Households' Information Rigidities

Student: Natalia Rostova

Supervisor: Prof. Luigi Paciello

Einaudi Institute for Economics and Finance (EIEF)

LUISS Guido Carli

RoME Program

Abstract

This paper studies the differences in information rigidities across income groups of households using the NY Fed Survey of Consumer Expectations. Following Coibion and Gorodnichenko (2012) approach, I consider conditional responses of inflation forecast errors to oil price and Gertler and Karadi (2015) monetary policy shocks. Since the survey does not specify to households a particular price index to forecast, I consider two cases: 1) CPI inflation and 2) income-specific inflation to account for the heterogeneity of households' consumption baskets. Studying the second case, I find that high-income households are subject to the lowest information rigidity, while middle-income households make the largest forecast errors. These results differ from the case of CPI inflation only for oil price shocks, for which the low-income households adjust their forecasts better than the other income groups.

2018/2019 Academic Year

1 Introduction

There are numerous papers in which researchers study the properties of agents' beliefs about macroeconomic variables. In this paper, I focus on households' inflation expectations. They have high importance for monetary policy for several reasons. First, inflation expectations affect saving and consumption decisions, since households make these decisions based on the real interest rate, and the latter is defined by inflation expectations. Moreover, inflation expectations are important for the effectiveness of forward guidance. The aim of forward guidance (Bernanke, 2013) is to provide transparency about future policies, shape agents' expectations and reduce the uncertainty about monetary policy. Thus, the effectiveness of forward guidance depends on to what extent agents' expectations are forward-looking. As Coibion et al. (2018) point out, the question about how economic agents form their expectations is still open and debatable.

Many researchers find the evidence that different economic agents - consumers, firms and professional forecasters - deviate from the full-information rational expectations (FIRE) model in a sense that they are inattentive to new information. The deviations are explained either by the cost of acquiring new information (Mankiw and Reis (2002) and Reis (2006)), or by agents' limited ability to monitor information (Sims (2003)) and observe the true values of economic variables (Woodford (2003)). For example, Bordo and Michaels (2007), Coibion (2010), Andrade and Bihan (2010) study inflation forecasts of professionals, Carroll (2003), Binder (2017) focus on households' expectations, Coibion et al. (2018) - on firms' forecasts. Coibion and Gorodnichenko (2012) study inflation forecasts of all the three groups of agents.

However, the majority of empirical papers, and particularly papers about households' expectations, focus on the measures of information stickiness derived from the aggregate forecasts. There are several studies documenting forecast errors of different demographic groups. For instance, Johannsen (2014) finds that low-income households disagree more about future inflation. Armantier et al. (2012) run an experiment evaluating how agents update their inflation forecasts when they are provided with information about past inflation rates. They find that low-income, low-education, and low financially literal households are more sensitive to new information than other groups of households. Bryan and Venkatu (2001) document that agents with high income perceive and expect lower inflation in comparison to low-income individuals. In addition, the authors show that even holding all the demographic characteristics constant, agents have highly heterogeneous expectations.

In this paper, based on the methodology of Coibion and Gorodnichenko (2012), I consider two questions. First, does the degree of information rigidity vary across income groups of households, and, second, what is the role of heterogeneity of households' consumption baskets? The latter is important to consider because surveys usually do not specify to households a particular price index to forecast. Households at different income percentiles have different consumption baskets, and, therefore, economic shocks may affect inflation rates faced by households heterogeneously. For example, Cravino et al. (2018) show that consumption baskets of high-income households are less volatile, and after a monetary policy shock the CPI of high-income households respond three times less than the CPI of middle-income households. Therefore, if when households are asked to forecast the inflation rate, they predict the changes in prices of their consumption baskets instead of the aggregate inflation rate, the estimates of the degree of information rigidity derived from the aggregate inflation rate can be overstated.

There are two main branches of models with information rigidities. The first is sticky information model of Mankiw and Reis (2002) and Reis (2006), and noisy information models of Woodford (2003), Morris and Shin (2002), Sims (2003), Patton and Timmermann (2006). According to the sticky information model, agents face fixed costs of acquiring new information and, as a result, update their forecasts infrequently. In this model, the probability of not updating the forecast is interpreted as the degree of information rigidity.

In noisy information models, agents are not able to observe the true values of fundamentals and form their expectations from noisy signals. Their forecast is a weighted average of new information and old beliefs. In that model, a weight which agents place on the old information each period is considered as the degree of information rigidity.

Coibion and Gorodnichenko (2012) show that both models make the same prediction that the average forecast responds more slowly to a shock than the forecasted variable.

In contrast, under the null of the full-information rational expectations (FIRE), agents track continuously new information and update forecasts immediately. Once a shock occurs, agents should adjust their forecast by the same amount as the forecasted variable, so that the response of the average forecast error is zero. Departing from these predictions, the authors suggest a way to test the FIRE hypothesis and estimate the degree of information rigidity.

I use micro dataset from the survey of Consumer Expectations conducted by the New York Fed to divide households into three groups depending on the level of pre-tax income and calculate the mean forecast error for each group. Then, based on the methodology of Coibion and Gorodnichenko (2012), I evaluate the degree of information rigidity for each group of households from the response of the mean forecast error to two types of shocks: oil price shocks and monetary policy shocks. Oil price shocks are calculated as in Hamilton (1996). To identify monetary policy shocks, I use high-frequency identification suggested by Gertler and Karadi (2015). That approach is suitable for the period considered (2013-2019) since it allows to capture shocks to forward guidance.

I consider two cases. First, I estimate the degree of information rigidity under the assumption that all agents forecast the same object - the aggregate inflation rate. Second, I construct income-specific inflation rates and repeat the analysis under the assumption that agents forecast changes in the prices of the consumption basket of their income group ("income-specific" inflation rate).

The results suggest that all groups of households are subject to information rigidities, but the degree varies across groups. At first, I find that low-income households adjust their forecast better, while the largest errors are made by high-income households. Replacing the aggregate inflation rate by the income-specific inflation rates, I find that the conclusion changes: the differences in the degree of information rigidity across income groups lessen, and high-income households are subject to the lowest degree of information rigidity among three groups.

However, controlling for the heterogeneity of consumption baskets seems not important in case of monetary policy shocks. Overall, I find that in both cases the null of FIRE cannot be rejected for the high-income group. In contrast, middle-income households are found to be completely inattentive to monetary policy shocks, the result which seems to be puzzling.

The rest of the paper is organized as follows: Section 2 describes the data and identification of aggregate shocks. Section 3 overviews the methodology to measure information rigidity. Section 4 presents the estimates of the degree of information rigidity for each income group of households. Section 5 concludes.

2 Data and Identification of Aggregate Shocks

2.1 Forecast Errors

To construct the "income-specific" mean forecast errors I use the publicly available micro dataset from the Survey of Consumer Expectations conducted by Federal Reserve Bank of New York. The survey is monthly and representative of the US nation. The survey data at households level is available from June 2013 to March 2018. The number of households is around 1300 each month, and each household can be present in the survey for up to 12 months. Among the numerous questions, the survey asks households to forecast the annual inflation rate and give information about a household's income, education, gender, etc. The question about the future inflation rate[1] is:

"What do you expect the rate of inflation/deflation to be over the next 12 months? Please give your best guess".

[1] In the questionnaire, this question is labeled as "Q8v2part2"

I divide households into three groups by the level of income: 1) low-income: if household's pre-tax annual income is less than \$50k, 2) middle-income: if a household earns from \$50k to \$100k, and 3) high-income: annual income is higher than \$100k. Then for each group, I calculate the average forecast error which is defined at time t as the actual (realized) inflation rate at time t minus the mean forecast of that inflation rate forecasted by households at time t-12.

The mean forecast errors for the period from June 2013 to March 2018 are presented in Fig.1. The average forecast error of low-income households is -5.35 p.p., of medium- and high-income households are -3.13 and -2.06 p.p. respectively. Standard deviations of the mean forecast errors are 1.55, 1.32 and 1.10 p.p. respectively. The plot shows that low-income households persistently make the largest forecast errors when predicting the one-year ahead inflation rate. High-income households make the smallest errors.

Figure 1: Mean Forecast Errors by Income Group

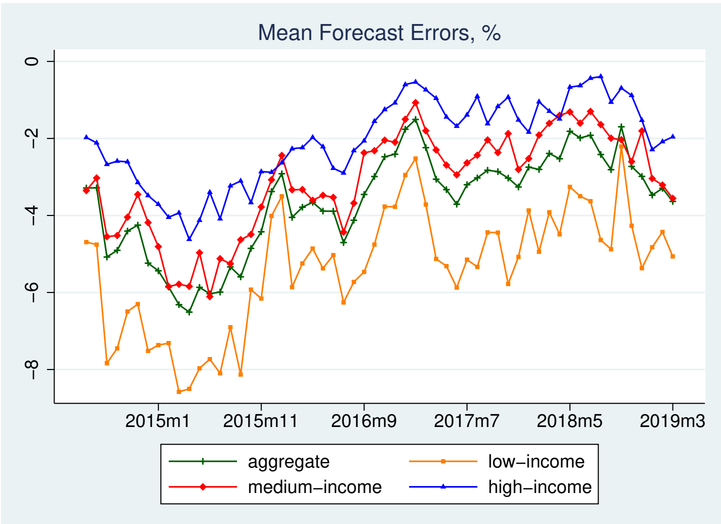

Note. The forecast error at time t is calculated as the actual one-year inflation rate minus the forecast of that inflation predicted at time t-12. Households are grouped by the level of income: 1) low-income: if reported pre-tax annual income is less than \$50k, 2) medium-income: \$50k-\$100k, 3) high-income: more than \$100k. Data Sources: households' inflation forecasts - New York Fed Survey of Consumer Expectations, the inflation rate - FRED.

2.2 Construction of Income Specific CPIs

To construct income specific CPIs, I merge two data sets: data on average consumer expenditures from the U.S. Consumer Expenditure Survey (CEX) and disaggregated CPI series published by the U.S. Bureau of Labor Statistics (BLS).

The first survey (CEX) is annual and contains information about average amounts of money spent by households on various products including detailed food items, housing, utilities, services, apparel items, transportation, entertainment, education, and healthcare. The total amount of items is 95. The survey is representative of the entire U.S. civilian noninstitutional population (which represents more than 98% of the total U.S. population). Expenditure shares are calculated for different households' groups by income, education, age, and other demographic characteristics. I focus on consumer expenditure shares grouped by the level of income.

The data is of yearly frequency and available until 2017. I use data for 2015-2017 since the data for 2018-2019 is not available yet and CEX definition of income groups before 2015 does not correspond to the grouping defined in Section 2.1. When calculating expenditure share for each item, I make several adjustments since not all of the CEX items are included in CPI. The adjustments are taken from Cravino et al. (2018) and reported in the Appendix.

I assume that the structure of the consumption basket of each income group is constant during the period considered (June 2013 - March 2019). For each income group and for each expenditure item I take the average share spent on that item in 2015-2017.

Table 1 presents some selective statistics on expenditure items by income group. For example, low-income households spend higher share of their income on food at home, utilities, while high-income households spend more on food away from home, entertainment, education and purchases of new vehicles than low- and medium-income households.

The second data set from BLS provides monthly disaggregated inflation rates. It contains detailed information on price changes for more than 300 categories. I match these categories with the ones in CEX data. Then, for each income group, I calculate the

Table 1: Statistics on expenditure shares

Item	Low	Medium	High
Food at home	9.0%	7.8%	6.8%
Food away from home	5.8%	6.7%	7.1%
Utilities	9.7%	8.5%	6.6%
Cars and trucks, new	2.8%	3.7%	4.5%
Education	2.2%	1.8%	3.9%
Entertainment	5.3%	6.1%	6.9%

Note. The table presents selective statistics on expenditure items by income group. Low income: less than $50k, medium income: $50-100k, high income: more than $100k. Each expenditure share is calculated as the average of 2015-2017 years.

income-specific inflation rate by weighing each item's inflation rate by the income share spent on that item by that income group.

The income-specific inflation rates are presented in Fig. 2. Table 2 presents means and standard deviations of the time series. Low-income households experienced the highest level of inflation - 1.52%. Standard deviations are 0.85, 0.91 and 0.72 respectively. These results correspond to the findings of Cravino et al. (2018) that prices of consumption baskets of middle-income consumers are more volatile than the ones of low- and high-income households. The authors find that volatility of inflation along the income distribution is hump-shaped, and the lowest volatility of prices is observed for the households at the top of the income distribution.

Table 2: Means and standard deviations of income-specific inflation rates, %

Income group	Mean	St. deviation
Low	1.52	0.85
Medium	1.37	0.91
High	1.39	0.72

Note. The table reports the means and standard deviations of income-specific inflation rates for the period June 2014 - March 2019.

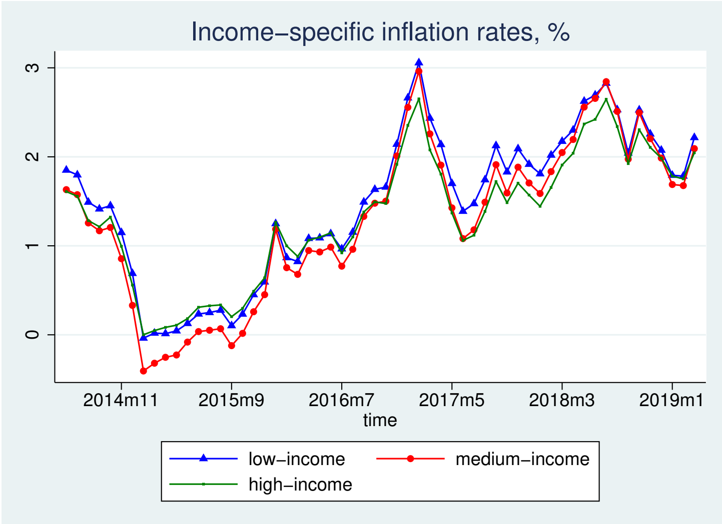

Figure 2: *Income-specific annual inflation rates for three groups of households (June 2014 to March 2019). The time series are constructed from two data sources: Consumer Expenditures Survey and disaggregated inflation rates for various prices categories published by Bureau of Labor Statistics.*

2.3 Identification of Oil Price Shocks

Similarly to Coibion and Gorodnichenko (2012), I use the following procedure to estimate oil price shocks: for each month I calculate the difference between the oil price in that month and the highest price in the last 12 months. If the difference is positive, this is considered as an oil price shock. If not, the oil price shock is set to be 0 in that month. I take logarithms of oil prices. Therefore, by construction, oil price shocks could be either positive or zero. That approach was proposed by Hamilton (1996).

Fig.3 presents the estimated oil price shocks for the period from June 2014 to March 2019. Oil shocks are calculated based on prices of West Texas Intermediate from FRED.

2.4 Identification of Monetary Policy Shocks

To identify monetary policy shocks, Coibion and Gorodnichenko (2012) use the approach of Christiano et al. (1999). Namely, the authors estimate the structural VAR model under recursive identification scheme. That approach is not suitable for the later period (June 2013 - March 2019) since, as Gurkaynak et al. (2005) point out, once during the crisis the short term interest rates reached the zero lower bounds, the Federal Reserve increased

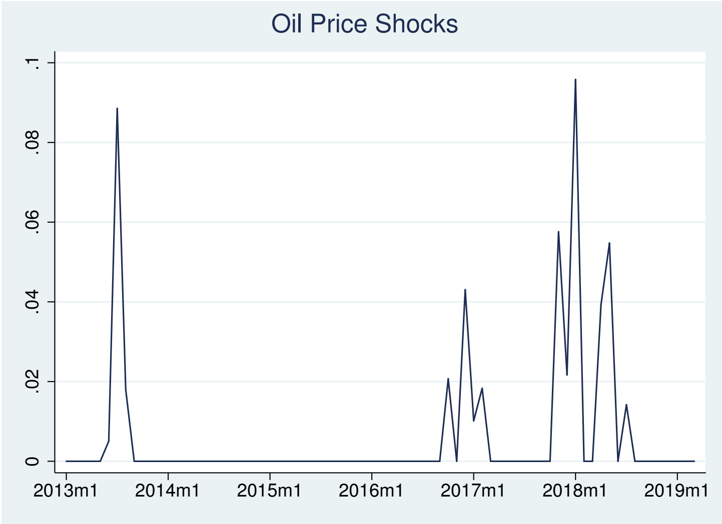

Figure 3: *Oil price shocks are identified as in Hamilton (1996). Oil prices (West Texas Intermediate) data source: FRED.*

its reliance on forward guidance to affect market beliefs about the future path of interest rates. Therefore, considering responses of agents' beliefs to monetary policy shocks, one should use an approach which captures shocks to forward guidance. As Gertler and Karadi (2015) argue, the high-frequency identification based on changes in future prices around FOMC announcements could capture these effects.

In addition, the high-frequency approach has an advantage (Stock and Watson, 2016) that monetary policy shocks are identified in a model-free framework. Moreover, in the VAR models monetary policy shocks are usually identified on monthly data, and, therefore, it is assumed that within a month the policy instrument responds to all the other variables in the VAR but not the opposite. Therefore, identified on a monthly frequency, shocks could be not fully exogenous to economic and financial variables, i.e. they may contain endogenous part of monetary policy.

The idea of the identification scheme is to use changes in prices of fed funds or eurodollar future contracts around the FOMC announcements as an unexpected part of monetary policy changes since market prices reflect all the beliefs of market participants about Fed policy.

Among the set of instruments, I use the prices of three month ahead monthly fed funds futures (hereafter FF4). In contrast to the current months fed funds futures (FF1), contracts on FF4 capture the effects of forward guidance since these contracts expire at a later date in the future (Gurkaynak et al., 2005).

Ideally, shocks are identified in a 30-minutes window around the FOMC announcements (Gertler and Karadi, 2015). However, because of data unavailability, I use daily changes in the federal funds futures rate, namely, the difference between the first and the last prices of a trading day. Daily data is also used by, for example, Kuttner (2001) and Hamilton (2009). That could be done for the following reasons.

First, the announcement is usually made at 2.15 p.m., and researches use 30-minutes window from 2.05 pm to 2.35 p.m. The US trading day ends at 03.00 p.m., therefore, the closing time is not far from the end point of 30-min. window (Kuttner, 2001). Second, Hamilton (2009) reviews the time-series properties of daily changes in prices of fed funds futures at 1- to 3-month horizons and concludes that they provide an excellent measure of markets expectation of changes in Fed policy. Also, he shows that the accuracy of fed funds futures as a measure of market expectations has increased over time (probably, in part because fed funds changes have become more modest).

Gertler and Karadi (2015) suggest the following steps to identify monthly monetary policy shocks:

- Calculate monetary policy surprises on each announcement day:

$$mp_t = \frac{D_1}{D_1 - d_1}(ff4_t - ff4_{(t-\Delta t)}) \tag{1}$$

  where D_1 - number of days in a month, d_1 - the date of announcement, $ff4$ - fed funds future rate (FF4).

  Monetary policy surprises are scaled by the factor of $\frac{D_1}{D_1-d_1}$ to account for the timing of an announcement in a month.

- for each day of the month, cumulate surprises on any FOMC days during the last month (Gertler and Karadi (2015) assume that shocks are persistent for one month).

For example, on March 20 sum all monetary policy surprises starting from February 20.

- Finally, average monthly surprises across days of the month.

Fig. 4 shows the estimated monetary policy shocks for the period of June 2014 - March 2019. I calculate shocks for the earlier period as well (from 2000 to 2012) to compare the results with the monetary policy shocks obtained by Gertler and Karadi (2015), and find that correlation is equal 0.96.

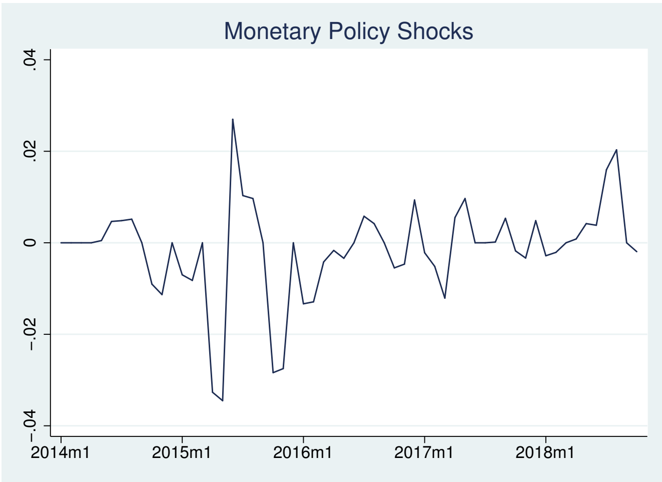

Figure 4: *The Figure presents monetary policy shocks for the period January 2014 - March 2019. Shocks are identified as in Gertler and Karadi (2015) using data on three-month ahead Fed Funds Future Rates (FF4). Future prices data source: Bloomberg.*

3 The Methodology to Measure Information Rigidity

I apply the methodology of Coibion and Gorodnichenko (2012) to analyze the responses of the mean forecast errors to shocks for each group of households. First, under the model of full-information rational expectations, agents continuously track economic variables and immediately incorporate news into their forecasts. Therefore, under the FIRE, the response of the mean forecast error to any shock should be zero.

Second, if the null hypothesis is rejected, then one could estimate the degree of information rigidity as the speed of convergence of the mean forecast error to zero over time.

There are two main types of models of information rigidity. The first is the sticky information model of Mankiw and Reis (2002) and Reis (2006). The second type of models is noisy information models such as Woodford (2003), Sims (2003), Morris and Shin (2002), Patton and Timmermann (2006).

To demonstrate how to test the presence of information rigidities in the survey data, I make an overview of the sticky information model of Mankiw and Reis (2002). Coibion and Gorodnichenko (2012) show that noisy information models predict similar patterns of the mean forecast error responses as the sticky information model, and the degree of information rigidity could be estimated in the same way.

According to the sticky information model, agents face a cost of acquiring new information and, because of that, update their forecasts infrequently. The degree of information rigidity is measured by the constant probability λ of not updating the forecast at time t.

Let the inflation rate at time t π_t follows the AR(1) process:

$$\pi_t = \rho\pi_{t-1} + \epsilon_t \tag{2}$$

where ϵ_t - shock to the inflation rate at time t.

Then the impulse response function of inflation at time $t+k$ to a shock at time t:

$$\frac{d\pi_{t+k}}{d\epsilon_t} = \rho^k \tag{3}$$

Let define the optimal forecast of the $t+h$ inflation rate given the information set of an agent at time t as $\pi_{t+h|t}(i) = E(\pi_{t+h}|I_{it})$. Then the mean forecast error across all agents at time t is:

$$\overline{\pi_{t+h|t}} = \overline{E}(\pi_{t+h|t}(i)) = (1-\lambda)\sum_{k=0}^{\infty}\lambda^k E_{t-k}\pi_{t+h} \tag{4}$$

Moreover, combining with Eq. 2 gives:

$$E_{t-k}\pi_{t+h} = \sum_{k=0}^{\infty}\rho^{k+h}\epsilon_{t-s} \tag{5}$$

Therefore, the mean forecast error could be rewritten as:

$$\overline{\pi_{t+h|t}} = \sum_{k=0}^{\infty}(1-\lambda^{k+1})\rho^{k+h}\epsilon_{t-k} \tag{6}$$

Therefore, if $\lambda < 1$, when a shock, which rises inflation, occurs, the mean forecast under-reacts to that shock when compared to the response of the inflation rate.

It could be seen that over time the mean forecast converges to the true actual inflation rate. Rewriting the last equation for the forecast error $\pi_{t+h} - \overline{\pi_{t+h|t}}$, we get:

$$\begin{aligned}\pi_{t+h} - \overline{\pi_{t+h|t}} &= \sum_{m=1}^{h}\rho^m\epsilon_{t+m} + \sum_{k=0}^{\infty}\rho^{k+h}\epsilon_{t-k} - \sum_{k=0}^{\infty}(1-\lambda^{k+1})\rho^{k+h}\epsilon_{t-k} = \\ &= \sum_{m=1}^{h}\rho^m\epsilon_{t+m} + \sum_{k=0}^{\infty}\lambda^{k+1}\rho^{k+h}\epsilon_{t-k}\end{aligned} \tag{7}$$

Therefore, the impulse response function of the forecast error to a shock is:

$$\frac{dFE_{t+h+j,t+j}}{d\varepsilon_t} = \frac{d(\pi_{t+h+j} - \overline{\pi_{t+h+j|t+j}})}{d\epsilon_t} = \rho^{j+h}\lambda^{j+1} \tag{8}$$

Consequently, if λ is zero (that correspond to full-information rational expectations), agents update their expectations each period and the response of the forecast error, on average across agents, is zero.

When λ is rising, the average response of the forecast error is becoming persistently different from zero. Therefore, to evaluate the degree of information rigidity one could run the following regression:

$$\pi_{t,t-12} - \overline{\pi_{t,t-12|t-12}} = c + \sum_{k=1,..,K} \beta_k (\pi_{t-k,t-12-k} - \overline{\pi_{t,t-12-k|t-12-k}}) + $$
$$+ \sum_{j=0,..,J} \gamma_j \hat{\varepsilon}^j_{t-j} + \mu_t \quad (9)$$

where - $\pi_{t,t-12} - \overline{\pi_{t,t-12|t-12}}$ - the mean forecast error at time t, $\hat{\varepsilon}^j_t$ - shock j at time t, J and K are the number of lags selected using BIC. Shocks are evaluated outside of the regression, and the regression is run separately for each type of shock.

In terms of Eq. 9, the null of FIRE is that any shock ($\hat{\varepsilon}$) happened more than 12 months ago does not affect the mean forecast error because information about that shock is incorporated into agents' information sets. Therefore, the impulse response of the mean forecast error to any shock should be zero after 12 months. In other words, when the shock happens at time t ($\hat{\varepsilon}_t$), agents update their forecasts of the one-year ahead inflation rate ($\pi_{t+12,t}$), but since the horizon of forecasting is fixed at 12 months, agents cannot update the forecast of the inflation rates $\pi_{t,t-12}$, $\pi_{t+1,t-11}$, ..., $\pi_{t+11,t-1}$. Therefore, according to Eq. 9, when a shock happens at time t, under the null of FIRE the impulse response of $\pi_{t+12,t} - \overline{\pi_{t+12,t|t}}$ and further must be zero:

Forecast Error	$FE_{t,t-12}$	$FE_{t+1,t-11}$	$FE_{t+2,t-10}$	...	$FE_{t+12,t}, \ldots$
IRF to $\hat{\varepsilon}_t$	γ_0	$\gamma_1 + \beta_1\gamma_0$	$\gamma_2 + \beta_1(\gamma_1 + \beta_1\gamma_0) + \beta_2\gamma_0$	...	$\underbrace{\ldots, \ldots}$ under the null, after 1 year the impact of a shock is not significant

Thus, in terms of coefficients from Eq. 9, the non-linear combination of β and γ should be such that IRF of the mean forecast error to any shock happened more than one year ago is zero. One of the particular examples of coefficients combination which makes the response to a shock $\hat{\varepsilon}_t$ zero after 12 months is when all γ_j are zero. Another example is when all β_k are zero and γ_j are zero for $k \geq 12$.

Fig. 5 presents the results of the estimated Eq. 9 for the aggregate mean forecast error in response to a unit oil price shock at the horizon of 36 months [2]. First 12 months are dropped, since agents do not observe the shock. Fig. 5 also presents the response of the annual inflation to a unit oil price shock from the following regression:

$$\pi_{t,t-12} = c + \sum_{k=1,..,K} \beta_k \pi_{t-k,t-12-k} + \sum_{j=0,..,J} \gamma_j \hat{\varepsilon}^j_{t-j} + \mu_t \tag{10}$$

where - $\pi_{t,t-12}$ - actual annual inflation rate at time t, $\hat{\varepsilon}^j_t$ - shock j at time t, J and K are the number of lags selected using BIC.

Figure 5: Impulse Response to the Oil Price Shock

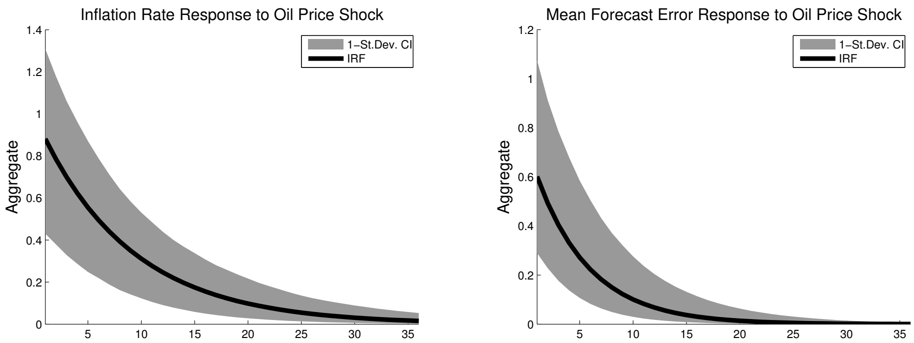

Note. The left plot shows the impulse response of the aggregate inflation rate to a unit oil price shock. IRF is estimated from Eq. 10. The right plot shows IRF of the mean forecast error estimated from Eq. 9. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

The response across all agents is similar to the one found by Coibion and Gorodnichenko (2012): the null hypothesis of no response of the mean forecast error is rejected at a one-standard-deviation level of significance.

Since the mean forecast error is calculated as the actual inflation rate minus the mean forecast, and the response of the inflation rate is positive (the left part of Fig. 5), the positive response of the mean forecast error means that households underestimate the rise of inflation in response to the oil price shock, exactly as the sticky information model predicts.

[2] Confidence intervals are calculated using parametric bootstrap.

According to the results, at time t+12 the mean forecast under-responds by 0.6 p.p. The mean forecast error response converges to zero over time. That also coincides with the prediction of the noisy and sticky information models.

Then I estimate Eq. 9 where instead of the oil price shock, I consider the monetary policy shock. Fig.6 reports the responses of the annual inflation rate and the mean forecast error across all agents to a unit contractionary monetary policy shock. As before, the responses to the shock at time t are plotted from the period t+12. The figure suggests that in response to the shock, the actual inflation rate drops by 0.8 p.p. The response of the forecast error is significant at a one-standard deviation level of significance, and the null of FIRE is rejected.

Figure 6: Impulse Responses to the Contractionary Monetary Policy Shock

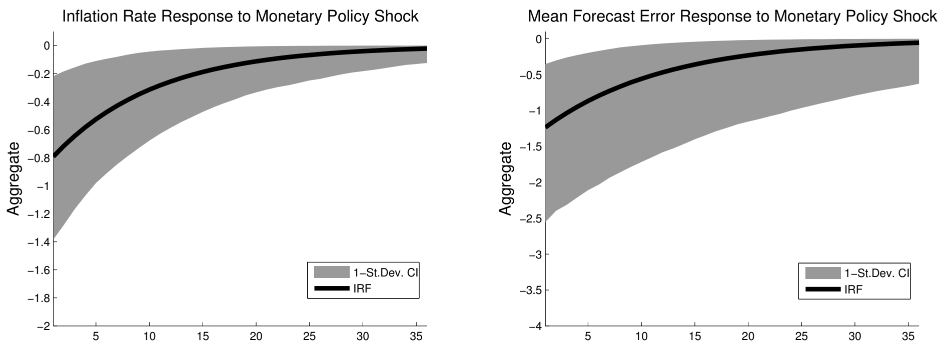

Note. The left plot shows the impulse response of the aggregate inflation rate to a unit monetary policy shock. IRF is estimated from Eq. 10. The right plot shows IRF of the mean forecast error estimated from Eq. 9. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

4 Measuring Income Specific Information Rigidity

Now I run the regression (Eq.9) replacing the mean forecast error across all agents by the mean forecast error of the specific income group. For each income group, I run the regression separately.

The impulse responses are presented in Fig.7. First, for all of the three income groups the null of FIRE can be rejected at a one-standard-deviation level. Second, there are differences in the absolute values of the responses across groups. According to the results, for low-income households, the median response is very small. On the other hand, responses of medium- and high-income households are significant: medium-income households make the error equal to 0.6 p.p., high-income households - around 0.7 p.p. Over time, the responses converge to zero.

Figure 7: Impulse Responses to the Oil Price Shock by Income Group

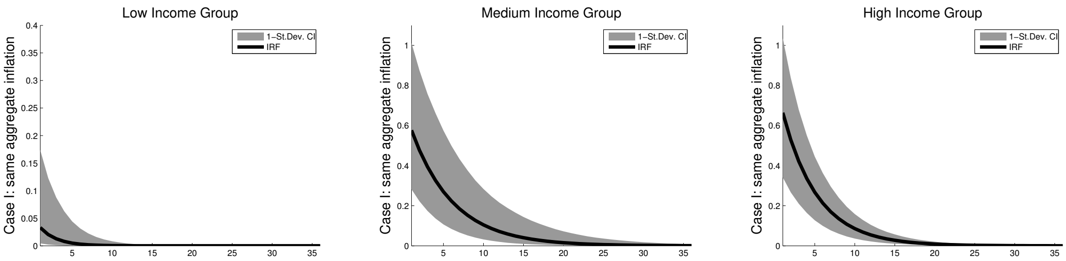

Note. Three plots report impulse responses of the mean forecast errors for different groups of households to the oil price shock. IRFs are estimated from Eq. 9. Households are grouped by the level of income: 1) low-income households: if reported pre-tax annual income is less than \$50k, 2) medium-income: \$50k-\$100k, 3) high-income: more than \$100k. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

The second step is to analyze the rates of convergence of the mean forecast errors. Eq.3 shows that the convergence rate of the inflation rate depends only on the persistence of the shock. Moreover, Eq. 8 could be rewritten as:

$$\frac{dFE_{t+h+j,t+j}}{d\varepsilon_t} = \frac{d(\pi_{t+h+j} - \overline{\pi_{t+h+j|t+j}})}{d\epsilon_t} = \rho^{j+h}\lambda^{j+1} = \frac{d\pi_{t+j+h}}{d\epsilon_t}\lambda^{j+1} \tag{11}$$

Therefore, the rate of convergence of the mean forecast error in response to a shock depends on two variables: the degree of information rigidity and the persistence of the shock. Hence, to estimate the degree of information rigidity one could normalize the impulse response of the mean forecast error by the impulse response of the inflation rate (taking the ratio of the estimated impulse response of the mean forecast error to the estimated impulse response of the inflation at the same horizon):

$$\frac{\frac{dFE_{t+h+j,t+j}}{d\varepsilon_t}}{\frac{d\pi_{t+h+j,t+j}}{d\varepsilon_t}} = \lambda^{j+1} \tag{12}$$

and then fit the AR(1) process to the normalized impulse response. The coefficient of the AR(1) process is the degree of information rigidity. In other words, according to the sticky information model, the normalized response of the mean forecast error to any shock should behave as the AR(1) process after h periods (for the one-year ahead inflation expectations $h = 12$).

Table 3 presents the estimates of λ for each income group. The numbers suggest that agents face significant information rigidities. The estimates for low-, medium-, and high-income households equal to 0.69, 0.93 and 0.89 respectively. That means that low-income households are subject to the lowest level of information rigidity, and, in terms of the noisy and sticky information models, that means that low-income households place the highest weight of 0.31 to new information relative to the old forecasts (noisy information model), or that low-income households update their belief with probability 0.31 each month (sticky information model). In contrast, for high-income households, the weight of new information is equal to 0.07.

Table 3: Convergence rates of the mean forecast errors

Income group	$\hat{\lambda}$
Low	0.69
	(.10)
Medium	.93
	(.06)
High	.89
	(.04)
p-value for equality	.000

Note. The table presents the estimates of the income-specific degree of information rigidity. For each group, the number is the AR(1) coefficient of the normalized mean forecast error response (Eq.12). The description of the procedures to calculate standard errors and p-value for the test of equality of estimates are presented in Appendix.

Finally, Table 3 presents the p-value for the test of equality of the estimates across three income groups [3]. The p-value suggests that the null of equality between three estimates is rejected at 1% level of significance.

Fig.8 plots the impulse responses of the mean forecast errors to the monetary policy shock by income group. Several observations could be made. First, in contrast to the "oil price shock" case, there is higher heterogeneity in the responses across groups. Thus, for low- and medium-households the null of the full-information rational expectations is rejected, but not for the high-income group. The non-significant response of the mean forecast error corresponds to zero degree of information rigidity ($\lambda = 0$). Second, middle-income households make huge errors when adjusting their forecasts to the monetary policy shock. In comparison to the low-income group whose median forecast error is around 0.3 p.p., the response of the medium-income group is around 1.3 p.p.

Figure 8: Impulse Responses to the Monetary Policy Shock by Income Group

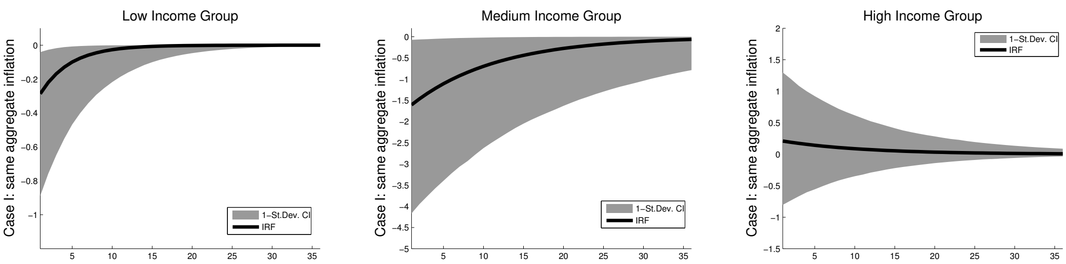

Note. Three plots report the impulse responses of the mean forecast errors for different groups of households to a unit positive (contractionary) monetary policy shock. IRFs are estimated from Eq.9. Households are grouped by the level of income: 1) low-income households: if reported pre-tax annual income is less than \$50k, 2) medium-income: \$50k-\$100k, 3) high-income: more than \$100k. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

Finally, Table 4 presents the estimates of the "income-specific" degrees of information rigidity estimated as the AR(1) coefficient of the normalized mean forecast error response. For high-income households, λ is set to be 0 since the null of FIRE is not rejected. For low-income households, $\hat{\lambda}$ is 0.85. For middle-income households, the estimate of λ is

[3] The description of the procedures to calculate λ, standard errors and the description of a test are presented in Appendix.

equal to 1. In terms of the sticky information model, it means that agents are completely inattentive to new information about monetary policy shocks, and their forecasts are based only on the old information. From Eq. 8 it means that the response of the mean forecast error converges over time to 0 only because the inflation rate converges to 0.

Table 4: Convergence rates of the mean forecast errors

Income group	$\hat{\lambda}$
Low	0.85
	(.13)
Medium	1.00
	(.11)
High	0
p-value for equality*	0.020

**the p-value is for the test of equality of estimates among low- and middle-income households.*
Note. The table presents the estimates of the income-specific degrees of information rigidity. For each group, it is estimated as the AR(1) coefficient of the normalized mean forecast error response (Eq.12). The description of the procedures to calculate standard errors and the p-value for the test of equality of estimates are presented in Appendix.

However, one concern about surveys of households' expectations (Coibion and Gorodnichenko, 2012) is that usually survey questions do not specify to households a particular price index to forecast. Therefore, the observed differences in Fig.7 and Fig.8 in the responses across groups may partially arise from that households at different income percentiles have different consumption baskets, and instead of forecasting the aggregate inflation rate, agents predict changes in prices of their consumption baskets.

Some evidence that agents forecast better the prices of their consumption baskets is presented in Fig.9. Survey of Consumer Expectations conducted by the New York Fed also asks households to forecast the price changes of some particular goods. For example, changes in the prices of gold and gas. I take the actual annual price changes of these two categories from FRED and calculate the forecast errors for each income group. Thus, fig. 9 presents the forecast errors of gold and gas prices by income group. It could be noticed that high-income households make persistently smaller errors of gold prices, while low-income households - the largest forecast errors. It is vice-versa for gas prices: most of the time, agents belonging to the high-income group make larger errors than the other two groups.

Figure 9: *The mean forecast errors of gas and gold prices by income group. Data sources: inflation expectations - New York Fed Survey of Consumer Expectations, gold and gas prices - from FRED.*

Moreover, products vary in price stickiness (Nakamura and Steinsson, 2008). Therefore, economic shocks may affect the inflation rates faced by different households heterogeneously. For example, Cravino et al. (2018) find that consumption baskets of high-income households are less volatile, and, because of that, monetary policy shocks affect CPI of high-income households three times less than CPI of middle-income households. Therefore, if when households are asked to forecast the inflation rate, they forecast the changes in prices of their consumption baskets instead of the aggregate inflation, the differences between groups observed in Fig. 7, 8 and Tables 3, 4 could be overstated because of the heterogeneity in households' consumption bundles.

In the next section, I construct income-specific inflation rates and repeat the analysis controlling for the differences in consumption baskets.

4.1 The Role of Household-Specific Consumption Basket

In this section, I estimate the degree of information rigidity under the assumption that agents forecast the changes in the prices of their consumption baskets instead of the aggregate inflation rate. Now, in Eq.9 the aggregate inflation rate $\pi_{t,t-12}$ is replaced with the income-specific inflation rate $\pi^j_{t,t-12}$ $(j = 1, 2, 3)$.

Fig.10 presents the mean forecast error responses to the oil price shock. The upper panel duplicates Fig.7 for the convenience. The results suggest that the null of the FIRE can still be rejected at a one-standard-deviation level of significance for all groups of households. However, there are clear changes in the sizes of the responses and the speed of convergence of the mean forecast errors. The median response starts from approximately 0.05 p.p. instead of 0.7 p.p. for the high-income group and 0.2 p.p. instead of 0.6 p.p. for the middle-income. The response of low-income households stays roughly the same. Therefore, in contrast to the upper panel, middle-income households are those who, adjusting forecasts to the oil price shock, make the largest errors.

Table 5 compares the degree of information rigidity estimated from the AR(1) process of the mean forecast error response normalized by the response of the income-specific inflation rate. Column 1 duplicates the results from Table 3.

Under the assumption that agents forecast changes in the prices of their consumption baskets, the estimates of λ for low-, middle-, and high-income groups are 0.76, 0.81 and 0.70 respectively versus 0.69, 0.93 and 0.89 before. Thus, controlling for the differences in consumption baskets, the results indicate that high-income households face the lowest degree of information rigidity among three groups.

In addition, the p-value for the test of equality of the estimates across groups suggests that the null hypothesis cannot be rejected. That is in contrast to what is observed in the first case. Finally, Table 5 also presents the p-values for the test of equality of the estimates across two cases (homogeneous versus heterogeneous inflation rates). For all the three groups, the null of identical estimates is rejected.

The evidence could be summarized as follows. First, controlling for the heterogeneity of

Figure 10: Impulse Responses to the Oil Price Shock

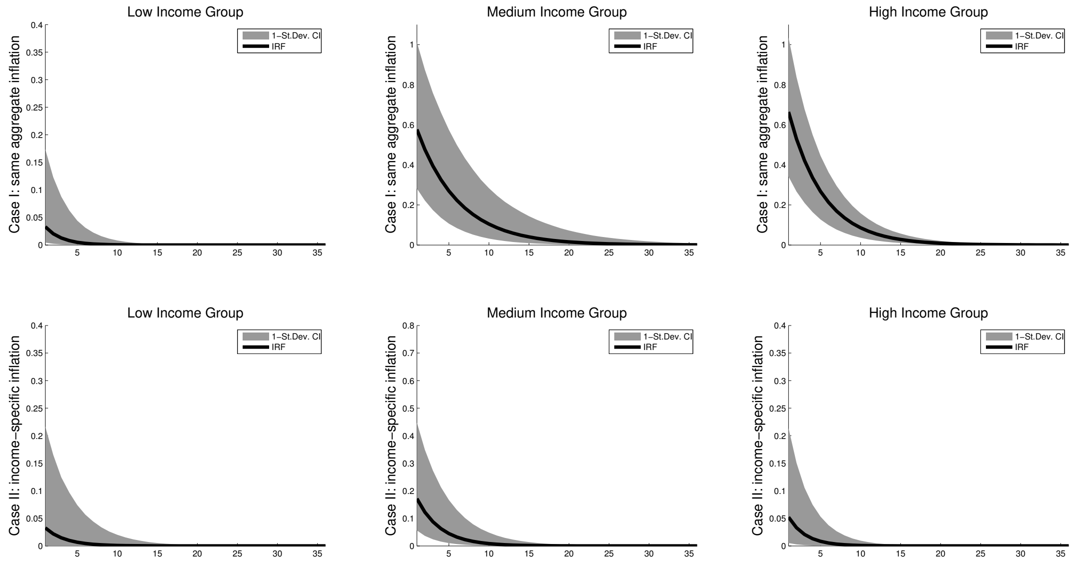

Note. Two panels report the impulse responses of the mean forecast errors for three income groups to the oil price shock. IRFs are estimated from Eq.9. The bottom panel: for each group, the mean forecast error is calculated as the income-specific inflation rate (instead of the aggregate inflation) minus the mean forecast error of that group. The upper panels duplicates Fig.7 for the convenience. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

Table 5: Convergence rates of the mean forecast errors (λ)

Income group	(1)	(2)	p-value for equality
Low	0.69	0.76	0.0006
	(.10)	(.12)	
Medium	.93	.81	0.0000
	(.06)	(.08)	
High	.89	.70	0.0003
	(.04)	(.10)	
p-value for equality	.000	0.379	

Note. Degree of information rigidity is estimated as the coefficient of the AR(1) process of the normalized mean forecast error response. (1) - normalized by the response of the aggregate inflation rate, (2) - by the income-specific inflation rates.

households consumption baskets is important: the differences between groups significantly lessen once controlling. Second, in response to the oil price shock, high-income households adjust their forecasts better than the other groups. Middle-income households are those who face the highest degree of information rigidity among groups.

Figure 11: Impulse Responses to the Monetary Policy Shock

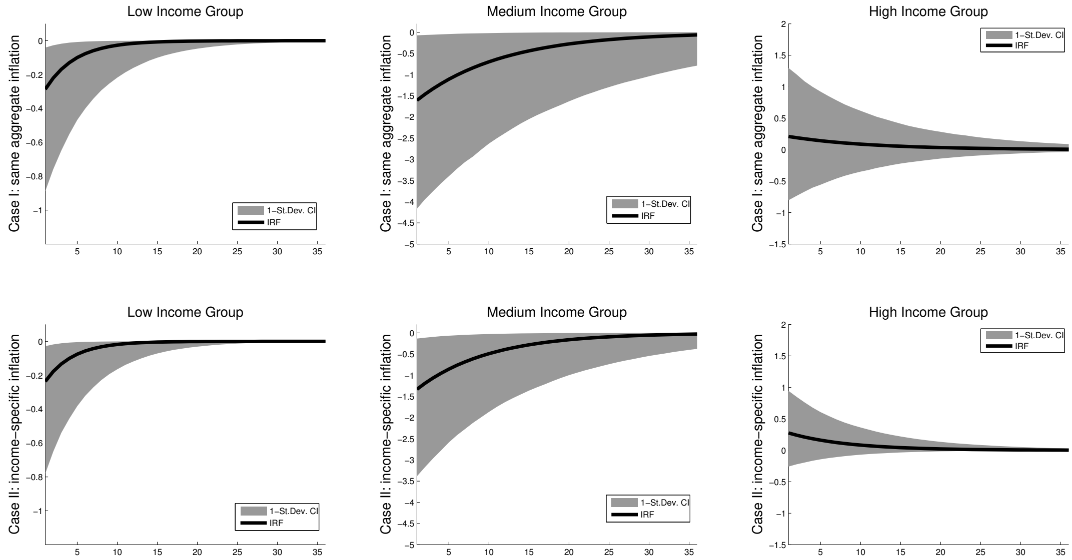

Note. Two panels report the impulse responses of the mean forecast errors for three income groups of households to a unit positive (contractionary) monetary policy shock. IRFs are estimated from Eq.9. The bottom panel: for each group, the mean forecast error is calculated as the income-specific inflation rate (instead of the aggregate inflation) minus the mean forecast error of that group of households. The upper panels duplicates Fig. 8 for the convenience. Confidence intervals are calculated with parametric bootstrap. The black line shows the median response.

Fig.11 reports the mean forecast error responses to the monetary policy shock. According to the results, controlling for the heterogeneity of consumption baskets seems less important than it is in the case of the oil price shock. The conclusion from the plots is the same as before: medium-income households still make the largest errors, while for the high-income group the null of FIRE is not rejected. The median sizes of the responses have changed just slightly: from 0.3 to 0.2 for the low-income group and from 1.6 to 1.3 for the middle-income. One potential explanation of the result that high-income households do not make significant errors in response to monetary policy shocks is that wealthy households spend a higher share of their income to make financial investments. The returns on financial instruments depend on market beliefs about monetary policy (for example, Gurkaynak et al. (2005)). Hence, high-income households may be motivated to track monetary policy announcements and how the monetary policy shocks affect their wealth. In contrast, the result that middle-income households are not attentive to monetary policy shocks seems to be puzzling.

Table 6 presents the estimates of the degree of information rigidity. The main conclusion from the table is that the difference in the estimates across cases (1) and (2) is not statistically significant for all the three income groups. The middle-income households are still completely inattentive to the monetary policy shocks. The difference in estimates across low- and middle-income groups is not statistically different from zero as well.

Table 6: Convergence rates of the mean forecast errors (λ)

Income group	(1)	(2)	p-value for equality across cases	p-value for equality across shocks[b]
Low	0.85	0.83	0.804	0.700
	(.13)	(.14)		
Medium	1.00	1.00	0.945	0.200
	(.11)	(.15)		
High	0	0	-	
p-value for equality[a]	0.020	0.229		

[a] p-value for the test of identical estimates across low- and middle-income households.
[b] p-value for the test of identical "income-specific" λ across two types of shocks.

Note. Convergence rates are estimated as the coefficient of AR(1) process of the normalized mean forecast error response. (1) - normalized by the response of the aggregate inflation rate, (2) - by the income-specific inflation rate.

The results from the two cases suggest that for both types of shocks - monetary policy shocks and oil price shocks - middle-income households make the largest errors when adjusting their forecasts, while the high-income group of households performs better than the other two groups. The last seems intuitive because households earning high income may have a higher level of education, and, therefore, the level of financial and economic literacy.

Finally, according to the sticky information model, the convergence rate of the mean forecast error (λ) should be the same for all types of shocks (Eq.12). In contrast, Coibion and Gorodnichenko (2012) show that the basic noisy information model (Woodford, 2003) allows for the different rates of convergence. The estimates for high-income households suggest that the degree of information rigidity varies across two types of shocks. That may suggest that the noisy information model provides a better characterization of how high-income agents form their expectations than the sticky information model does. For low- and middle-income households the p-values for the test of equality of estimates across shocks are 0.7 and 0.2 (the last column of Table 6). Therefore, for these two groups, we can not distinguish between the sticky and noisy information models.

5 Conclusions

The majority of empirical papers which derive the measures of households' information rigidity focus on the aggregate expectations. In this paper, using the Survey of Consumer Inflation Expectations from the New York Fed, I divide households into three income groups and estimate households' "income-specific" degrees of information rigidity.

One concern about the survey is that it does not specify which price index agents should forecast. Since households at different income percentiles have different consumption baskets, and economic shocks affect prices of different goods heterogeneously, households may forecast the changes in prices they face instead of the aggregate inflation. Therefore, I also consider responses of the mean forecasts errors controlling for the differences in consumption baskets across income groups of households.

Conditional responses of the mean forecast errors to the oil price shock suggest that all income groups are subject to some extent of information rigidity, but the estimates significantly vary across groups: surprisingly, low-income households adjust their forecasts better than the other two groups. However, controlling for the heterogeneity of consumption bundles significantly reduces the differences in the degree of information rigidity across households. Moreover, in the second case, the high-income group is found to be the one with the least degree of information rigidity.

In contrast, controlling for the heterogeneity seems to be not important in the case of monetary policy shocks: I find the estimates of the degree of information rigidity to be roughly the same in both cases. However, there is high heterogeneity in the mean forecast error responses and the degree of information rigidity across groups: for the high-income group, I cannot reject the null of FIRE which corresponds to zero degree of information rigidity. One potential explanation is that wealthy households make more investments in financial instruments, and, therefore, they are more attentive to monetary policy announcements. For low- and middle-income households, the null of FIRE is rejected. Moreover, I find middle-income households to be fully inattentive to monetary policy shocks. That puzzling result requires further investigation.

One potential improvement of the current paper is to consider other households' characteristics which may explain the differences in the mean forecast error responses, and, particularly, why middle-income households make the largest errors. For example, both surveys (Survey of Consumer Expectations and Survey of Consumer Expenditures) ask households about age, level of education or employment status. Another improvement is to use more disaggregated data on consumer expenditures (at the micro level) and construct the inflation rates for more detailed income groups (for instance, for each income percentile). That will allow to assign the particular inflation rate to each household from the Consumer Expenditure Survey and calculate the forecast error not for a group of households, but for each household. That may eliminate the heterogeneity in consumption baskets which could be still present when I use the inflation rates for three income groups.

The result that information rigidity of households at different income percentiles is not the same may be important for macroeconomic models which incorporate information rigidities. For example, one could study how the implications of monetary policy change if wealthy and poor households have heterogeneous degrees of information rigidity.

References

Andrade, P. and Bihan, H. L. (2010). Inattentive Professional Forecasters. Working Paper 307, Banque de France, Paris.

Armantier, O., Nelson, S., Topa, G., van der Klaauw, W., and Zafar, B. (2012). The Price is Right: Updating Inflation Expectations in a Randomized Price Information Experiment. Manuscript.

Bernanke, B. (2013). Communication and Monetary Policy. Speech delivered at the National Economists Club Annual Dinner Herbert Stein Memorial Lecture Washington, D.C., November 19, 2013.

Binder, C. (2017). Consumer forecast revisions: Is information really so sticky? *Economics Letters*, 161 (12):112–115.

Bordo, Michael, C. E. A. L. and Michaels, R. (2007). Three Great American Disinflations. Fed. Reserve Bank San Francisco Proceedings.

Bryan, M. F. and Venkatu, G. (2001). The Curiously Different Inflation Perspectives of Men and Women. Federal Reserve Bank of Cleveland Economic Commentary Series.

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics*, 118 (2):269–98.

Christiano, L., Eichenbaum, M., and Evans, C. (1999). Monetary Policy Shocks: What Have We Learned, and to What End? *In Handbook of Monetary Economics, edited by John B. Taylor and Michael Woodford, Amsterdam: Elsevier Sci.*, 1 (Part A):65–148.

Coibion, O. (2010). Testing the Sticky Information Phillips Curve. *Review of Economics and Statistics*, 92 (2):87–101.

Coibion, O. and Gorodnichenko, Y. (2012). What Can Survey Forecasts Tell Us about Information Rigidities? *Journal of Political Economy*, 120 (1):116 – 159.

Coibion, O., Gorodnichenko, Y., and Kumar, S. (2018). How Do Firms Form Their Expectations? New Survey Evidence. *American Economic Review*, 108 (9):2671–2713.

Cravino, J., Lan, T., and Levchenko, A. (2018). Price stickiness along the income distribution and the effects of monetary policy. *Journal of Monetary Economics, forthcoming.*

Gertler, M. and Karadi, P. (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics*, 7 (1):44–76.

Gurkaynak, R., Sack, B., and Swanson, E. (2005). Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking*, 1 (1):55–93.

Hamilton, J. D. (1996). This is what happened to the oil price-macroeconomy relationship. *Journal of Monetary Economics*, 38 (2):215–220.

Hamilton, J. D. (2009). Daily Changes in Fed Funds Futures Prices. *Journal of Money, Credit and Banking*, 41 (4):567–582.

Johannsen, B. K. (2014). Inflation Experience and Inflation Expectations: Dispersion and Disagreement Within Demographic Groups. Finance and Economics Discussion Series 2014–89, Board of Governors of the Federal Reserve System (US).

Kuttner, Kenneth, N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47 (3):523–544.

Mankiw, N. G. and Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117 (November):1295–1328.

Morris, S. and Shin, H. S. (2002). Social Value of Public Information. *American Economic Review*, 92 (December):1521–34.

Nakamura, E. and Steinsson, J. (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models. *The Quarterly Journal of Economics*, 123 (4):1415–1464.

Patton, A. J. and Timmermann, A. (2006). Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion. *Journal of Monetary Economics*, 57 (October):803–20.

Reis, R. (2006). Inattentive Producers. *Review of Economic Studies*, 73 (July):793–821.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics*, 50 (3):665–90.

Stock, J. H. and Watson, M. W. (2016). Chapter 8 - Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 415–525.

Woodford, M. (2003). Imperfect Common Knowledge and the Effects of Monetary Policy. *In Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps, edited by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford. Princeton, NJ: Princeton Univ. Press.*

6 Appendix

6.1 Data Sourses

- **Inflation expectations**: Survey of Consumer Expectations, New York Fed (monthly survey, June 2013 - March 2019);
- **Oil prices**: West Texas Intermediate, FRED;
- **Futures data**: FF4 - three month ahead fed funds future rate, daily data (open and close prices), Bloomberg;
- **Disaggregated CPI time series**: U.S. Bureau of Labor Statistics (BLS);
- **Expenditure shares by income groups**: U.S. Consumer Expenditure Survey (CEX), BLS.

6.2 Estimation of Impulse Response Functions

For each income group, Impulse Responses (Fig.5,7,6,8) are calculated using parametric bootstrap as follows.

1. Estimate Eq.10 using Newey-West HAC standard errors. The number of lags is chosen by BIC. Let define the vector of coefficients as $\hat{\alpha} = [\hat{\beta}\ \hat{\gamma}]'$ and variance-covariance matrix as $\hat{\Omega}$.
2. Bootstrap sampling. Generate a vector $randn_M$ consisting of M numbers from the standard normal distribution, where M - the number of coefficients in $\hat{\alpha}$. Generate bootstrap coefficients as $\hat{\alpha}_{boot} = \hat{\alpha} + \hat{\Omega}^{1/2} * randn_M$.
3. Using $\hat{\alpha}_{boot}$ and Eq.10, calculate IRF of the mean forecast error to a shock at the selected horizon (for example, 5 years). Drop the first 12 periods of IRF since agents do not observe the shock in the first 12 months.
4. Repeat steps 2) and 3) 10000 times.
5. To plot the median response to a shock (black line in Fig.5,7,6,8), choose the 50th percentile of 10000 bootstrap IRF. To get the one-standard deviation confidence intervals (grey shaded area), plot the 16th and 84th percentiles of 10000 bootstrap IRF.

6.3 Estimation of Degree of Information Rigidity

For each income group, the degree of information rigidity (λ) is calculated as follows.

1. Estimate separately Eq.9 and Eq.10.
2. Using the coefficients from the estimated regressions, calculate the IRF of the mean forecast error and the IRF of the inflation rate to a unit shock.
3. Divide the IRF of the mean forecast error by the IRF of the inflation rate. According to the sticky information model, the ratio between two IRFs should converge to 0 as the AR(1) process.
4. Fit the AR(1) coefficient to the normalized mean forecast error response. The coefficient is the estimate of degree of information rigidity ($\hat{\lambda}$).

For each income group, standard errors for the degree of information rigidity from Tables 3 and 4 are calculated with parametric bootstrap as follows.

1. Using Eq.10 and the parametric bootstrap procedure from Section 8.1, obtain 10000 bootstrap IRFs of the inflation rate to a unit shock. In total, there are 10000 bootstrap IRFs of the mean forecast error and 10000 bootstrap IRFs of the inflation rate.
2. For each bootstrap sampling j, divide the j-th IRF of the mean forecast error by the j-th IRF of the inflation rate, where j=1,...,10000. Let define the ratio as IRF_j^{norm}. According to the sticky information model, IRF_j^{norm} should converge to 0 as the AR(1) process.
3. For each j, fit the AR(1) coefficient. Thus, there will be a vector $\hat{\lambda}_{boot}$ of 10000 bootstrap estimates of λ.
4. Take the standard deviation of $\hat{\lambda}_{boot}$. That is the standard error of the degree of information rigidity.

Steps to calculate the p-value for the test of equality of $\hat{\lambda}$ across N income groups or across N cases (tables 3, 4, 5, 6):

1. Calculate the NxN variance-covariance matrix $\hat{\Lambda}$ of $(\hat{\lambda}^1_{boot}...\hat{\lambda}^N_{boot})$ obtained with bootstrap. For example, for the test of identical estimates across three income groups, $\hat{\Lambda}$ is 3x3 matrix.
2. Calculate the test statistics $\chi = (\hat{\lambda} - \overline{\lambda})\hat{\Lambda}^{-1}(\hat{\lambda} - \overline{\lambda})'$, where $\hat{\lambda}$ - $Nx1$ vector of the estimates of the degree of information rigidity, $\overline{\lambda}$ - the mean of these estimates. The obtained statistics χ is distributed as χ^2_{N-1}, where N - the number of groups.

6.4 Construction of income-specific CPIs from the Consumer Expenditure Survey

Description of expenditure shares adjustments:

- According to BLS, expenditures on the following items: "homeowner insurance", "maintenance", and "major appliances" should be adjusted by multiplying by a factor of 0.43. That is needed because these items include both consumption and investment components, and, when calculating CPI, only consumption expenditures should be counted.
- Similarly to Cravino et al. (2018), I adjust expenditures on the item "used cars" by multiplying by the factor of 0.5. That is needed because CPI should reflect only the dealer value added. Therefore, we have to subtract trade-in value of vehicles and other sales of consumer-owned vehicles. Cravino et al. (2018) found that the portion is around one half.
- Such survey items as cash contributions, pensions and social security were excluded when calculating consumption bundles.

6.5 Other Figures

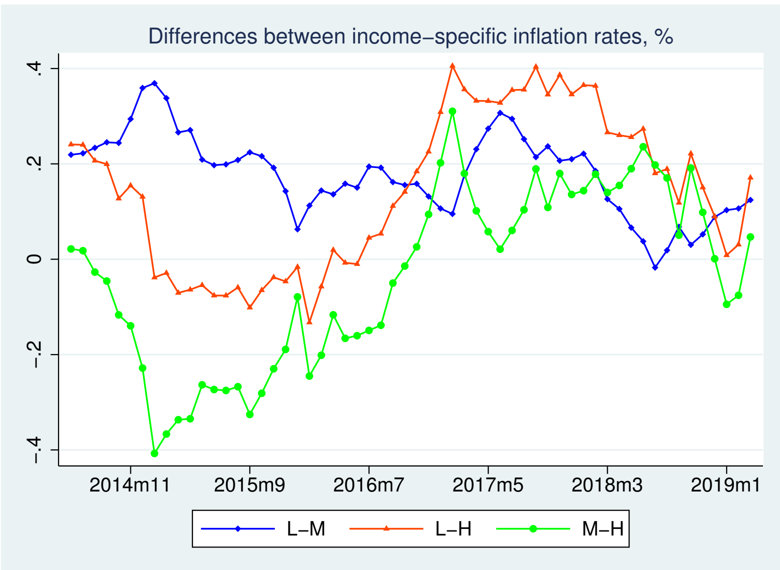

Figure 12: *The figure reports the differences between the income-specific inflation rates. L-M: the difference between the "low-income" and the "middle-income" inflation rate, L-H: the "low-income" minus the "high-income" inflation rate, M-H: "middle" minus "high".*