quantiles of household inflation expectations refer to quantiles of the distribution of household inflation expectations conditional on demographic and socio-economic characteristics as well as some macroeconomic variables.

respondent with more correct answers to the numeracy questions has lower inflation expectations than those with fewer correct answers. As with demographic and socio-economic characteristics, the effect of economic literacy is stronger in upper quantiles than in lower quantiles. The survey respondents who answer correctly to none of the five numeracy questions predict higher inflation by 4.220%p at the 75% quantile than those with five correct answers.

In response to a contractionary monetary policy shock, we find that households adjust their inflation expectations downward. The monetary policy shock has generally bigger impacts on upper quantiles of inflation expectations than on lower quantiles except for near extreme quantiles like 5% and 95%, whose estimates are not precise. That is, the upper quantiles of household inflation expectations are adjusted by more than the lower quantiles. Specifically, the 25%, 50%, and 75% conditional quantile of household inflation expectations decrease by 0.548%p, 0.940%p, and 1.366%p, respectively. As emphasized by Bauer and Swanson (2023), orthogonalizing high-frequency financial market responses with respect to past macro and financial information is crucial for a well identified monetary policy shock. When we use the unorthogonalized measure of a monetary policy shock in Bauer and Swanson (2023) that does not remove the predictability by past macro and financial information, we do not obtain a significant and negative coefficient on the monetary policy shock measure in the quantile regression. To check the robustness of our findings to additional heterogeneities, we run the conditional quantile regression of household inflation expectations with the monetary policy shock interacted with one of the demographic and socio-economic characteristics at a time. Although the uncertainty surrounding the regression coefficients are somewhat larger in this alternative version, we find similar patterns. Using the orthogonalized monetary policy shock measure tends to recover a significant and negative coefficient while using the unorthogonalized measure does not. Quite interestingly, the contractionary monetary policy shock is more effective in stabilizing the upper quantile of households in the low income group when we use the orthogonalized monetary policy shock measure. This may be because households who are confident or more knowledgeable about the central bank's systematic reaction to stabilize inflation would not respond to short-term economic news much as pointed out by Blinder et al. (2024).

Related Literature: This paper is related to a growing literature on exploring survey-based household inflation expectations. Weber et al. (2022) and Blinder et al. (2024) provide a comprehensive review of the literature and we will focus on papers examining the heterogeneity in inflation expectations specifically, which are most closely related to our paper.

Madeira and Zafar (2015) show a vast degree of the heterogeneity in household inflation expectation across demographic and socio-economic characteristics based on the MSC data. They find female respondents with a low level of education have a higher degree of heterogeneity in inflation expectations after controlling for publicly available information. Our finding of a higher within-group heterogeneity in this subgroup is consistent with their finding but we consider the treatment effect of a monetary policy shock in a quantile regression framework, which they do not consider. Armantier et al. (2020) analyze the SCE data to examine how the COVID-19 pandemic affected inflation beliefs across different households. They find a polarization in inflation beliefs at the onset of the pandemic with some households expecting high inflation and others expecting low inflation or deflation. Although highly educated (college diploma and above) and high numeracy respondents saw the pandemic largely as a deflationary demand shock, lowering their inflation expectations, the polarization in belief was rather uniform along other socio-demographic dimensions. They do not investigate the effect of monetary policy on inflation expectations as we do in this paper. Since the measure of monetary policy shock from Bauer and Swanson (2023) that we use is not available for the pandemic period, we cannot include the pandemic period in our analysis. Ahn, Xie, and Yang (2024) is most closely related to our paper. They investigate the heterogeneous effect of monetary policy on household expectations across homeowners and renters based on both the MSC and the SCE data. They find that homeowners' inflation expectations are more responsive to monetary policy and explain their finding using a rational inattention model with two types of households: homeowners and renters. Unlike Ahn, Xie, and Yang (2024) we control many more demographic and socio-economic characteristics beyond homeownership and consider the distribution of inflation expectations not just the conditional mean.[3]

Our choice of relevant demographic and socio-economic characteristics are motivated by other papers. Burke and Manz (2014) emphasize the connection between economic literacy and inflation expectations. Kim and Binder (2023) find a "survey tenure" effect from the SCE data, which may contaminate the interpretation of the analysis based on the SCE data when this effect is not controlled. We include these variables on top of income, education, and gender discussed in Madeira and Zafar (2015).

Our paper is also related to D'Acunto et al. (2022), who study the effect of unconventional fiscal and monetary policies on managing household expectations using German survey data. While they find that the unexpected announcement of a value-added tax (VAT) increase in Germany in 2005 to be implemented in 2007 significantly affected household inflation expectations and the willingness to purchase durable goods, they do not find a similar effect for the forward guidance on the monetary policy adopted by the ECB in 2013. In contrast, our result suggests that policy communications including forward guidance on the rate path can be effective in influencing household inflation expectations, especially at upper quantiles of the distribution. Our sample period includes the period when the federal funds rate was constrained by the effective lower bound when the policy shock reflects solely the effect of unconventional policies such as forward guidance and asset purchases.[4] Since D'Acunto et al. (2022) do not use a policy shock measure orthogonalized in the way of Bauer and Swanson (2023), their results are not directly comparable. If we use the unorthogonalized version of

[3]We do not allow the quantile coefficient on a monetary policy shock to be different across groups unlike Ahn, Xie and Yang (2024) in the baseline specification but we interact each characteristic one by one with the monetary policy shock in the alternative specification to check the robustness of our findings.

[4]The federal funds rate target was at the effective lower bound during about 30% of our sample period. Even after the federal funds rate target was lifted off the bound, our measure partially captures forward guidance but we do not separately identify the forward guidance factor during this off-the-bound period. To isolate the effect of unconventional monetary policies, we estimate the model using the MSC data for the period when federal funds rate was constrained by the effective lower bound and report the result in the online appendix.

the monetary policy shock, our result is not much different from what they find, suggesting that the proper identification of a monetary policy shock can be important for assessing policy effectiveness in changing household inflation expectations.

In multiple waves of large online surveys that randomizes information treatment on monetary policy, Knotek et al. (2024) find that numerical literacy is associated with the probability of hearing monetary policy news. The finding suggests that an important role of numeracy score in our study might be associated with different degrees of attention for monetary policy news or inflation.

Our paper is related to another growing literature on using a quantile regression framework to estimate tail risks in macroeconomic outcomes (Adrian et al. (2019) , López-Salido and Loria (2020) among others). These papers use a time series of realized aggregate data and do not include cross-sectional distribution information, which we do by leveraging the panel structure of the survey data. Including the cross-sectional information in household survey data brings a challenge because most responses are likely to be rounded and ignoring this feature might distort the sampling uncertainty. We address this issue by using the "jittering" method in Machado and Silva (2005) originally developed for running quantile regression for discrete counting data. The jittering allows us to transform an integer response in the original data to a continuous real variable so that we can apply a standard quantile regression framework. We extend the standard jittering method by group-specific heteroskedasticity in the jittering noise, which we calibrate from the subjective probability distribution of an individual household from the SCE data.

Our paper is organized as follows. Section 2 discusses our empirical strategy by describing the underlying data and quantile regression framework. Section 3 provides empirical results. Section 4 discusses potential policy implications. Section 5 concludes.

2 Empirical strategy

2.1 Data

Before we describe the specification of our empirical analysis, let us first explan the data for empirical analysis. We seek to estimate how demographic and socio-economic characteristics determine the household inflation expectations and also how changes in macroeconomic variables affect the same distribution using the conditional quantile regression.

For household inflation expectations, we primarily use the SCE data collected by the Federal Reserve Bank of New York along with their demographic characteristics. Every month, SCE interviews approximately 1,300 household heads via the Internet. It has a rotating panel design where respondents participate in the panel for up to twelve months, with a roughly equal number rotating in and out of the panel each month. We do not utilize this rotating panel design in the baseline analysis. The data is available since June 2013 and we use the sample up through December 2019, right before the coronavirus pandemic. Following the literature, we winsorize the dataset by dropping observations less than or equal to the 3% quantile and greater than or equal to 97% quantile for each of the major characteristics. The winsorizing results in dropping about 5% in the lower and right tails, respectively. Table 1 provides the detailed information on demographic and socio-economic characteristics in the SCE data.

Burke and Manz (2014) conduct a lab experiment to find that economic literacy contributes to success of forecasting inflation as more literate subjects are better able to make use of given data and more likely to select highly relevant information. To replicate this result in our setting, we control for economic literacy of survey respondents by including the numeracy score as a measure of economic literacy in the conditional quantile regression. SCE asked five numeracy questions until April 2015 when it started to ask two more numeracy questions. We use only the first five numeracy questions for analysis to maximize the sample period. We consider the numeracy score as a measure for economic literacy of survey respondents.

Kim and Binder (2023) find that survey respondents lower their inflation expectations as they repeatedly participate in the survey using the SCE data and attribute their result to the learning-through-survey effect. Following them, we control for the survey tenure, which is the number of times each respondent has finished the survey including the last one he or she participates in. As in Kim and Binder (2023), we include dummy variables for each round of the survey per household to capture potential nonlinearity in the effect of the learning-through-survey.

We include a monetary policy shock in the quantile regression, which is a measure of the monetary policy shock identified by Bauer and Swanson (2023) using high-frequency financial market responses (Eurodollar futures contracts) around FOMC announcements. Their monetary policy shock addresses a concern raised for the monetary policy shocks identified using high-frequency surprises previously in the literature by removing the predictability of the monetary policy shock based on information in past macroeconomic and financial data. In the literature, the monetary policy shock identified using high-frequency surprises was shown to be predictable with macroeconomic and financial data that pre-dates the FOMC announcement, which raised a doubt about its exogeneity. Bauer and Swanson (2023) propose to orthogonalize the monetary policy surprises in terms of information available in the past macroeconomic and financial data to make it exogenous. They also provide the monetary policy shock series that was not orthogonalized in terms of predictability, with which we do an extra analysis to check how its effect is different from the baseline effect of the orthogonalized shock. We use the first lag of the monetary policy shock to make sure that all the survey respondents have a chance to observe a new FOMC announcement.

Lastly, we control for some macro variables. CPI inflation is the year-on-year rate of change in the CPI. The unemployment gap of a given month is the gap between the unemployment rate of the month and the twelve-month average of the unemployment rate up through the previous month. We include the second lag for CPI inflation and the unemployment rate gap to make sure that all the survey respondents have a chance to observe the announcement of new information on inflation and unemployment. Gasoline price inflation is the year-on-year rate of change in the gasoline price (US All Grades Conventional Gasoline Price).

2.2 Empirical model

The conditional quantile regression allows for a flexible modeling of the conditional distribution of household inflation expectations. We can explore the characteristics of the conditional distribution with the conditional quantile regression beyond the conditional mean that the ordinary least squares regression provides (Koenker and Bassett 1978).

Specifically, for $0 < \tau < 1$, the τ-th conditional quantile of household inflation expectations is described by the equation

$$Q_{y_{it}}\left(\tau | x_{it}, z_t, w_t\right) = \beta_{0,\tau} + x'_{i,t}\beta_{1,\tau} + z_t\delta_\tau + w'_t\gamma_\tau, \tag{1}$$

where $y_{i,t}$ is one-year ahead inflation expectations for household i in period t; $x_{i,t}$ includes dummy variables for income groups (low, middle, and high income groups), homeownership (owners and renters), education (at most high school diplomas and more than high school diplomas), gender (male and female), age groups (young $<$ 40 years old, middle-aged $\geq$ 40 years old and $<$ 60, and old $\geq$ 60 years old), and the region of primary residence (the West, Midwest, Northeast, and South), the number of kids and adults in a household, survey tenure, and numeracy score;[5] z_t is an externally-identified monetary policy shock; and w_t includes aggregate CPI inflation, the unemployment rate gap, and the gasoline price inflation rate. To capture the potentially non-linear effects of survey tenure and numeracy score, we include dummy variables for each value of the survey tenure and the numeracy score.

Among the various demographic groups, we use the following group as a base group:

[5] SCE explains that the questions on these characteristics are asked only to the first-time respondent. Hence, these variables should be time-invariant. However, the region of primary residence changes for a small set of the households so we allow $x_{i,t}$ to vary over time.

high income, homeowners, more education (more than high school diplomas), male, young generation and the West. That is, the dummy variables for this base group are omitted. In order to make sure that the survey respondents observe the macroeconomic variables at the time of survey, we properly lag the macroeconomic variables. We provide the detail on the lag assumption after we describe the data below.

One problem with using the individual survey data in quantile regression is that solicited values of inflation expectations are rounded. SCE solicits one-year ahead inflation expectations by first asking whether a respondent thinks that there will be inflation or deflation over the next 12 months and then asking what he or she expects the rate of inflation or deflation over the next 12 months.[6] The survey respondent can answer any numbers, integer or non-integer, but more than 90% of the answers are an integer in our sample. This is problematic when running conditional quantile regressions as well as computing unconditional quantiles since the quantiles do not exhibit so much variation. More importantly, it violates the sufficient condition for asymptotically valid inference of the conditional quantile regression that the conditional probability density function be continuous.

To address this problem, we *jitter* the integer data of inflation expectations, or add a random noise to the integer data, and construct a continuous variable with the conditional quantiles that have a one-to-one relationship with the conditional quantiles of the integer data. Explicitly or implicitly, the respondents are likely to round their expectations to the nearest integer to answer. Therefore, the jittering process can be thought of replicating this mental process of approximation. It is also similar to the linear interpolation of the empirical distribution of inflation expectations used by SCE to compute the median and the quantiles.

The jittering method was originally proposed by Machado and Silva (2005) to run the conditional quantile regression on count data. For our exercise, we use the random noise uniformly distributed over $(-0.5, 0.5)$. We also generate 500 jittered samples and take the average of the estimates across the jittered samples to improve efficiency of our estimator.

[6] SCE also asks the probability distribution of one-year ahead inflation but we do not use that information in this study.

The standard jittering method assumes homoskedastic jittering noises. However, inflation expectations of households whose numeracy scores are lower may be subject to larger variances of jittering noises. The SCE data provide information on the subjective probability distribution of inflation by individual respondents. We allow the group-specific heteroskedasticity in the jittering noise by calibrating the jittering noise variance of household with the same numeracy score based on the standard deviation or the interquartile range of the subjective probability distribution of inflation averaged across households with the same numeracy score.

The baseline specification controls the individual heterogeneity only in the level of inflation expectations but assumes that households share the common quantile regression coefficient on the monetary policy shock. In principle, they may react differently to the same monetary policy shock depending on demographic and socio-economic characteristics. To explore how these features missing in the baseline specification influence quantile regression results, we run the quantile regression after allowing for a heterogenous quantile response to a monetary policy shock across different demographic and socio-economic characteristics. Specifically, we interact each characteristic one by one with the monetary policy shock as in

$$Q_{y_{it}}(\tau|x_{it}, z_t, w_t) = \beta^r_{0,\tau} + x'_{1,i,t}\beta^r_{1,\tau} + z_t x'_{2,i,t}\delta^r_\tau + w'_t\gamma^r_\tau, \qquad (2)$$

where x_{it} is partitioned into $x_{1,i,t}$ and $x_{2,i,t}$ and $x_{2,i,t}$ is interacted with z_t to obtain the quantile response coefficient.

Conditional quantile regression is estimated using the R package `quantreg` (Koenker 2022). Jittering introduces additional sampling uncertainty and we address this issue by computing the confidence interval by xy-pair resampling.

3 Empirical results

This section reports the estimation results of the conditional quantile regressions and discusses the results.

3.1 Coefficient estimates

Table 2 reports the coefficient estimates and their 90% confidence intervals at the 25%, 50%, and 75% quantile together with the OLS estimates and their two-standard deviation confidence intervals. Figure 2 presents the coefficient estimates for a set of the quantiles ranging from 5% to 95%. The bands around the point estimates indicate the 90% and 95% confidence intervals. For reference, we also estimate the OLS of the same specification and report its coefficient estimates using the horizontal dashed lines.

Let us first look at the coefficient estimates on the monetary policy shock. In response to a contractionary monetary policy shock, we find that most of the households adjust their inflation expectations significantly downward. The estimates suggest that monetary policy is effective in stabilizing household inflation expectations. The magnitude of the effect of the monetary policy shock is however substantially different across the quantiles. Except for the left and right extreme tail (5% and 95%), the monetary policy shock has bigger impacts on upper quantiles of inflation expectations than on lower quantiles. That is, the upper quantiles of household inflation expectations are adjusted by more than the lower quantiles. For example, the 25%, 50% and 75% conditional quantile of household inflation expectations decreases by 0.548%p, 0.940%p, and 1.366%p, respectively. Since the dependent variable of the conditional quantile regression is the level of inflation expectations, not the revision in inflation expectations, the evidence does not tell us that individual households with inflation expectations at each quantile adjust their inflation expectations by the reported estimates. However, the result implies that overall the conditional distribution of household inflation expectations shifts to the left and, at the same time, the upper quantiles shrink in response to a contractionary monetary policy shock. Overall, households adjust their inflation expectations to be substantially lower.

An increase in the past realized inflation is associated with an increase of household inflation expectations in lower quantiles but not in upper quantiles though the quantitative magnitude is small for all the quantiles. An increase in the unemployment gap is associated with a downward adjustment of household inflation expectations across the quantiles. The magnitude of the downward adjustment is significant overall and bigger in upper quantiles. For instance, a 1% increase in the unemployment gap reduces household inflation expectations by 0.379%p, 0.656%p, and 0.884%p respectively at the 25%, 50%, and 75% quantile. It is estimated that a rise in gasoline price inflation over the past year is associated with an upward adjustment of household inflation expectations, which is consistent with the findings in the literature such as Coibion and Gorodnichenko (2015) but the quantitative magnitude is smaller in our sample.[7] In sum, monetary policy shock and unemployment rate gap have quantitatively significant impacts on household inflation expectations while the past inflation and gsoline price inflation have rather small impacts.

Now we discuss the estimation results of the coefficients on the demographic and socio-economic factors. Income is estimated to be an important demographic characteristic that determines the group-specific distribution of household inflation expectations. Households with low income tend to predict higher inflation than households with high income across all the quantiles. However, the difference between income groups is much bigger in upper quantiles than in lower quantiles. Compared to the households in the high income group, the households in the low and middle income group predict higher inflation by 0.207%p and 0.080%p, respectively, at the 25% quantile of inflation expectations while by 1.619%p and 0.595%p, respectively, at the 75% quantile. The OLS estimate of the coefficient on the low and middle income group is 1.318 and 0.660, respectively, which would ignore these differences in the between-income group difference across the quantiles.

Education also plays an important role in determining the conditional distribution of household inflation expectations. Except for the left tail, households with less education (at most high school diplomas) are estimated to predict higher inflation than households with more education (more than high school diplomas). Again, the difference between the less

[7] In particular, a substantial decline in the oil price during 2014-2016 period did not lower household inflation expectations significantly, weakening the correlation between gsoline price inflation and household inflation expectations in our sample period.

educated and more educated households is bigger in the upper quantiles than in the lower quantiles. The households with less education predict higher inflation by 0.077%p, 0.491%p, and 1.537%p at the 25%, 50%, and 75% quantile, respectively, than the households with more education. The OLS estimate of the coefficient on the dummy variable for the less educated group is 0.866, which is bigger than the coefficient at the 25% and 50% quantile but smaller than the coefficient at the 75% quantile.

The same pattern of the heterogeneity in household inflation expectations also appears for gender. Females tend to predict higher inflation than males but the difference between females and males is much bigger in the upper quantiles than in the lower quantiles. The female respondents predict higher inflation by 0.105%p, 0.440%p, and 1.214%p at the 25%, 50% and 75% quantile, respectively, than the male respondents. That is, some female respondents predicting relatively high inflation one year ahead are responsible for most of the difference in inflation expectations observed between the male and female group.

Though a bit weaker than for the other demographic characteristics we discussed above, we observe the same pattern of heterogeneity for homeownership. Except for the left tail, renters are estimated to predict higher inflation than homeowners. However, the difference between the renters and homeowners is bigger in the upper quantiles than in the lower quantiles. The renters predict higher inflation by 0.028%p, 0.114%p, and 0.331%p at the 25%, 50%, and 75% quantile, respectively, than the homeowners.

It is estimated that the number of kids and adults in a household and age have a similar, though a bit weaker, pattern of the heterogeneity in household inflation expectations. The results regarding the other characteristics are presented in the online appendix.

Another important characteristic associated with the group-specific distribution of household inflation expectations is the economic literacy of the survey respondents. As the coefficient estimates on the numeracy score in Figure 2 show, a survey respondent with a high numeracy score has lower inflation expectations than those with a low numeracy score. However, the effect is not linear. Compared to the survey respondents who do not answer correctly to any of the numeracy questions, those respondents with one or two correct answers do not predict significantly lower inflation expectations while those respondents with three or more correct answers predict significantly lower inflation expectations. Interestingly, as with demographic and socio-economic characteristics, the effect of economic literacy is stronger in upper quantiles than in lower quantiles. The effect can be quite large. Compared to those with no correct answers, those households with five correct answers have one-year inflation expectations lower by about 4%p at the 75% quantile. This result can be understood based on the finding of Burke and Manz (2014) that more literate respondents are better able to make use of given data and more likely to select highly relevant information. Similarly, Knoteck et al. (2024) find that respondents with higher numeracy score are likely to pay more attention to monetary policy news.

Consistent with the finding of Kim and Binder (2023), the respondents with more rounds of the survey participation tend to have lower inflation expectations. Interestingly, the learning-through-survey effect is stronger in upper quantiles than in lower quantiles. Compared to the fresh participants in the survey, the households in the second round of the survey predicts lower inflation by 0.135%p, 0.317%p, and 0.667%p at the 25%, 50%, and 75% quantile, respectively, while those in the twelfth round of the survey predicts lower inflation by 0.290%p, 0.709%p, and 1.459%p at the 25%, 50%, and 75% quantile, respectively. The marginal effect of one more round of survey participation is somewhat diminishing in the survey tenure, which is aligned with the declining learning-through-survey effects found by Kim and Binder (2023). The diminishing effect is especially visible at the 75% quantile.

3.2 Predicted conditional quantiles

By applying the conditional quantile regression to a fine grid of probabilities, we can approximate the conditional distribution of household inflation expectations with a high degree of precision. This approximation provides valuable insights into the structure of the conditional distribution of household inflation expectations. To highlight this, we connect the predicted conditional quantiles of inflation expectations across 24 demographic groups based on the previously estimated conditional quantile regressions.

Figure 6 presents the predicted conditional quantiles in December 2019 for the groups. We find that the heterogeneity in household inflation expectations is not large in the lower quantiles but much larger in the upper quantiles. Compared to the group of the male households in the high income group and with more education and homeownership (denoted as Group 24 in the figure), the group of the female renter in the low income group and with less education (denoted as Group 1 in the figure) has higher inflation expectations across most part of the distribution. However, their difference is dramatically enlarged in the upper quantiles.

The same pattern that the between-group heterogeneity is small in lower quantiles but large in upper quantiles is also observed in the other time periods. Figure 7 presents the time series of the predicted conditional quantile of one-year ahead inflation expectations across the demographic groups for selected probabilities. As we discussed above, the heterogeneity across the groups is much larger for the 75% quantiles than for the 25% quantiles. We also find that the high income group's inflation expectations are strongly correlated (correlation coefficient of 0.71) with the five year inflation compensation measure from the Treasury Inflation-Protected Securities (TIPS) market data while the low income group's inflation expectations are less so (correlation coefficient of 0.32).[8] The finding suggests that households who can hedge inflation risk through TIPS may have better anchored inflation expectations, though the direct evidence for this hypothesis will require the analysis of TIPS holdings by different income groups. Relatedly, Blinder et al. (2024) suggest that households who are more confident about the central bank's ability to keep inflation near the target inflation rate are less likely to respond to short-term economic news.

Since the predicted conditional quantiles in Figure 6 can be considered as the inverse function of the empirical cumulative distribution function for each group, we can use the

[8]The difference in the correlation between groups is smaller at 25%p.

uniform distribution inversion to generate random numbers from the conditional distribution of household inflation expectations for each group. The conditional density function can be estimated based on the generated random numbers using the kernel density estimation.

Figure 8 presents the conditional density function of household inflation expectations for different demographic and socio-economic groups. In all groups, the peak of the group-specific distribution of household inflation expectations aligns closely with the 2% target by the Federal Reserve. However, there is substantial heterogeneity in both the left and right tails. Most notably, the heterogeneity between groups mainly occurs in the upper quantiles of household inflation expectations. This implies that while many of the households of each group hold inflation expectations consistent with the Federal Reserve's inflation target, some groups have a higher number of households that predict elevated inflation, leading to the between-group heterogeneity. There is a considerable difference in the length of the right tails among the groups, whereas the difference in the left tails is relatively minor.

3.3 Robustness checks

We carry out several exercises to check robustness of our main results and report the results in this section.

3.3.1 MSC

As a robustness check, we estimate the conditional quantile regression on the MSC data with a similar specification to the baseline one on the SCE data. To conserve space, we describe the exact specification for the MSC data in the online appendix and also report the results on the MSC data there.

The main result is qualitatively similar to the baseline result. Especially, we find that the between-group heterogeneity is mostly driven by the difference in upper quantiles of household inflation expectations.

3.3.2 Jittering

Recall that in the baseline empirical analysis, we jitter integer inflation expectations by adding a random noise from a uniform distribution over $(-0.5, 0.5)$. However, depending on their demographic characteristics, households are likely to differ in the degree of uncertainty they have about future inflation. In particular, as found by Burke and Manz (2014), economic literacy could affect how households choose and use information, and thus how uncertain they are about future inflation. Failure to account for such heterogeneity could be a source of bias in the coefficient estimates. We thus account for heterogeneity in the degree of uncertainty in jittering using the density forecast of future inflation elicited by SCE. Due to a concern on noises in the density forecast, we do not use the household-level density forecast but the dispersion of the density forecast for each numeracy score group.

Specifically, we use two types of measures of the dispersion of the density forecast for each group of households with identical numeracy scores (0 through 5): the interquartile range and the standard deviation. Both measures are computed and provided by SCE and we compute the group median of these two measures. We then jitter integer forecasts by adding a random noise from a uniform distribution over $(-IQR, +IQR)$, where IQR is the group-specific median interquartile range or by adding a random noise from a normal distribution with mean 0 and standard deviation equal to the group-specific median standard deviation.

In both cases, the coefficient estimates are quantitatively very similar to those of the baseline analysis. The results are provided in the online appendix.

3.3.3 Monetary policy treatment effects

Bauer and Swanson (2023) find that it is important to remove the predictability in the monetary policy surprises identified with the high-frequency data in order to isolate exogenous variations in the monetary policy surprises. To check how the predictability in the high-frequency-identified monetary policy surprises affect the conditional distribution of household inflation expectations, we replace the orthogonalized monetary policy shock in the baseline specification with the unorthogonalized monetary policy shock estimated by Bauer and Swanson (2023) and run the conditional quantile regression again. Figure 9 compares the coefficient estimate on the orthogonalized and unorthogonalized monetary policy shock. In response to a positive unorthogonalized monetary policy shock, we find that the conditional distribution of household inflation expectations does not respond significantly. We can infer that households do not adjust their inflation expectations significantly since they understand that an expansionary economic condition warrants the policy rate raise by the FOMC.

We then investigate whether the effect of the monetary policy shock is heterogeneous across demographic and socio-economic characteristics. The baseline specification does not allow for a heterogeneous quantile response to a monetary policy shock across different groups sorted by demographic and socio-economic characteristics. We now run the quantile regression with the monetary policy shock interacted with demographic and socio-economic characterstics one at a time for a robustness check.[9]

Figure 10 shows the estimated quantile regression coefficient on both the orthogonalized monetary policy shock and the unorthogonalized shock interacted with the household income. As in the baseline specification, we observe a significant and negative response of the household inflation expectation only to an orthogonalized monetary policy shock. Interestingly, the negative response is most pronounced at the upper quantile of the low income group. The finding suggests that a contrationary monetary policy shock is most effective in lowering inflation expectations of the low income households who tend to have higher inflation expectations than others in the same income group. Overall, our analysis supports the view that a contractionary monetary policy shock either through conventional (e.g., change in the federal funds rate target) or unconventional (e.g., forward guidance) policies can lower inflation expectations of the households who are most vulnerable to the loss of inflation anchor.[10]

[9] We interact the monetary policy shock with the characteristics one at a time because of the concern on the sample size.

[10] The response of the high income group to a monetary policy shock is relatively muted. Since the high income group's inflation expectations are well correlated with the inflation compensation from TIPS and better anchored, we suspect that their inflation expectations may be less sensitive to any news than other groups along the observation made by Blinder et al. (2024).

3.3.4 Unconventional monetary policy

The monetary policy shock identified by Bauer and Swanson (2023) covers both conventional and unconventional monetary policy shocks. We do another exercise to check whether unconventional monetary policy shocks have comparable effects on household inflation expectations to conventional monetary policy shocks. To that end, we run the conditional quantile regression on the sample where the Federal Funds Rate was at the zero lower bound (ZLB). To use a sufficiently large sample for estimation, we instead estimate the conditional quantile regression on the MSC data for the period from December 2008 through December 2015.[11]

It is estimated that the coefficients on the orthogonalized monetary policy shock on the full sample and the ZLB sample, reported in the online appendix, are qualitatively similar. The magnitude of the coefficient estimate is a bit smaller on the ZLB sample. Therefore, we can conclude that unconventional monetary policy, for example forward guidance during the ZLB period, was also effective in moving household inflation expectations. It is found that the coefficient estimates on the unorthogonalized monetary policy shock are significantly positive, reinforcing the importance of controling for the predictability in the high-frequency monetary policy surprises.

4 Policy Implications

Distinguishing the source of the heterogeneity in household inflation expectations that we document above is important for monetary policy to the extent that we can design polices in order to reduce the dispersion in inflation expectations centered around the central bank's target. We list several factors that can potentially account for the heterogeneity in household inflation expectations.

The estimation result suggests that economic literacy is a powerful source of the heterogeneity

[11] The SCE data was not available before June 2013. Since the MSC does not provide information on numeracy or economic literacy, the robustness exercise does not control it. Also, it includes a dummy variable for repeat participants as survey respondents can participate in the survey at most twice in the MSC. The exact specification for this robustness exercise is described in the online appendix.

in household inflation expectations. Especially, the group-specific distribution of economic literacy is a mirror image of the group-specific distribution of household inflation expectations as shown in Figure 11. That is, the difference in the numeracy score is rather small in the upper quantiles but large in the lower quantiles. As can be anticipated from the estimation result above, the groups with a high numeracy score are those groups with low inflation expectations while the groups with a low numeracy score are those groups with high inflation expectations. Therefore, the within-group and between-group heterogeneity in economic literacy can generate a large within-group and between-group heterogeneity in household inflation expectations.[12]

However, even after the numeracy score is controlled for in the conditional quantile regression, there still remains a considerable heterogeneity in household inflation expectations within each group and across the groups. The female respondents with low income, less education, and no homeownership have a very long right tail in the conditional distribution of inflation expectations while the male respondents with high income, more education, and homeownership have a relatively short right tail. This differential degree of the within-group heterogeneity in inflation expectations leads to a large between-group heterogeneity in upper quantiles. In other words, there is a substantial difference in the fraction of the households predicting elevated inflation across the groups.

What drives the remaining heterogeneity?[13] We first point out that the left tail of the conditional distribution in Figure 8 is much shorter than the right tail. For many of the groups, the left tail merely covers zero and does not extend significantly into negative values. Actually, in the SCE data, most of the households predict inflation and only about 5% of the households predict deflation before the early period of the coronavirus pandemic. Probably because of their lack of deflation experiences, the households tend to predict inflation. When

[12]Note that the predicted conditional quantiles in Figure 6 were computed assuming that the numeracy score is fixed at the median of the numeracy score for each group. If we take into account the within-group distribution of the numeracy score, we will observe a much more amplified heterogeneity than that observed in Figure 6.

[13]The scale effect that there is a large variation for a forecast with a large magnitude cannot explain the between-group heterogeneity.

they forecast future inflation, with all the other factors fixed, even those who make wild guesses on inflation tend to predict inflation, not deflation. This tendency is likely to contribute to the large heterogeneity in the upper quantiles and the small heterogeneity in the lower quantiles. Our finding is consistent with Gorodnichenko and Sergeyev (2021) who document evidence for zero lower bound on inflation expectations from the MSC data.

The next potential source of the heterogeneity in inflation expectations we consider is the heterogeneity in household-specific inflation that arises due to the differences in the item weights of the consumption basket. There is evidence that households with high inflation experiences expect high inflation. Johannsen (2014) also finds that groups with greater dispersion in experienced inflation also disagree more about future inflation. However, the magnitude of inflation inequality experienced by individual households does not appear to be big enough to explain the substantial heterogeneity in inflation expectations.[14]

In addition, where households glean information about monetary policy and/or inflation plays a role in determining their inflation expectations. Blinder et al. (2024) mention that the differential media coverage of economic news based on the bias due to political ideologies can influence households' inflation expectations. For example, the recent MSC data show that the perception of the economy diverge across households by political affiliations.[15] We do not control this factor in our current study given the lack of sufficient data to address this but it will certainly be an interesting research question.

Hence, all of these factors are likely to play some role in generating the heterogeneity in household inflation expectations. From a monetary policy perspective, understanding how much of this heterogeneity is tied up with the realized outcome (experienced inflation) could provide valuable insights. If certain groups have higher and dispersed inflation expectations

[14] Inflation inequality measured by Jaravel (2021) using full consumption basket from the consumer expenditure survey (CEX) shows that the difference in the experienced inflation across 10 different income groups is less than 0.4% on average and the difference in the experienced inflation of households conditional on the income group is even smaller. We also find that the quantitative magnitude of heterogeneity in the experienced inflation is relatively small to account for the substantial portion of the heterogeneity that we find from household survey data.

[15] See https://news.umich.edu/consumer-sentiment-solidifies-sharp-gains/

due to a greater price dispersion in their consumption baskets, this may have implications for the choice of price index targeted by the central bank.[16] On the other hand, if cognitive differences across households are mainly responsible for the heterogeneity in inflation expectations, targeted communication policies that aim to improve the economic and financial literacy of certain groups could help anchor inflation expectations close to the target inflation rate. In addition, increasing the accessibility to the TIPS market broadly might be helpful because households who can hedge inflation risk through TIPS likely have better anchored inflation expectation. Our finding of the strong correlation between the high income group's inflation expectations and the inflation compensation measure from TIPS is supportive of this view.

Although we do not quantitatively decompose the heterogeneity in household inflation expectations into these separate factors, our analysis has an important positive implication for the current practice of monetary policy. Since a contractionary monetary policy shock leads to a significantly downward adjustment of household inflation expectations across the quantiles, monetary policy is estimated to be effective in stabilizing inflation expectations after controlling group-specific heterogeneity in the level of inflation expectations. In response to a contractionary monetary policy shock, the conditional distribution of household inflation expectations shifts to the left and its right tail shrinks to the left. In spite of a vast degree of the heterogeneity in the level of inflation expectations across different groups, our analysis suggests that the current practice of monetary policy is effective in stabilizing inflation expectations. Further stabilization of inflation expectations may benefit from a more comprehensive study on the source of the heterogeneity in inflation expectations, which could inform different policy prescriptions.

[16] Pedemonte et al. (2023) argue that when the heterogeneity in household inflation expectations is shaped by the past belief, the optimal policy should be more aggressive to inflationary shocks to prevent agents from having memories of high inflation. In this case, not the difference in the current consumption basket but the difference in the past inflation experience has a persistent impact on inflation expectations.

5 Conclusion

Whether or not household inflation expectations are well anchored at the central bank's target is an important issue for monetary policy. The heterogeneity of inflation expectations across different demographic and socio-economic groups poses a challenge in assessing the degree of anchoring. We empirically characterize the heterogeneity in the conditional distribution of household inflation expectations across the demographic groups using the SCE data and also investigate how a monetary policy shock affects this conditional distribution. Our findings are somewhat encouraging for the current practice of monetary policy. We find that, across all the groups, the mode of household inflation expectations aligns closely with the 2% target by the Federal Reserve.

However, there is substantial heterogeneity in both within and across groups, primarily on the right tail. Nonetheless, in response to a contractionary monetary policy shock, households overall adjust their inflation expectations significantly downward at every quantile. This finding implies that monetary policy is effective in stabilizing inflation expectations to some degree in spite of significant heterogeneities in the level of inflation expectations among households. Further improving the degree of stability in inflation expectations may involve closing the gap between different households, which could benefit from further research on the sources of the heterogeneity.

References

[1] Adrian et al. 2019. "Vulnerable Growth," *American Economic Review,*, 109(4).

[2] Ahn, Xie, and Yang. 2024. "Effects of Monetary Policy on Household Expectations: The Role of Homeownership," *Journal of Monetary Economics*, Forthcoming.

[3] Armentier et al. 2020. "How Economic Crises Affect inflation Beliefs: Evidence from the COVID-19 Pandemic," *Federal Reserve Bank of New York Staff Reports*, 949.

[4] Bauer and Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual 2022* edited by Eichenbaum, Hurst and Ramey.

[5] Blinder, Ehrmann, de Haan, and Janssen. 2024. " Central Bank Communication with the General Public: Promise or False Hope?"*Journal of Economic Literature,*, 62(2).

[6] Burke and Manz. 2014. "Economic Literacy and Inflation Expectations: Evidence from a Laboratory Experiment," *Journal of Money, Credit and Banking*, 46(7).

[7] Coibion and Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 7(1).

[8] D'Acunto F, Hoang D, Weber M. 2022. "Managing households? expectations with unconventional policies," *The Review of Financial Studies.* 35(4):1597-642.

[9] D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber. 2023."What do the data tell us about inflation expectations?." In *Handbook of economic expectations*, pp. 133-161. Academic Press, 2023.

[10] D'Amico et al. 2018. "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices," *Journal of Financial and Quantitative Analysis*, 53(1).

[11] Gorodnichenko, Yuriy, and Dimitriy Sergeyev. 2021. "Zero lower bound on inflation expectations," *Working Paper 29496*, National Bureau of Economic Research.

[12] Jaravel, 2021. "Inflation Inequality: Measurement, Causes, and Policy Implications," in *Annual Review of Economics*, 13:599-629.

[13] Johannsen. 2014. "Inflation Experience and Inflation Expectations: Dispersion and Disagreement Within Demographic Groups," *Finance and Economics Discussion Series*, Federal Reserve Board, 2014-89.

[14] Knotek II, Mitchell, Pedemonte, and Shiroff. 2024. "The Effects of Interest Rate Increases on Consumers' Inflation Expectations: The Roles of Informedness and Compliance," *Working Paper 24-01*, Federal Reserve Bank of Cleveland.

[15] Kim and Binder. 2023. "Learning-through-Survey in Inflation Expectations," *American Economic Journal: Macroeconomics*, 15(2).

[16] Koenker and Bassett. 1978. "Regression Quantiles," *Econometrica*, 46(1).

[17] Koenker. 2022. *quantreg: Quantile Regression*. R package version 5.88. https://CRAN.R-project.org/package=quantreg

[18] López-Salido and Loria. 2020. "Inflation at Risk," *Finance and Economics Discussion Series*, Federal Reserve Board, 2020-013.

[19] Machado and Silva. 2005. "Quantiles for Counts," *Journal of the American Statistical Association*, 100(472).

[20] Madeira and Zafar. 2015. "Heterogeneous Inflation Expectations and Learning," *Journal of Money, Credit, and Banking*, 47(5).

[21] Pedemonte, Toma, and Verdugo. 2023. "Aggregate Implications of Heterogeneous Inflation Expectations: The Role of Individual Experience," *Working Paper 23-04*, Federal Reserve Bank of Cleveland.

[22] Reis. 2020. "The people versus the markets: A parsimonious model of inflation expectations." CFM, Centre for Macroeconomics, 2020.

[23] Reis. 2021. "Losing the Inflation Anchor," *Brookings Papers on Economic Activity*, Fall.

[24] Weber, D'Acunto, Gorodnichenko and Coibion. 2022. "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications," *Journal of Economic Perspectives*. 36(3).

Table 1: Descriptive Statistics of the SCE Data

(a) Income ($)	<10K	<20K	<30K	<40K	<50K	<60K	<75K	<100K	<150K	<200K	≥200K
Density (%)	2.8	6.4	8.0	8.3	9.1	9.0	12.5	14.6	15.7	7.0	6.5
Income Group			Low income				Middle income			High income	
Group density (%)			34.7				36.1			29.2	

	(b) Homeownership		(c) Gender		(d) Education	
	Homeowners	Renters	Male	Female	High school or less	Some college or more
Density (%)	73.8	26.2	52.1	47.9	11.3	88.7

	(e) Number of kids in a household					(f) Number of adults in a household				
	1	2	3	4	≥5	1	2	3	4	≥5
Density (%)	68.3	13.6	12.3	4.2	1.6	26.1	55.1	12.5	4.2	2.0
Cumulative density (%)	68.3	81.9	94.2	98.4	100.0	26.1	81.2	93.7	98.0	100.0

	(g) Age		
	Young (≤40 years old)	Mid-aged (≥40 and ≤60 years old)	Old (> 60 years old)
Density (%)	29.5	40.1	30.4

	(h) Region of primary residence			
	Midwest	Northeast	South	West
Density (%)	23.3	19.6	34.2	23.0

(i) Numeracy score	0	1	2	3	4	5
Density (%)	0.4	2.6	8.1	17.8	30.9	40.3
Cumulative density (%)	0.4	2.9	11.0	28.9	59.7	100.0

(j) Survey tenure	1	2	3	4	5	6	7	8	9	10	11	12
Relative density (%)	100.0	91.3	87.5	84.0	78.5	73.8	70.1	65.9	61.4	55.9	48.3	36.0

Notes: All respondents are asked in their first survey about demographic characteristics such as (a) income, (b) homeownership, (c) gender, (d) education, (e) number of kids and (f) adults in their households, (g) age, and (h) region of primary residence, and asked to answer (i) numeracy questions. The descriptive statistics on characteristics from (a) income through (i) numeracy score is at the respondent level, whose total number is 12,600. The total number of observations is 86,961. Since the attrition rate of the respondents is similar across the demographic characteristics, the composition of the monthly sample is similar over time. Some respondents change the region of their primary residence, in which case each of their answers is counted separately, despite that the SCE questionnaire reports that the question on the primary residence is asked only at the first interview. There are only 127 of such cases (about 1.0% of all the respondents).

Table 2: Estimation results: baseline specification

Quantiles	25%	50%	75%	OLS
(Intercept)	**1.690**	**4.119**	**7.645**	**4.953**
	(1.156, 2.301)	(3.581, 5.488)	(6.419, 9.762)	(4.452, 5.454)
Monetary policy shock (L1)	**-0.548**	**-0.940**	**-1.366**	**-1.331**
	(-0.984,-0.109)	(-1.545,-0.388)	(-2.358,-0.325)	(-2.255,-0.408)
CPI inflation (L2)	**0.054**	**0.017**	**0.000**	**0.006**
	(0.027, 0.080)	(-0.016, 0.051)	(-0.054, 0.063)	(-0.047, 0.060)
Unemployment rate gap (L2)	**-0.379**	**-0.656**	**-0.884**	**-0.739**
	(-0.445,-0.312)	(-0.736,-0.568)	(-1.040,-0.739)	(-0.866,-0.612)
gasoline price inflation	**0.003**	**0.005**	**0.004**	**0.005**
	(0.002, 0.005)	(0.003, 0.006)	(0.001, 0.007)	(0.002, 0.007)
Low income	**0.207**	**0.646**	**1.619**	**1.318**
	(0.172, 0.243)	(0.596, 0.700)	(1.498, 1.730)	(1.245, 1.390)
Mid income	**0.080**	**0.243**	**0.595**	**0.660**
	(0.052, 0.111)	(0.208, 0.277)	(0.526, 0.653)	(0.595, 0.725)
Renters	**0.028**	**0.114**	**0.331**	**0.352**
	(-0.009, 0.064)	(0.067, 0.163)	(0.242, 0.418)	(0.286, 0.417)
High school or less	**0.077**	**0.491**	**1.537**	**0.866**
	(0.016, 0.135)	(0.400, 0.578)	(1.324, 1.745)	(0.779, 0.953)
Female	**0.105**	**0.440**	**1.214**	**1.076**
	(0.078, 0.132)	(0.397, 0.479)	(1.139, 1.299)	(1.022, 1.130)
Number of kids	**0.023**	**0.068**	**0.121**	**0.093**
	(0.007, 0.038)	(0.050, 0.084)	(0.094, 0.147)	(0.064, 0.121)
Number of adults	**0.006**	**0.020**	**0.050**	**0.026**
	(-0.002, 0.014)	(0.011, 0.027)	(0.026, 0.078)	(0.013, 0.038)
Middle-aged	**0.269**	**0.422**	**0.629**	**0.629**
	(0.238, 0.304)	(0.379, 0.463)	(0.561, 0.705)	(0.563, 0.695)
Old	**0.512**	**0.671**	**0.882**	**0.911**
	(0.474, 0.550)	(0.624, 0.712)	(0.804, 0.967)	(0.836, 0.986)
Midwest	**-0.082**	**-0.133**	**-0.239**	**-0.239**
	(-0.117,-0.038)	(-0.182,-0.083)	(-0.322,-0.154)	(-0.314,-0.164)
Northeast	**-0.117**	**-0.243**	**-0.424**	**-0.310**
	(-0.154,-0.079)	(-0.283,-0.198)	(-0.512,-0.346)	(-0.388,-0.231)
South	**-0.028**	**0.009**	**0.076**	**0.079**
	(-0.064, 0.010)	(-0.033, 0.059)	(-0.006, 0.175)	(0.010, 0.149)

Notes: We report the coefficient estimates on the dummy variables for the survey tenure and the numeracy score only in Figures 4 and 5 to save the space. For the number estimates on these dummy variables, see the online appendix.

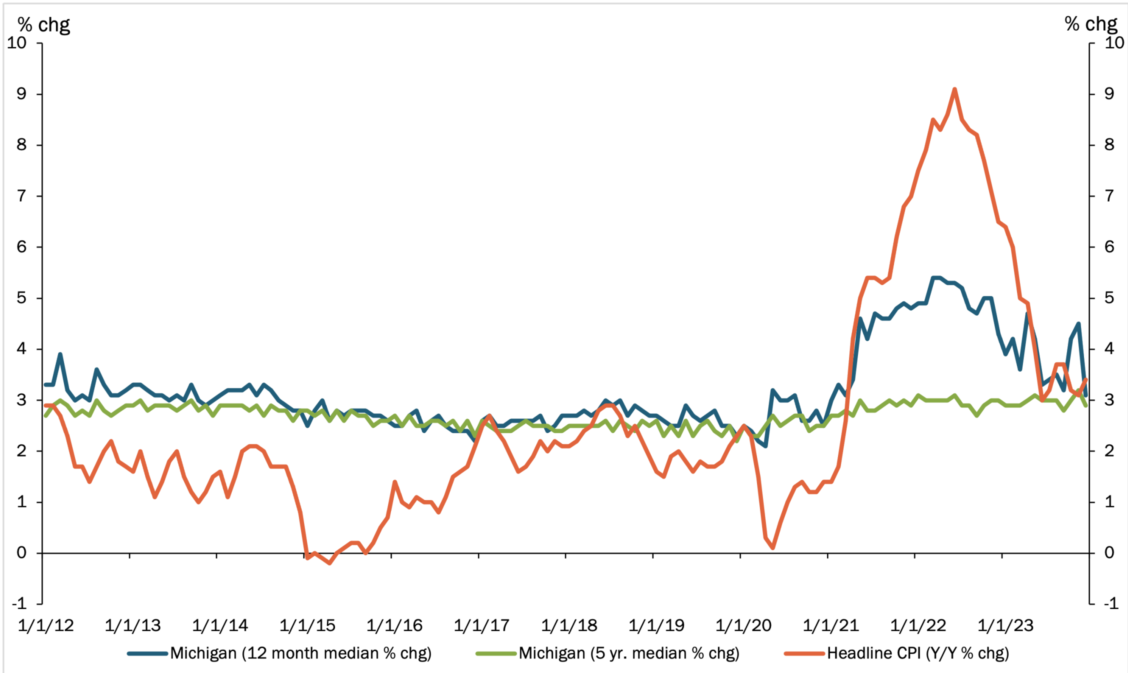

Figure 1: Michigan Survey and Headline CPI Inflation

Notes: Headline CPI inflation is measured by the 12-month change in the headline CPI.

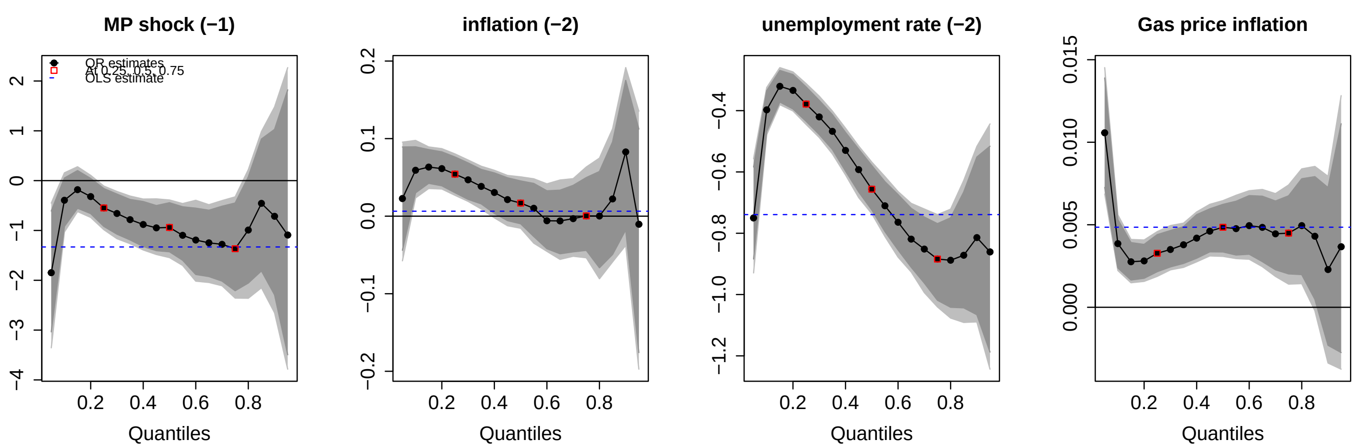

Figure 2: Coefficient estimates on macro variables across quantiles: baseline specification

Notes: The OLS estimates are the coefficient estimates of the same specification in the conditional mean regression by OLS. The bands represent the 90% and 95% confidence intervals estimated by bootstrapping.

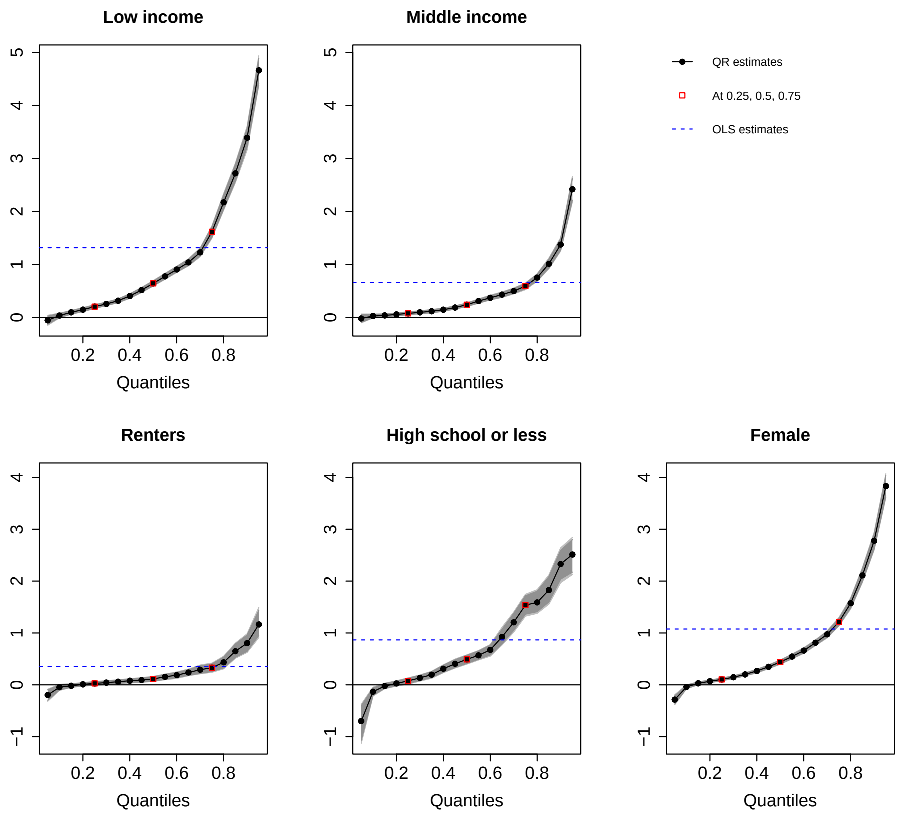

Figure 3: Coefficient estimates on demographic characteristics across quantiles: baseline specification

Notes: The base (omitted) group is a group of households with highest income quartile, more education, and homeownership and who are male and a young generation living in the West. See the notes in Figure 2 as well.

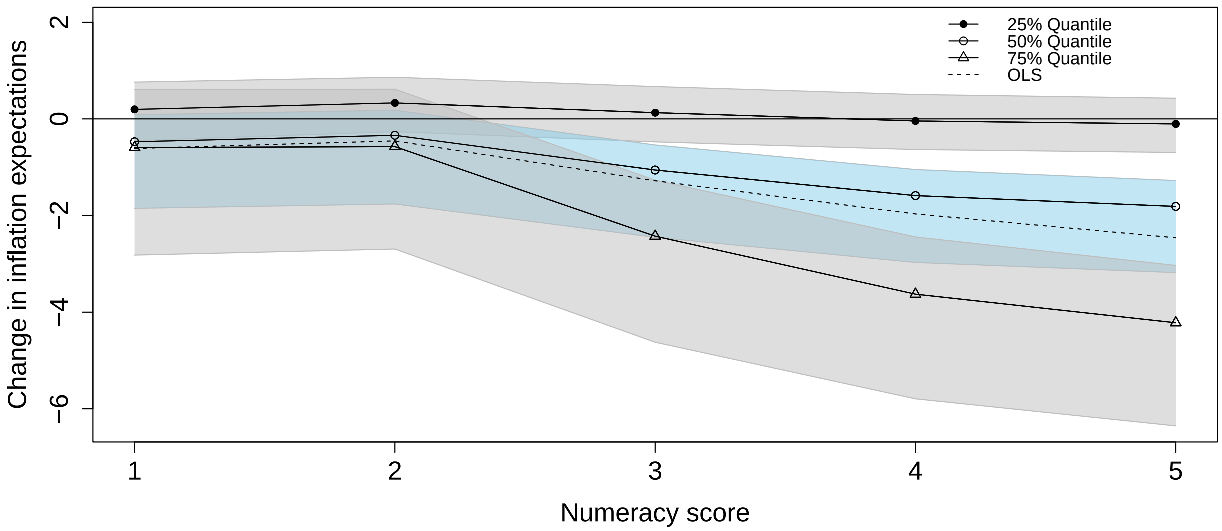

Figure 4: Estimated changes in inflation expectations over numeracy score: baseline specification
Notes: See the notes in Figure 3.

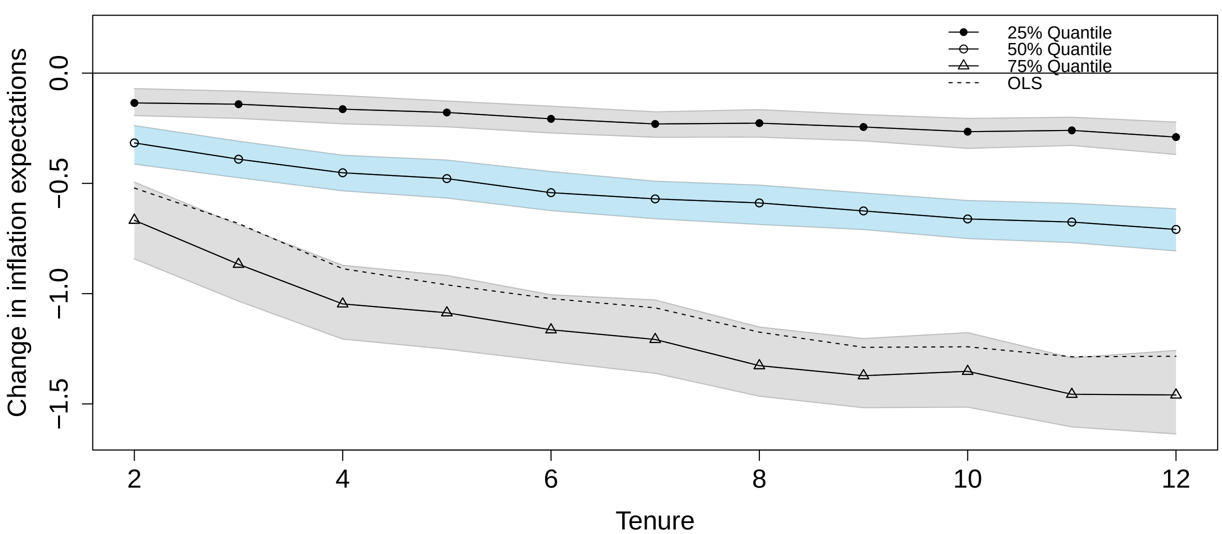

Figure 5: Estimated changes in inflation expectations over survey tenure: baseline specification
Notes: See the notes in Figure 2.

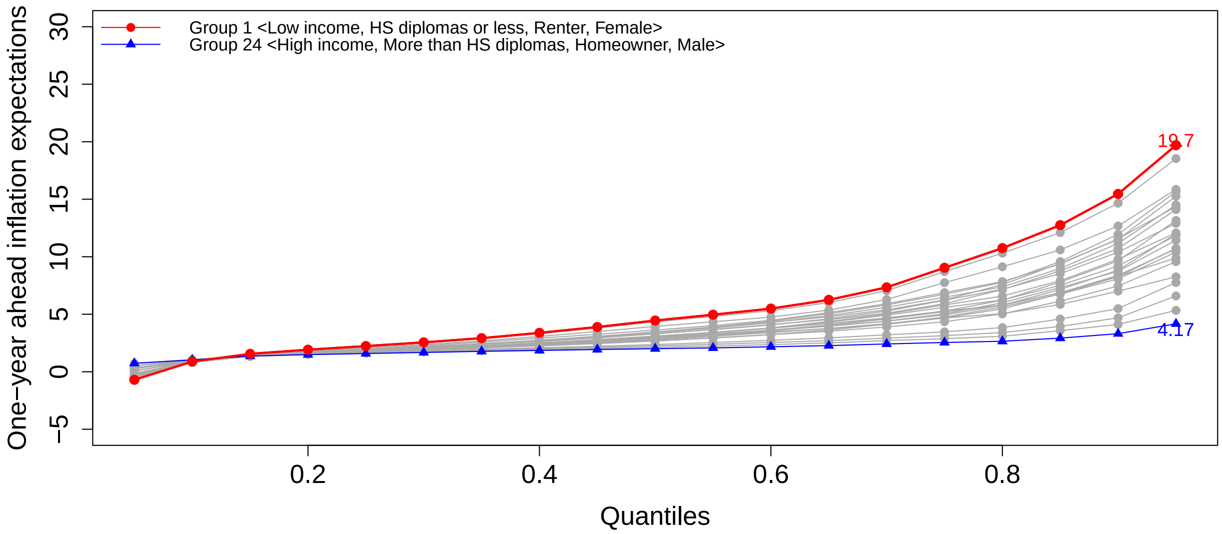

Figure 6: Predicted conditional quantiles of household inflation expectations across demographic groups: December 2019

Notes: With the orthogonalized monetary policy shock. There are 24 groups in total: 3 income groups, 2 education groups, 2 homeownership groups, 2 gender groups. For the other variables in the conditional quantile regression in (1), when computing the predicted conditional quantiles, it is assumed that the survey tenure, the numeracy score, the number of kids and adults are equal to the respective medians for each group, the age group is the young generation, and the region of primary residence is the West.

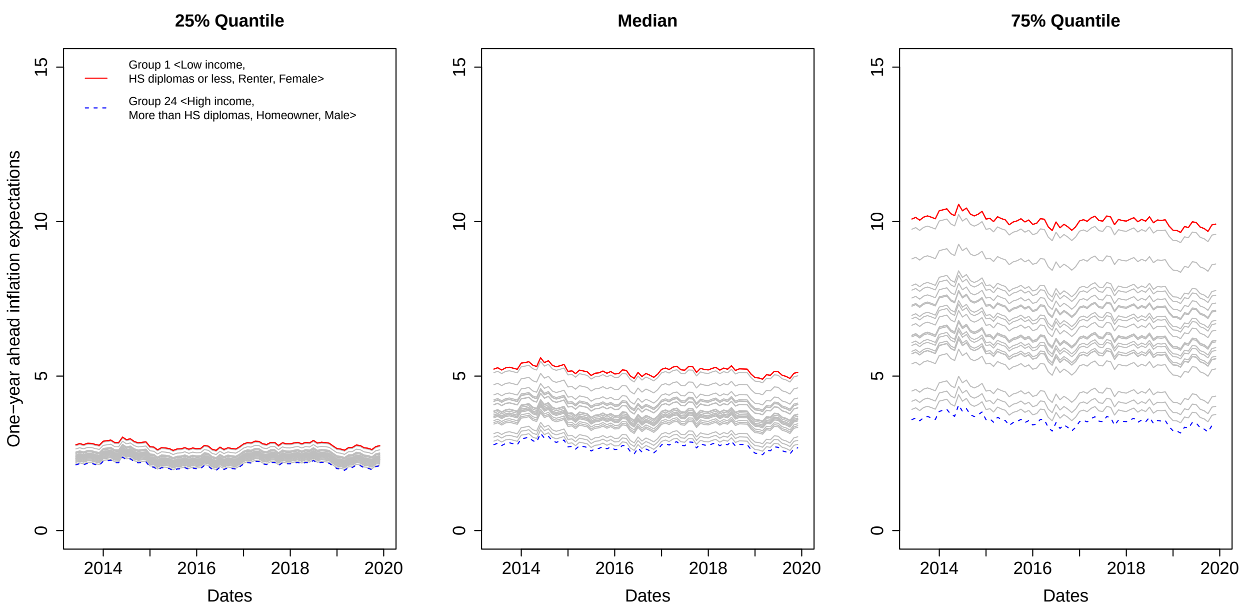

Figure 7: Predicted conditional quantiles of household inflation expectations over time: 25%, 50%, and 75% quantiles

Notes: With the orthogonalized monetary policy shock. See the notes in Figure 6 as well.

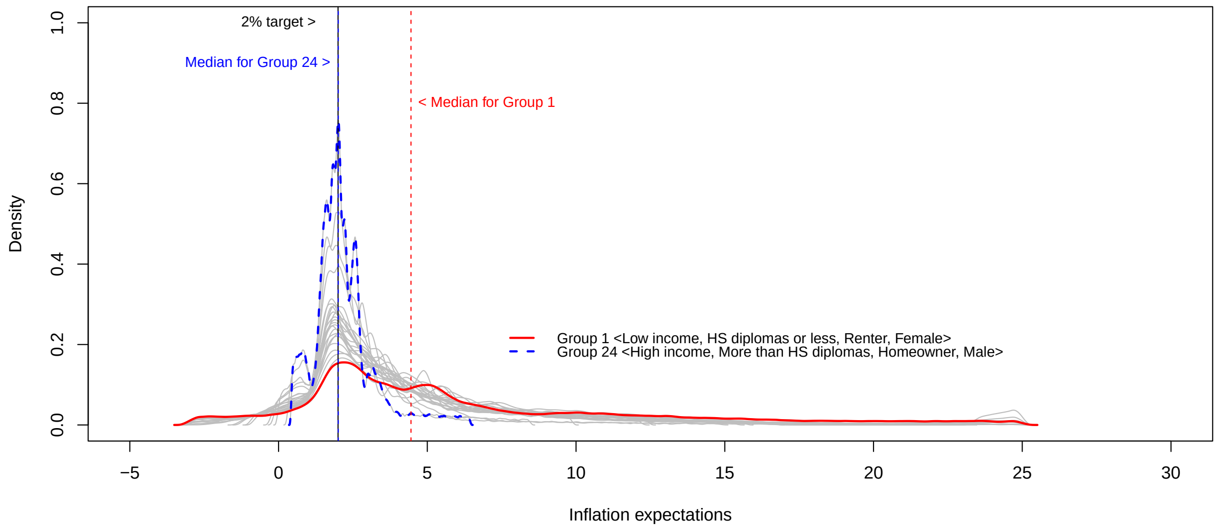

Figure 8: Predicted conditional density function of household inflation expectations across demographic groups: December 2019

Notes: With the orthogonalized monetary policy shock. See the notes in Figure 6 as well. See the online appendix on the method of computing the predicted conditional distribution.

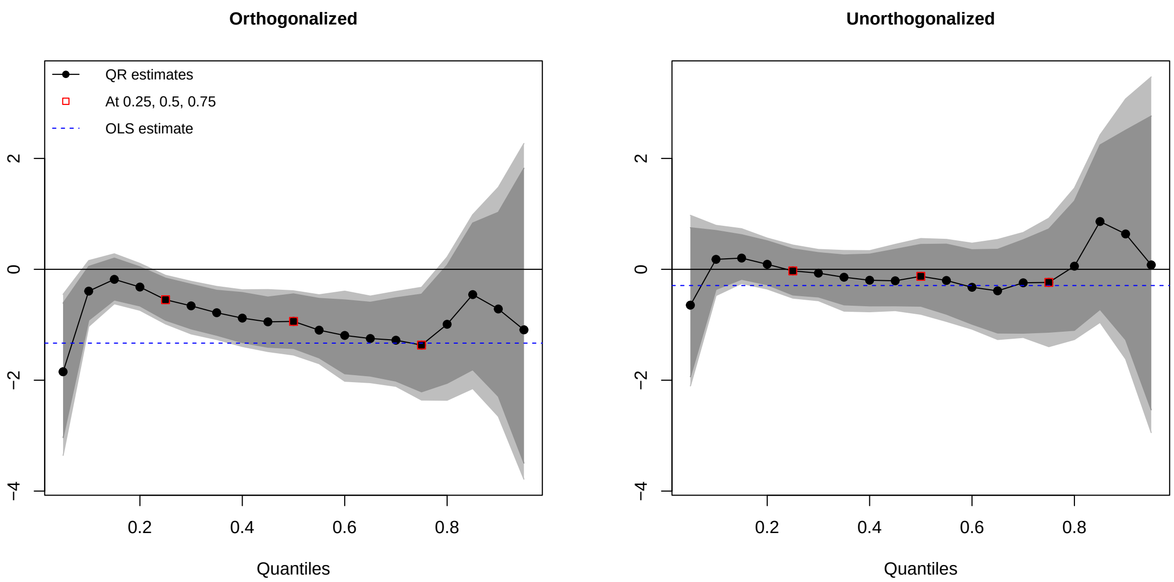

Figure 9: Coefficient estimates on the orthogonalized and unorthogonalized monetary policy shock by Bauer and Swanson (2023)

Notes: The plots present the coefficient estimates with their 90% and 95% confidence intervals of the baseline specification with the orthogonalized monetary policy shock and of an alternative specification with the unorthogonalied monetary policy shock estimated by Bauer and Swanson (2022). The alternative specification uses the same specification as the baseline except for the monetary policy shock and is estimated on the same sample. The horizontal dashed lines represent the OLS estimate of the coefficient on the monetary policy shock in each specification.

(a) On the orthogonalized monetary policy shock

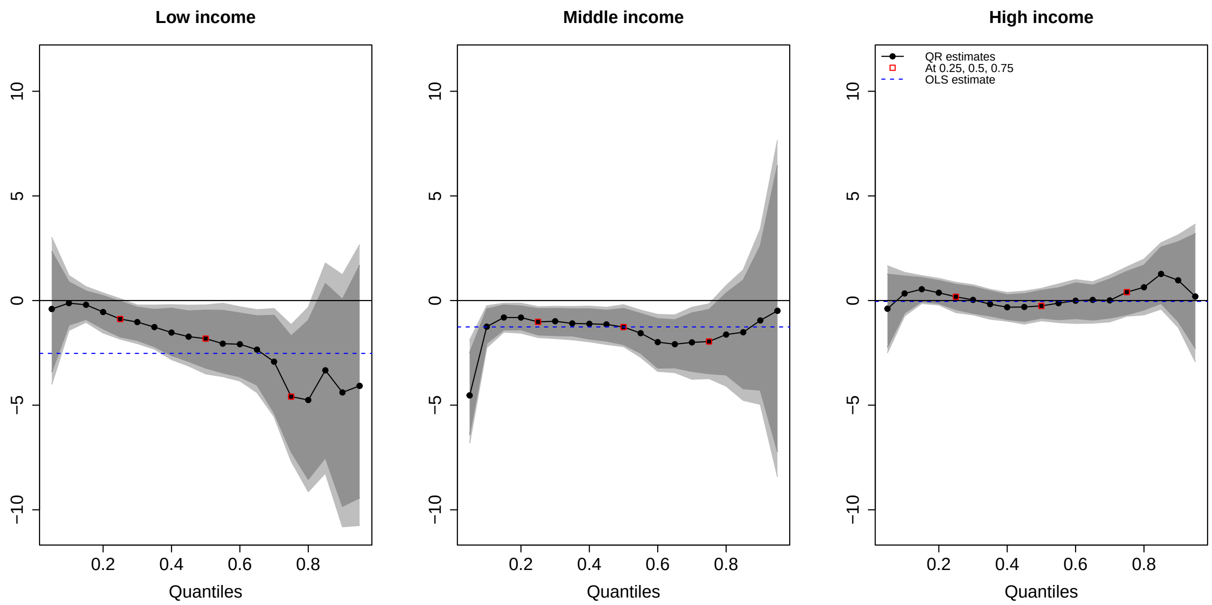

(b) On the unorthogonalized monetary policy shock

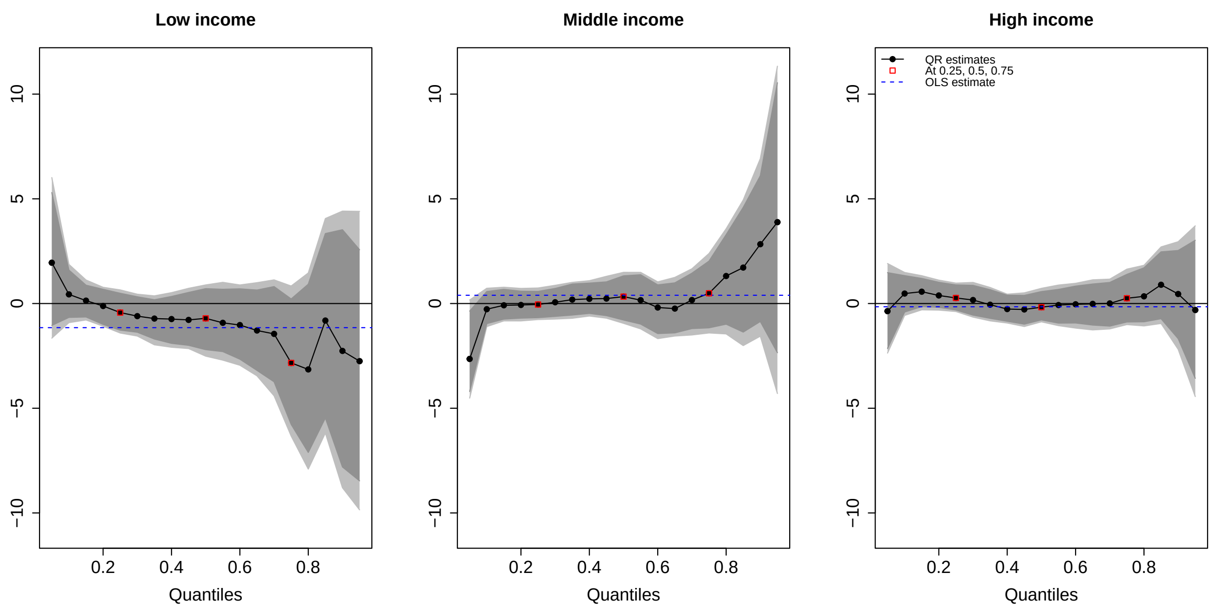

Figure 10: Coefficient estimates on the monetary policy shock across quantiles: alternative specification

Notes: Panel (a) shows the quantile coefficient on the orthogonalized monetary policy shock interacted with the household income while panel (b) shows the quantile coefficient on the unorthogonalized monetary policy shock interacted with the household income. The OLS estimates in the blue dash lines are the coefficient estimates of the same specification in the conditional mean regression by OLS. The bands represent the 90% and 95% confidence intervals estimated by bootstrapping.

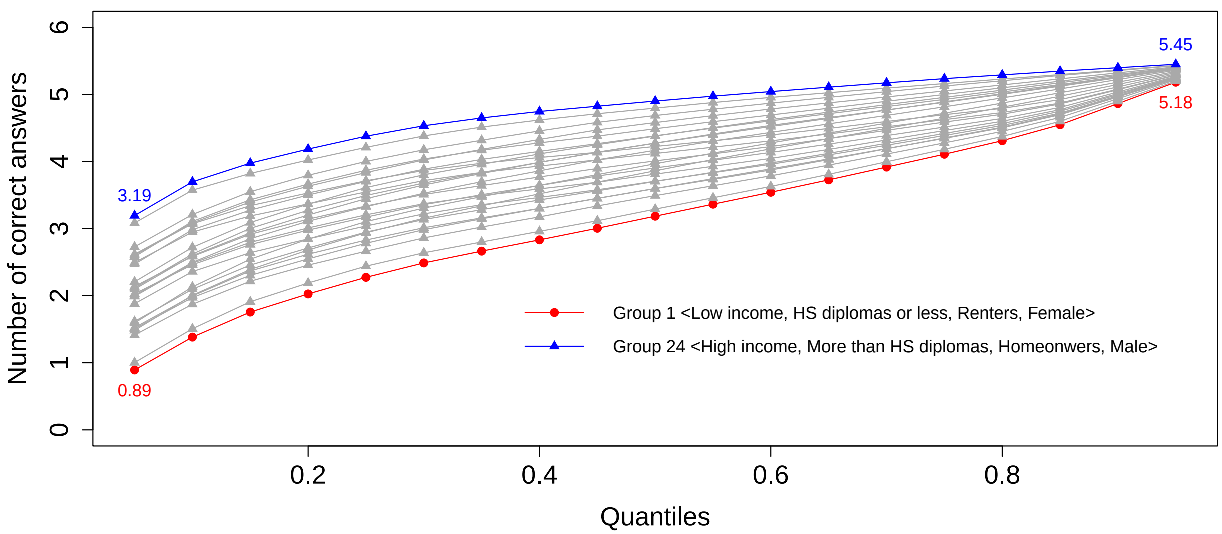

Figure 11: Predicted conditional quantiles of the numeracy score across demographic groups

Notes: There are 24 groups in total: 3 income groups, 2 education groups, 2 homeownership groups, 2 gender groups.

Online Appendix for

Heterogeneity in Household Inflation Expectations: Policy Implications

by Doh, Lee and Park

(Not for publication)

A Additional results

A.1 Additional coefficient estimates

Figure A1 reports the coefficient estimates on other demographic characteristics including the family size, age, and region of primary residence in the baseline empirical analysis.

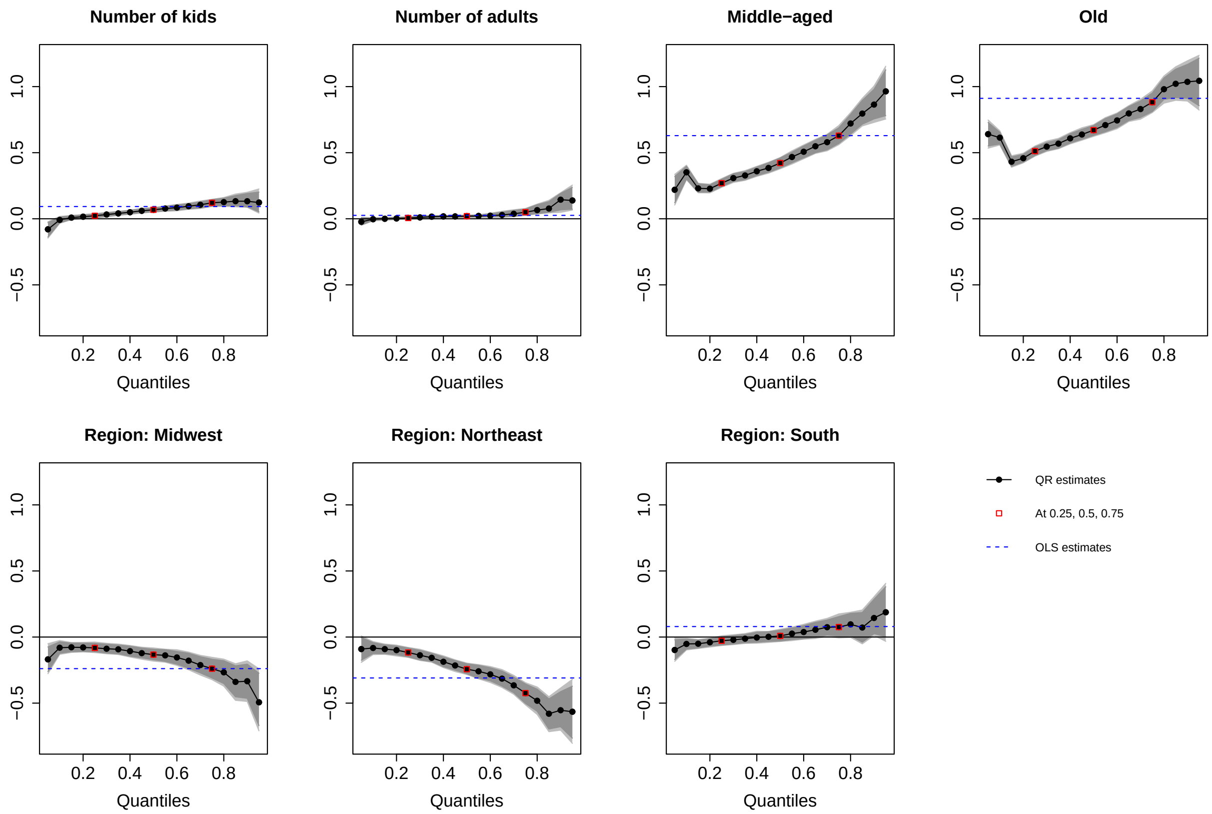

Figure A1: Coefficient estimates on other demographic characteristics across quantiles: baseline specification

Notes: The base (omitted) group is a group of households with highest income quartile, more education, and homeownership and who are male and a young generation living in the West. See the notes in Figure 2 as well.

Table A1 reports the coefficient estimates on the dummy variables for the survey tenure and the dummy variables for the numeracy score at the 25%, 50%, and 75% quantile, in the baseline empirical analysis.

Table A1: Estimation results on the survey tenure and the numeracy score: baseline specification

Quantiles	25%	50%	75%	OLS
tenure 2	**-0.135**	**-0.317**	**-0.667**	**-0.521**
	(-0.192,-0.070)	(-0.413,-0.238)	(-0.842,-0.494)	(-0.629,-0.413)
tenure 3	**-0.141**	**-0.390**	**-0.867**	**-0.682**
	(-0.205,-0.082)	(-0.474,-0.309)	(-1.034,-0.689)	(-0.791,-0.573)
tenure 4	**-0.163**	**-0.452**	**-1.047**	**-0.887**
	(-0.230,-0.102)	(-0.534,-0.372)	(-1.206,-0.872)	(-0.997,-0.777)
tenure 5	**-0.178**	**-0.478**	**-1.087**	**-0.960**
	(-0.243,-0.127)	(-0.566,-0.394)	(-1.252,-0.918)	(-1.073,-0.848)
tenure 6	**-0.207**	**-0.542**	**-1.164**	**-1.023**
	(-0.271,-0.150)	(-0.623,-0.447)	(-1.308,-1.005)	(-1.137,-0.908)
tenure 7	**-0.231**	**-0.570**	**-1.208**	**-1.064**
	(-0.291,-0.176)	(-0.660,-0.490)	(-1.361,-1.029)	(-1.180,-0.948)
tenure 8	**-0.227**	**-0.589**	**-1.327**	**-1.175**
	(-0.290,-0.166)	(-0.686,-0.508)	(-1.465,-1.152)	(-1.293,-1.056)
tenure 9	**-0.244**	**-0.625**	**-1.372**	**-1.243**
	(-0.307,-0.188)	(-0.709,-0.543)	(-1.517,-1.204)	(-1.364,-1.123)
tenure 10	**-0.266**	**-0.661**	**-1.352**	**-1.241**
	(-0.341,-0.206)	(-0.750,-0.578)	(-1.515,-1.177)	(-1.365,-1.116)
tenure 11	**-0.260**	**-0.675**	**-1.456**	**-1.286**
	(-0.328,-0.200)	(-0.768,-0.590)	(-1.604,-1.291)	(-1.417,-1.156)
tenure 12	**-0.290**	**-0.709**	**-1.459**	**-1.284**
	(-0.369,-0.222)	(-0.806,-0.615)	(-1.636,-1.258)	(-1.429,-1.139)
score 1	**0.196**	**-0.473**	**-0.594**	**-0.613**
	(-0.451, 0.761)	(-1.852, 0.085)	(-2.819, 0.606)	(-1.119,-0.106)
score 2	**0.331**	**-0.340**	**-0.574**	**-0.456**
	(-0.264, 0.861)	(-1.761, 0.179)	(-2.693, 0.614)	(-0.942, 0.030)
score 3	**0.128**	**-1.058**	**-2.427**	**-1.280**
	(-0.476, 0.668)	(-2.450,-0.540)	(-4.622,-1.269)	(-1.760,-0.800)
score 4	**-0.043**	**-1.588**	**-3.628**	**-1.967**
	(-0.634, 0.505)	(-2.971,-1.049)	(-5.793,-2.444)	(-2.446,-1.489)
score 5	**-0.105**	**-1.812**	**-4.220**	**-2.461**
	(-0.695, 0.428)	(-3.181,-1.276)	(-6.353,-3.033)	(-2.940,-1.982)

Notes: See the notes in Table 2.

A.2 Comparison of inflation expectations and TIPS

Figure A2 shows median forecasts from the low-income group and the high-income group together with the five-year inflation compensation measure from TIPS as well as the five-year expected inflation from D'Amico et al. (2018).

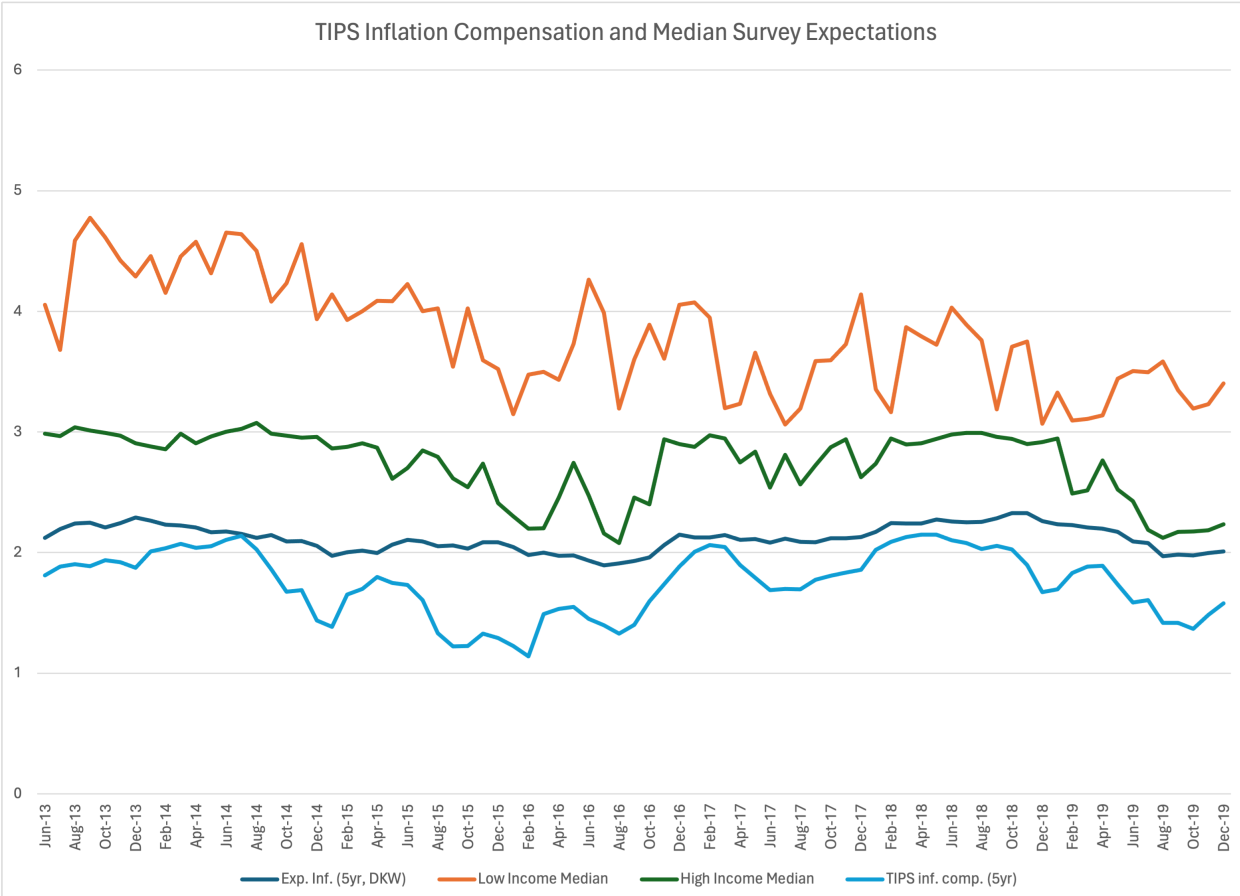

Figure A2: Median forecasts of households belonging to different income groups in each point of time are plotted together with the five-year inflation compensation measure from TIPS and the five-year expected inflation from D'Amico et al. (2018).

A.3 Robustness checks for jittering

We tried two cases: 1) uniform distribution with mean 0 and length equal to the numeracy-score-specific IQR, 2) normal distribution with mean 0 and standard deviation equal to the numeracy-score-specific standard deviation. Both cases generate similar results as the baseline analysis except for the coefficient estimates on the numeracy score. Therefore, we only report the coefficient estimates on the macroeconomic variables and the numeracy score here. The coefficient estimates on the numeracy score slightly got bigger in terms of the magnitude than those in the baseline analysis.

Figures A3 and A4 report the coefficient estimates of the first case, where the random noise for jittering is drawn from a uniform distribution with mean 0 and length equal to the numeracy-score-specific IQR.

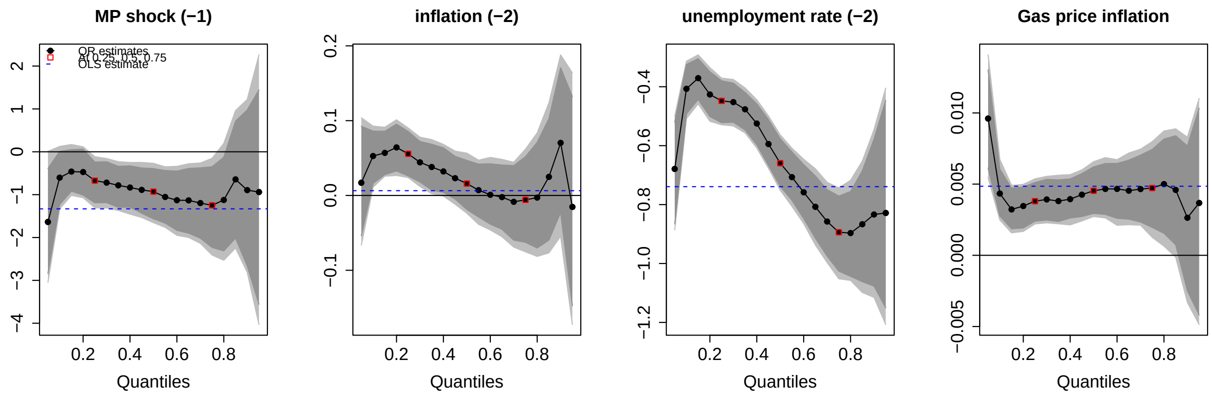

Figure A3: Coefficient estimates on macro variables across quantiles: baseline specification with jittering using numeracy score-specific jittering, case 1

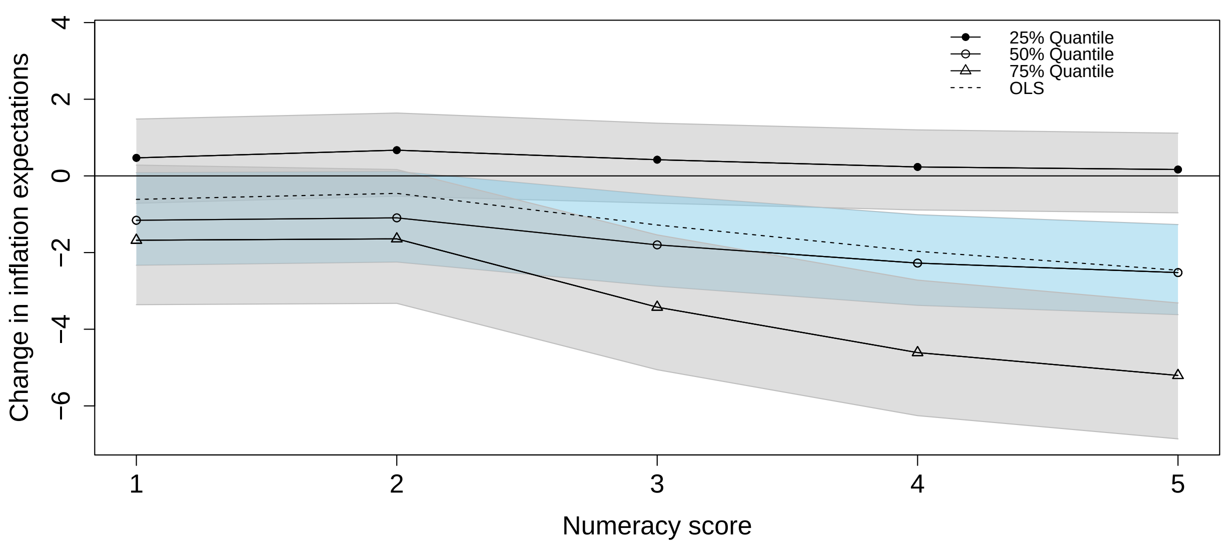

Figure A4: Estimated changes in inflation expectations over numeracy score: baseline specification with jittering using numeracy score-specific jittering, case 1

Figures A5 and A6 report the coefficient estimates of the first case, where the random noise for jittering is drawn from a normal distribution with mean 0 and standard deviation equal to the numeracy-score-specific standard deviation.

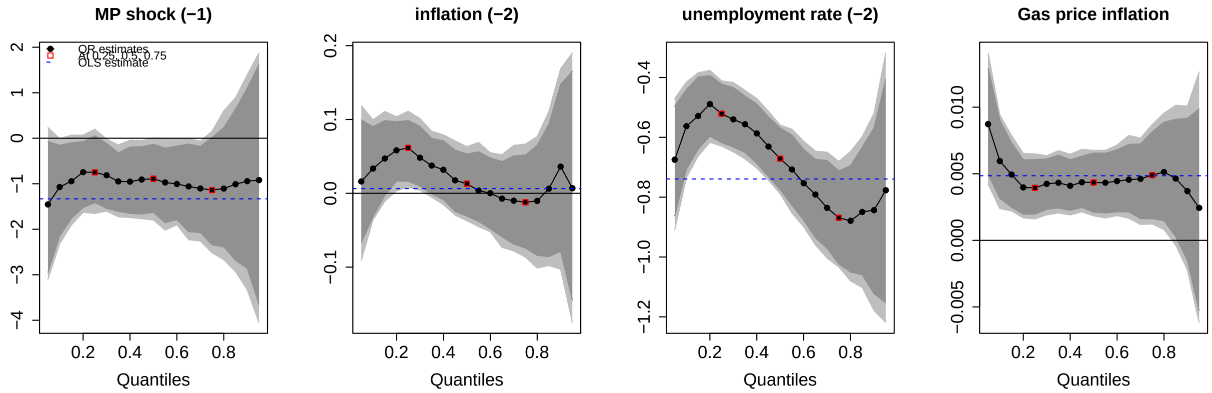

Figure A5: Coefficient estimates on macro variables across quantiles: baseline specification with jittering using numeracy score-specific jittering, case 2

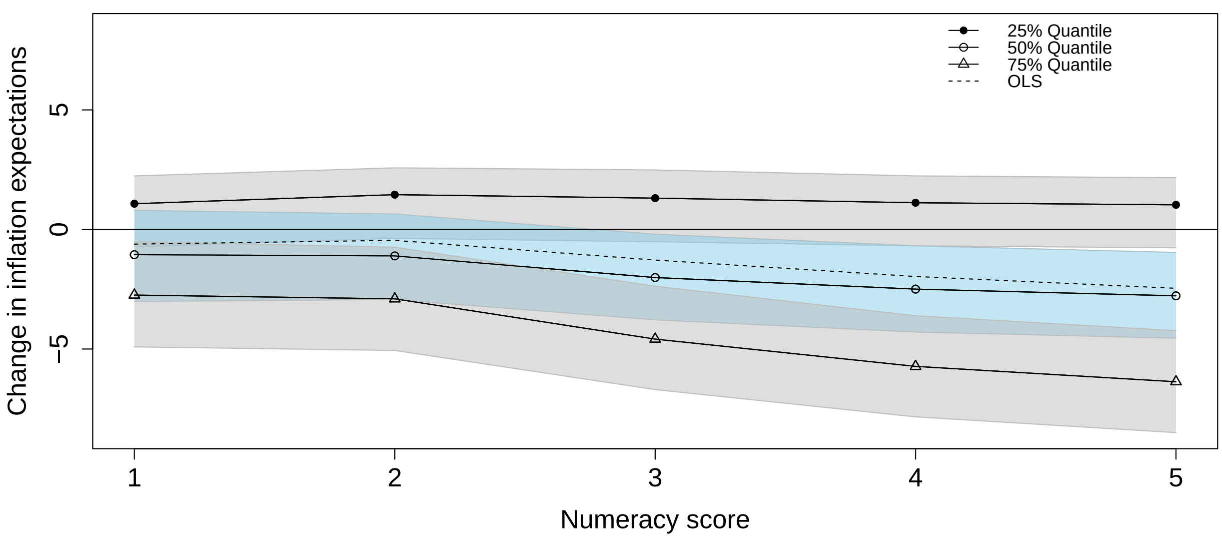

Figure A6: Estimated changes in inflation expectations over numeracy score: baseline specification with jittering using numeracy score-specific jittering, case 2

A.3.1 Robustness checks for the unconventional monetary policy shock

Here, we present the result of the robustness exercise where the conditional quantile regression is estimated on the ZLB sample from December 2008 through December 2015. To have a sufficiently large sample, we use the MSC data instead of the SCE data. The coefficient estimates on the orthogonalized and unorthogonalized monetary policy shock are reported in Figure A7. For reference, we also report the coefficient estimates on the full sample from September 1992 through December 2019.

The details on estimation of quantile regressions with the MSC data are explained in the next section.

(a) On the full sample (September 1992 through December 2019)

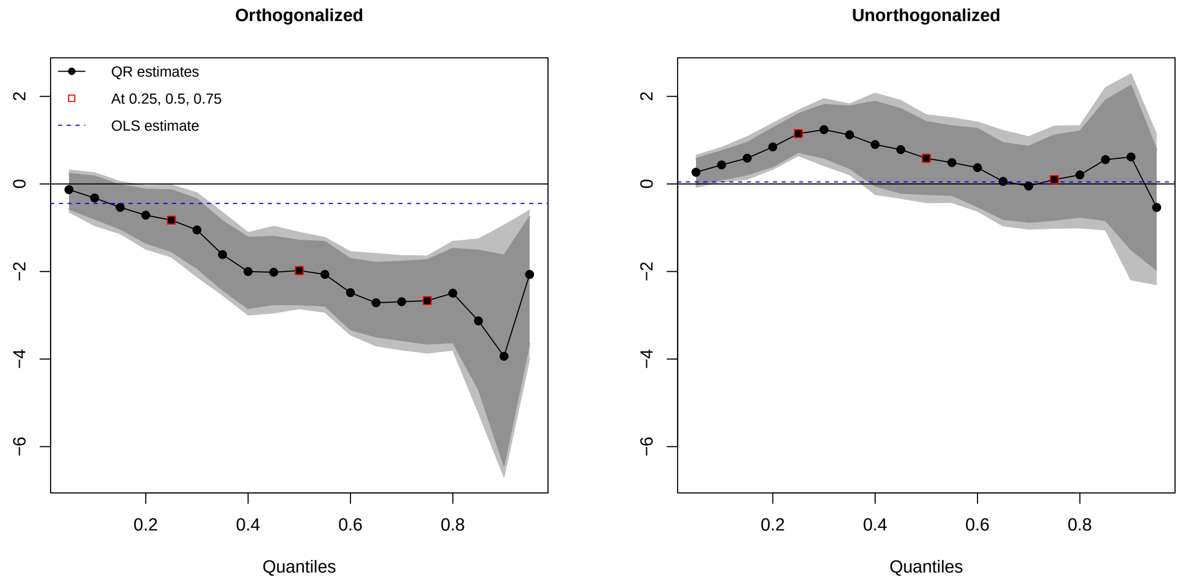

(b) On the ZLB sample (December 2008 through December 2015)

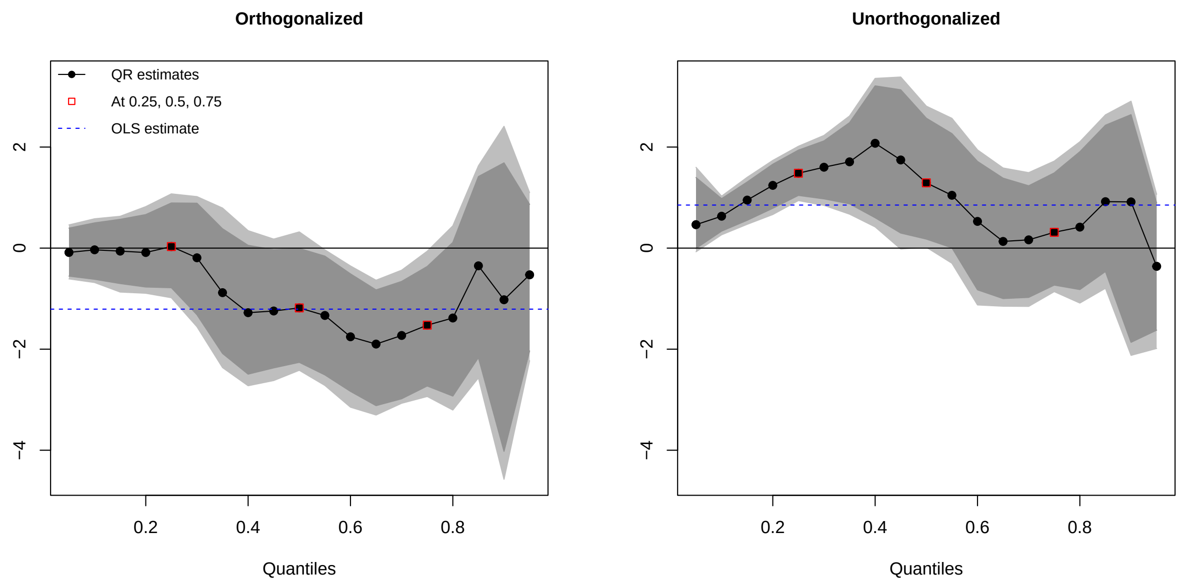

Figure A7: Coefficient estimates on the orthogonalized and unorthogonalized monetary policy shock by Bauer and Swanson (2023): full sample vs. ZLB sample

B Results from the MSC

Though the SCE has a bigger survey panel and provides more information on the survey respondents such as their numeracy, it only started in June 2013. In contrast, while the MSC has a smaller survey panel and lacks information on numeracy of the survey respondents, the MSC has a longer sample. We check the robustness of our empirical results by estimating the conditional quantile regression on the MSC data in this section. The MSC data dates back to 1978 but we use the data since September 1992 because of the availability of some data. MSC only records integer inflation expectations so we jitter the whole sample of inflation expectations.

The specification of the conditional quantile regression is the same as the baseline specification for the SCE data except for a few variables. MSC interviews households at most twice so we replace the survey tenure with a dummy variable for repeat survey. MSC lacks information on numeracy of the survey respondents, which is thus missing in the alternative regression. We use the same specification on the full sample (1992M9 through 2019M12) and on the ZLB sample (2008M12 through 2015M12).

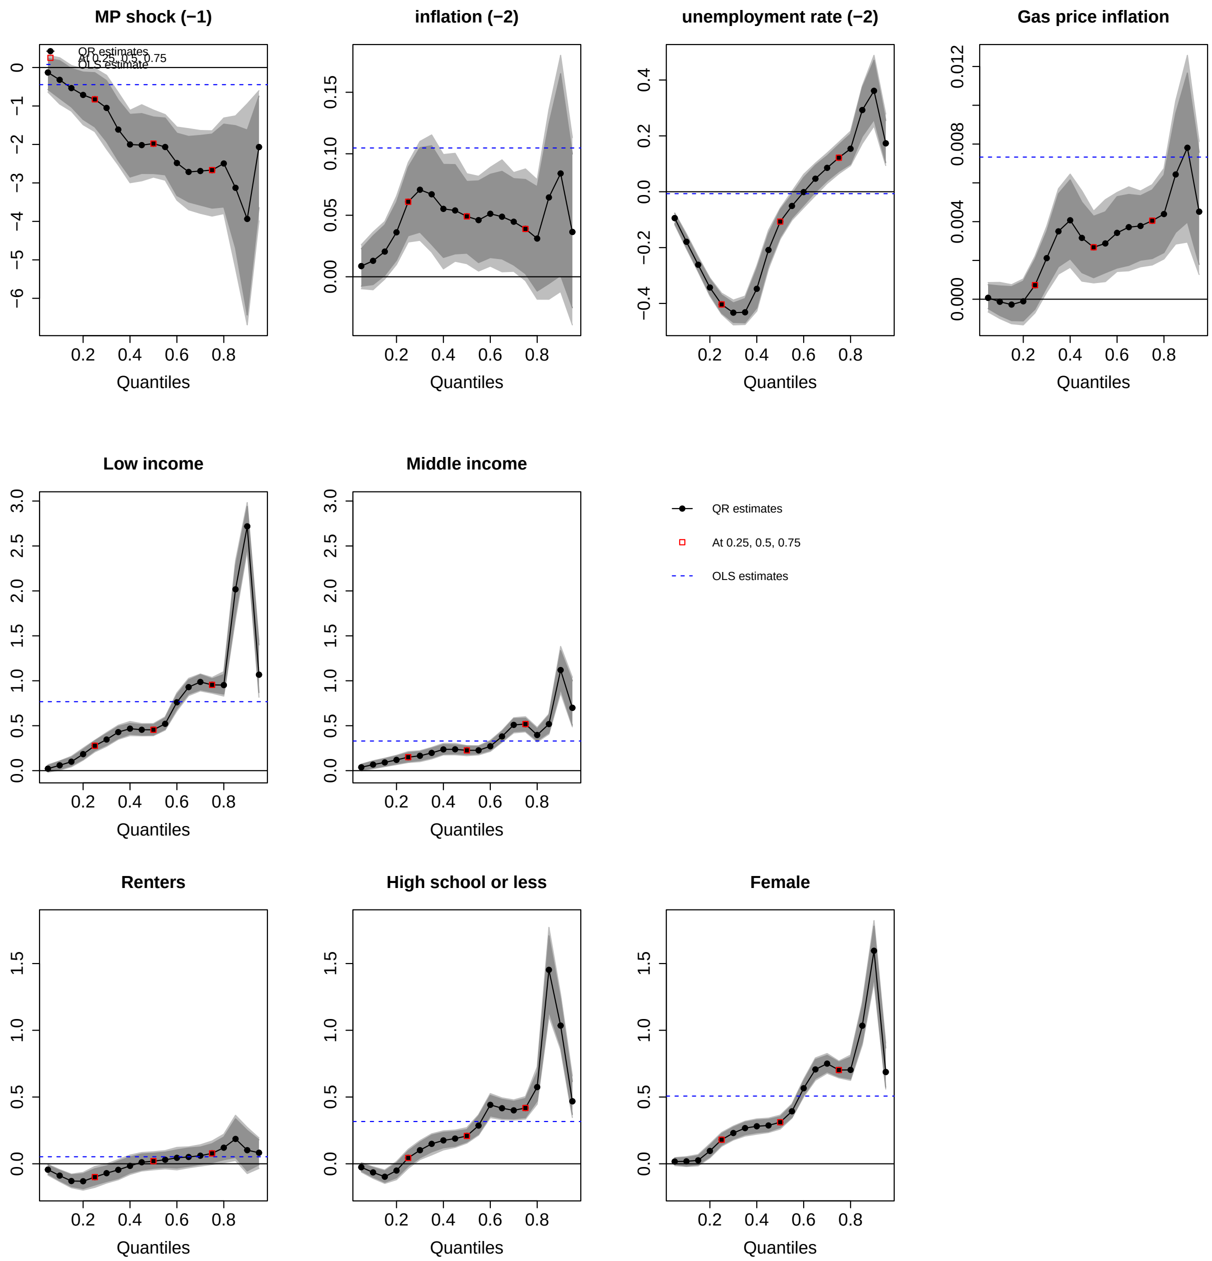

Figure A8: [Michigan Survey of Consumers] Coefficient estimates across quantiles

Notes: The base (omitted) group is a group of households with highest income quartile, more education, and homeownership and who are male and a young generation living in the West. The OLS estimates are the coefficient estimates of the same specification in the conditional mean regression by OLS. The bands represent the 90% and 95% confidence intervals estimated by boostrapping. The sample period is from September 1992 through December 2019.

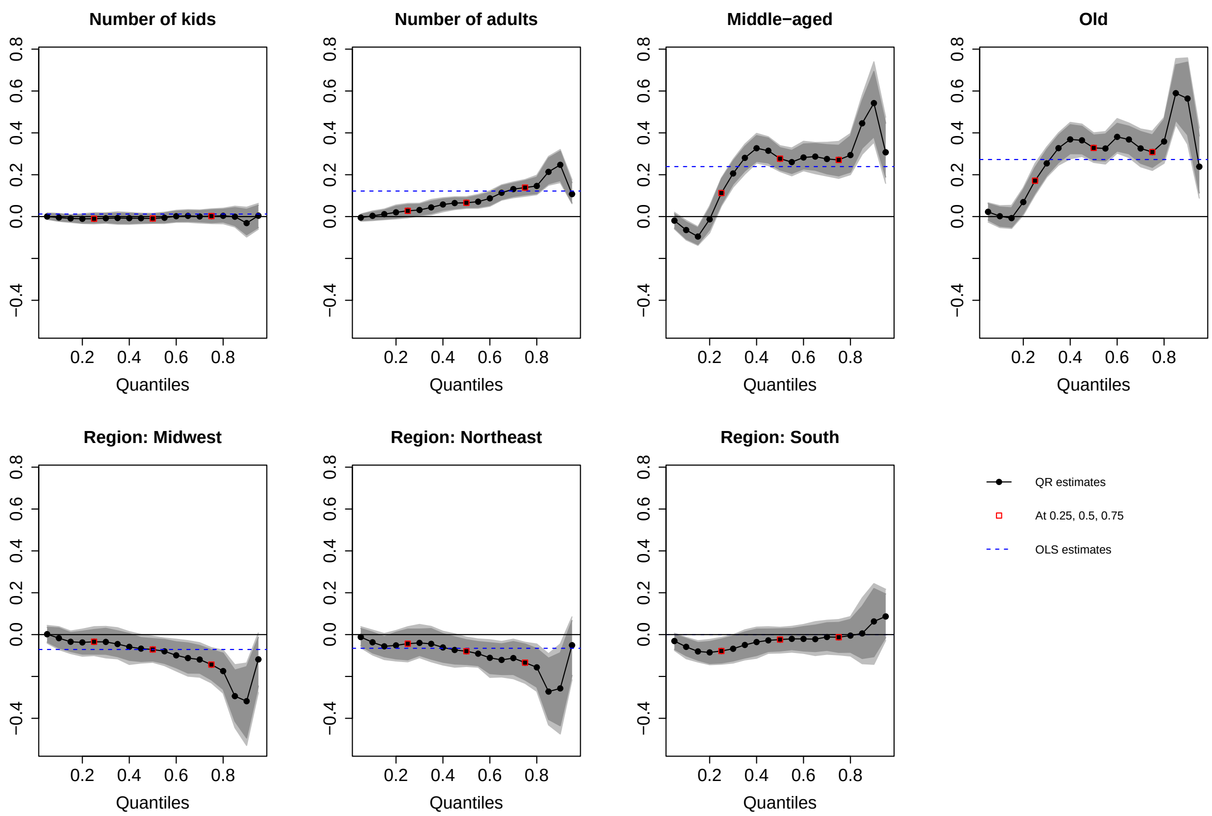

.

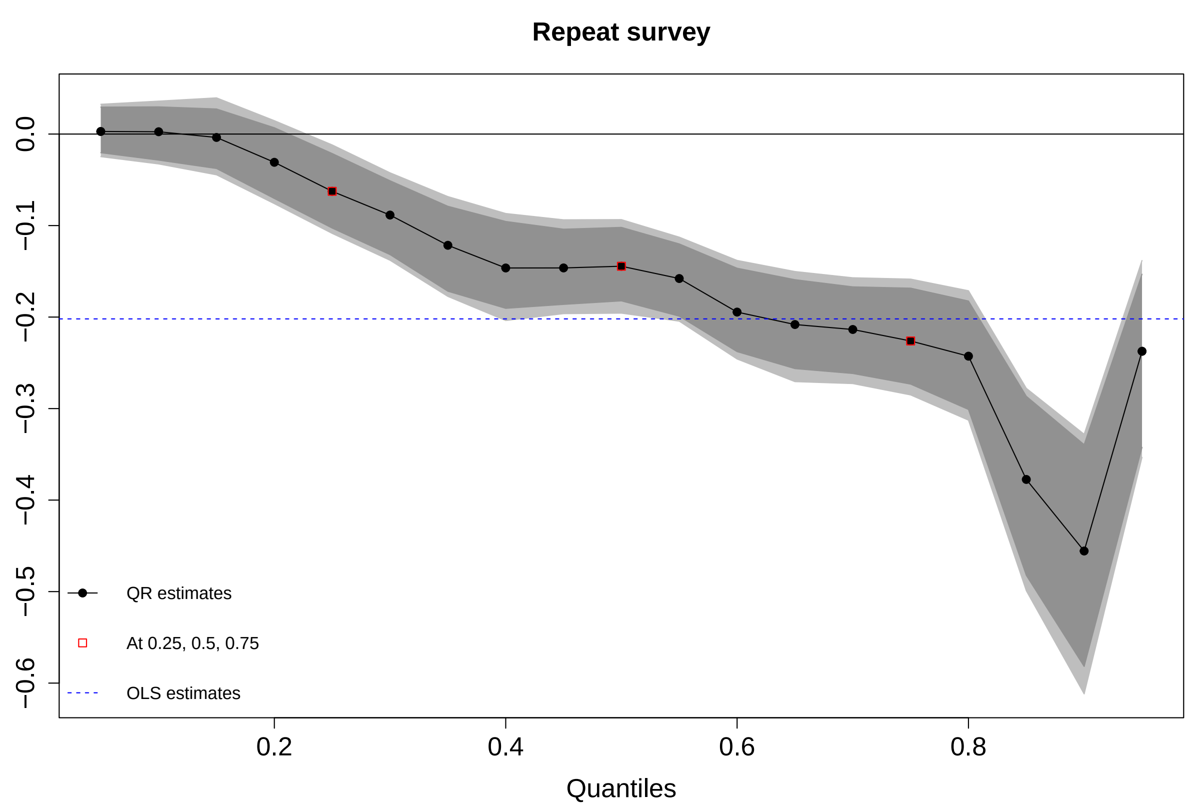

Figure A8: [Michigan Survey of Consumers] Coefficient estimates across quantiles (continued)

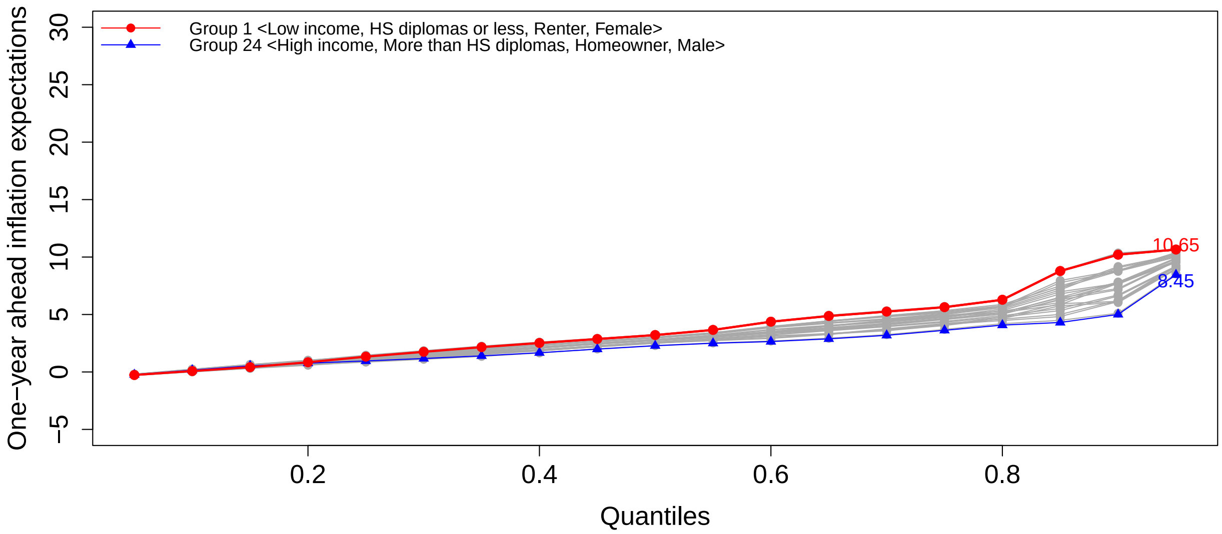

Figure A9: [Michigan Survey of Consumers] Predicted conditional quantiles of household inflation expectations across demographic groups: December, 2019

Notes: With the orthogonalized monetary policy shock. There are 32 groups in total: 4 income quartiles, 2 education groups, 2 homeownership groups, 2 gender groups. To compute the predicted conditional quantiles, it is assumed that the number of kids and adults are equal to their medians for each group, respectively, the age group is the young generation, the region of primary residence is the West, and the survey respondent participates for the first time.

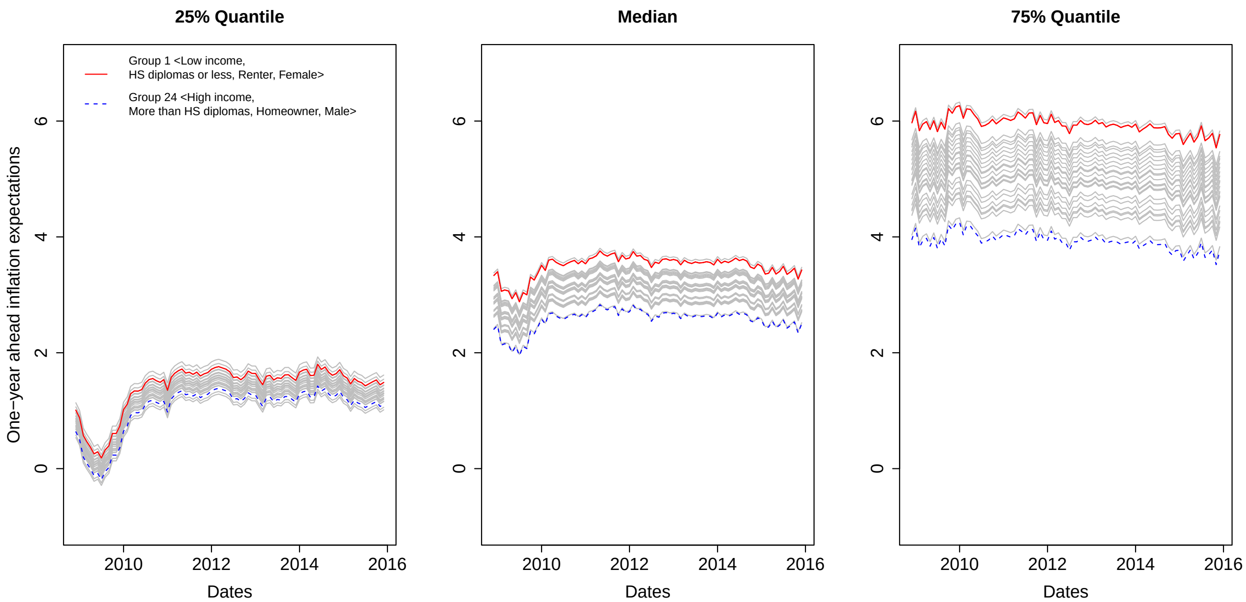

Figure A10: [Michigan Survey of Consumers] Predicted conditional quantiles of household inflation expectations over time: 25%, median, and 75% quantiles

Notes: With the orthogonalized monetary policy shock. There are 32 groups in total: 4 income quartiles, 2 education groups, 2 homeownership groups, 2 gender groups. See the notes in Figure 6 as well.

It is striking that the coefficient estimates on the demographic and socio-economic characteristics based on the MSC data in Figure A8 are qualitatively similar to those on the SCE data. Coefficients on time-varying aggregate variables are also broadly similar except for the fact that the effect of the lagged inflation or the oil price inflation becomes quantitatively more significant. For income, homeownership, education, gender, the number of kids and adults, and age, the coefficient estimates are small for the lower quantiles but increase in the probability to compute the quantile. We also find that the second-time survey respondents have lower inflation expectations than the first-time respondents and the difference is bigger in the upper quantiles. This pattern we observe for the coefficient estimates, which is similar to the one we observe on the SCE data, leads to conditional distributions of household inflation expectations that are qualitatively similar to the ones based on the SCE data. The lower quantiles are very close among groups but the upper quantiles are substantially different among groups.

C A parsimonious model of inflation expectations

Our empirical analysis suggests that a monetary policy tightening orthogonalized with respect to information before a FOMC meeting lowers inflation expectations especially for the low income and less education group. We provide a parsimonious model of household inflation expectations similar to the one in Reis (2021) that can be consistent with the finding.

Let $\pi^e_{h,g,t}$ be the inflation expectation of a household h that belongs to the group g at time t. The household-specific inflation expectations is determined by fundamental inflation (e.g., the core measure of consumer price inflation), the group-specific inflation bias and the group-specific treatment effect of a monetary policy shock as follows:

$$\pi^e_{h,g,t} = \pi^\star_t + \pi_{g,t} + \theta_g(z_t - \pi^\star_t - \pi_{g,t}) + \epsilon_{g,h,t}, \tag{3}$$

where $\pi^\star_t$ and $\pi_{g,t}$ denote fundamental inflation and the group-specific inflation bias, and z_t is the policy rate. It is assumed that $\epsilon_{g,h,t} \sim \mathcal{N}(0, \sigma_{g,t})$. Therefore, the group-dependent heteroskedasticity in $\epsilon_{g,h,t}$ generates the heterogeneous distribution of household inflation expectations across demographic and socio-economic groups. The coefficient θ_g measures the monetary policy treatment effect on the household inflation expectation, which can vary across different groups. We let $\theta_g < 0$ for $g \in \{L, H\}$ so that a rise in the policy rate puts downward pressures on inflation expectations for the households in all the groups.

The policy rate is determined as follows:

$$z_t = \gamma\pi^\star_t + \delta x_{t-} + \epsilon^{\text{orth}}_{mp,t}, \tag{4}$$

where x_{t-} is information available in the financial markets before a FOMC meeting. It implies that the policy rate is adjusted in response to changes in fundamental inflation and pre-FOMC information available to the financial markets. It is assumed that $\text{Corr}\,(x_{t-}, \pi^*_t) = 0$ so information in x_{t-} is relevant for the rate decision above and beyond fundamental inflation.

We further assume that x_{t-} is positively correlated with the group-specific inflation bias for group L but orthogonal to the group-specific inflation bias for group H: $\text{Corr}\,(x_{t-}, \pi_{g,t}) > 0$ if $g = L$ but $= 0$ if $g = H$. For example, a commodity price indicator might be more relevant for L group's inflation expectation but not for H group.

The monetary policy shock $\epsilon^{\text{orth}}_{mp,t}$ is orthogonal to $\pi^{\star}_t$ and x_{t-}. We consider an alternative measure of the monetary policy shock, consistent with the unorthogonalized monetary policy shock series in Bauer and Swanson (2023). The unorthogonalized shock $\epsilon^{\text{unorth}}_{mp,t}$ still has some predictability by information available in the financial markets before an FOMC meeting so it can be written as

$$\epsilon^{\text{unorth}}_{mp,t} = \delta x_{t-} + \epsilon^{\text{orth}}_{mp,t}. \tag{5}$$

With the model specified above, we can obtain two different OLS estimates of θ_g depending on the measure of monetary policy treatment used. The population of each OLS estimate can be expressed as follows:

$$\widehat{\theta}^{\text{orth}}_g = \frac{E\left\{\left[\theta_g \epsilon^{\text{orth}}_{mp,t} + ((1-\theta_g(1-\gamma))\pi^{\star}_t + (1-\theta_g)\pi_{g,t})\right]\epsilon^{\text{orth}}_{mp,t}\right\}}{E(\epsilon^{\text{orth}}_{mp,t})^2} = \theta_g,$$

$$\widehat{\theta}^{\text{unorth}}_g = \frac{E\left\{\left[\theta_g \epsilon^{\text{unorth}}_{mp,t} + ((1-\theta_g(1-\gamma))\pi^{\star}_t + (1-\theta_g)\pi_{g,t})\right]\epsilon^{\text{unorth}}_{mp,t}\right\}}{E(\epsilon^{\text{unorth}}_{mp,t})^2} = \theta_g + \frac{\delta(1-\theta_g)E(\pi_{g,t}x_{t-})}{E(\epsilon^{\text{unorth}}_{mp,t})^2} > \theta_g.$$

Hence, the model explains the treatment effect can be negative in response to an orthogonalized monetary policy shock but close to zero in response to an unrothogonalized monetary policy shock.