Department of Economics Discussion Paper Series

Heterogeneity in imperfect inflation expectations: theory and evidence from a novel survey

Alistair Macaulay and James Moberly

Number 970
April, 2022

Department of Economics
Manor Road Building, Manor Road
Oxford, OX1 3UQ

ISSN 1471-0498

Heterogeneity in imperfect inflation expectations: theory and evidence from a novel survey*

Alistair Macaulay† James Moberly‡

April 20, 2022

Abstract

A given observation of uncertainty in expected inflation could be consistent with many different beliefs about how inflation is formed, with different implications for the aggregate transmission of shocks. We add novel questions to the Bundesbank Survey of Consumer Expectations to elicit (i) how persistent households perceive inflation to be, and (ii) how certain they are in their inflation perceptions. Combining these with existing survey questions, we infer laws of motion for expectations at the individual level. Based on averages alone, a standard model calibrated to our data predicts shocks to inflation generate small and transitory responses in expectations and consumption. Accounting for the large heterogeneity in expectation laws of motion across households, however, increases the transmission of inflation shocks to aggregate consumption through expectations by an order of magnitude, and substantially increases the persistence of the consumption response.

*This paper uses data from the Bundesbank-Online-Panel-Households. The results published and the related observations and analysis may not correspond to results or analysis of the data producers. We are grateful Alex Haas and Susanne Helmschrott for help with designing the survey questions. We thank Guido Ascari, Sergio di Ferra, Lena Dräger, Martin Ellison, Andrea Ferrero, Michael McMahon, Sebastian Rast, Francesco Zanetti, and seminar participants at the University of Oxford for valuable feedback

†Department of Economics & St. Anne's College, University of Oxford. Email: alistair.macaulay@economics.ox.ac.uk

‡Department of Economics & University College, University of Oxford. Email: james.moberly@economics.ox.ac.uk

1 Introduction

Households' information and subjective models of inflation shape how their inflation expectations, and hence macroeconomic aggregates, respond to shocks and policy changes. However, commonly used survey data on expectations is consistent with various combinations of information and subjective models, with contrasting implications. The well-known regressions in Coibion and Gorodnichenko (2015a), for example, are consistent with models in which information is noisy but households know the true law of motion for inflation (Coibion and Gorodnichenko, 2015a), or models with full information but misspecified forecasting rules (Gabaix, 2020; Hajdini, 2020). How then do households form their expectations?

This paper uses novel survey data to answer that question. We add new questions to the Bundesbank's Survey on Consumer Expectations to elicit the uncertainty in household perceptions of current inflation, and how persistent households perceive inflation to be. When combined with responses to other standard survey questions, this allows us to infer subjective laws of motion and details of information processing at the individual level.

We find that, on average, uncertainty about current inflation is low, and the perceived persistence of inflation is close to that of realised inflation. However, these averages mask a great deal of heterogeneity, which is important for understanding aggregate behaviour. Calibrating a standard consumption-saving model to our data, heterogeneity in the expectations process amplifies the aggregate consumption response to inflation shocks by an order of magnitude, relative to the representative-agent case. Expectations may therefore play a substantially larger role in business cycle fluctuations than implied by models based on aggregate expectations alone.[1]

How do these questions pin down households' subjective models of inflation? Suppose that each household, indexed i, believes that inflation follows the AR(1) process:[2]

$$\pi_t = \hat{\rho}_i \pi_{t-1} + \varepsilon_t \tag{1}$$
$$\varepsilon_t \sim N(0, \tilde{\sigma}^2_{\varepsilon,i})$$

Uncertainty in the inflation forecast can then be decomposed into two components: (i) uncertainty about the current inflation rate, and (ii) uncertainty arising from future

[1] e.g. Fuster et al. (2010), Bhandari et al. (2019), Angeletos et al. (2020), and many others.

[2] While this is likely a simplification of the true data generating process for inflation, numerous studies suggest that household forecasts are better characterized by simple forecasting rules than by rationality (e.g. Adam, 2007).

shocks:[3]

$$\tilde{Var}_{i,t}(\pi_{t+1}) = \hat{\rho}_i^2 \tilde{Var}_{i,t}(\pi_t) + \tilde{\sigma}^2_{\varepsilon,i} \tag{2}$$

Existing survey questions measure subjective uncertainty in the forecast $\tilde{Var}_{i,t}(\pi_{t+1})$. However, different combinations of the two components of this uncertainty imply very different aggregate dynamics, for the same observed forecast uncertainty. To illustrate this, suppose we know $\tilde{Var}_{i,t}(\pi_{t+1})$, and consider three different ways in which this forecast variance could be decomposed.

1. $\hat{\rho}_i = 1$, $\tilde{\sigma}^2_{\varepsilon,i} = 0$, $\tilde{Var}_{i,t}(\pi_t) = \tilde{Var}_{i,t}(\pi_{t+1})$. Agents are uncertain in their perceptions, suggesting they are imperfectly informed about current inflation. This implies their inflation perceptions respond sluggishly to shocks. However, agents also overextrapolate from their perceptions to expectations, leading to 'delayed overshooting' (as in Angeletos et al., 2020).
2. $\hat{\rho}_i = 1$, $\tilde{\sigma}^2_{\varepsilon,i} = \tilde{Var}_{i,t}(\pi_{t+1})$, $\tilde{Var}_{i,t}(\pi_t) = 0$. Agents still overextrapolate, but they are certain about their perception of inflation, so there is no delay in the response of perceptions and expectations to realised inflation.
3. $\hat{\rho}_i = 0$, $\tilde{\sigma}^2_{\varepsilon,i} = \tilde{Var}_{i,t}(\pi_{t+1})$, $\tilde{Var}_{i,t}(\pi_t)$ undefined. Agents underextrapolate (as in Gabaix, 2020). Forecasts only respond to signals about future inflation.

A given level of uncertainty in inflation forecasts could therefore be consistent with a continuum of models, which have very different implications for how expectations, and hence macroeconomic aggregates, respond to shocks. Our questions pin down the correct model at the household level by eliciting the variance of the inflation perception $\tilde{Var}_{i,t}(\pi_t)$, and the perceived persistence of the forecast $\hat{\rho}_i$. Combined with existing questions eliciting $\tilde{Var}_{i,t}(\pi_{t+1})$, we can then calculate the perceived variance of the innovations $\tilde{\sigma}^2_{\varepsilon,i}$.

To measure uncertainty in perceptions, we start with an existing survey question eliciting a point estimate of the respondent's inflation perception. We then add a new question, asking for the probability that the inflation rate was within a specified range around the respondent's point estimate. Fitting a triangular distribution to these responses as in Coibion et al. (2021) yields $\tilde{Var}_{i,t}(\pi_t)$. To measure $\hat{\rho}_i$, we present respondents with hypothetical scenarios of macroeconomic shocks, as in Andre et al. (2022). We specify that the shock has increased current inflation by one percentage point, and ask how the respondent would update their expected future inflation as a result.

Our finding that, on average, respondents are less uncertain about current inflation than future inflation suggests that much of the uncertainty in expectations comes from

[3]We abstract here from uncertainty about $\hat{\rho}_i$.

perceived noise in the inflation process, not a lack of information about current inflation. However, a minority of households are very uncertain in their inflation perceptions. Perceived persistence is similarly heterogeneous: while two-thirds of households perceive no persistence in the inflation process at the one-year horizon, those who do update their expectations after the hypothetical shock often do so by a large amount.

This heterogeneity has a large impact on how aggregate consumption responds to shocks. In a standard partial equilibrium consumption-saving model, the individual consumption response to an inflation shock is highly convex in perceived inflation persistence. The same path of aggregate inflation expectations is therefore associated with much larger fluctuations in aggregate consumption if the individual-level subjective laws of motion are heterogeneous. Calibrating a representative agent model to average parameters from our data, a 1% shock to inflation implies a 0.08% increase in aggregate consumption on impact, slightly below the response under full information and rational expectations. Taking account of the heterogeneity we observe, however, implies aggregate consumption rises by 1.28% on impact. The persistence of aggregate consumption also rises by nearly one-half.

Moreover, elements of household expectation formation correlate systematically with each other, and with household characteristics, which further distorts aggregate consumption behaviour away from the representative agent case. In particular, households with little liquid wealth, who are likely to be hand-to-mouth, are more uncertain about current inflation relative to future inflation, and believe inflation to be substantially more persistent. The overextrapolation of aggregate expectations found in other studies (e.g. Angeletos et al., 2020) may therefore be partly driven by hand-to-mouth households who cannot react to their expectations. Accounting for hand-to-mouth households reduces the aggregate partial equilibrium consumption response to inflation in our model, but the response remains an order of magnitude larger than the case with homogeneous expectation formation.

Related Literature: Angeletos et al. (2020) similarly study the reaction of expectations to shocks, directly estimating impulse responses of aggregate expectations using local projections. We view our approach as complementary. While their aggregate impulse responses can capture richer subjective models than the linear approximations we identify, our approach allows us to study heterogeneity in households' subjective laws of motion and information processing, which cannot be observed with local projections using average inflation forecasts.[4]

There is a large literature testing noisy information models. Badarinza and Buchmann

[4]Ryngaert (2018) estimates perceived persistence and signal precision among professional forecasters, but similarly cannot uncover heterogeneity in these parameters.

(2009) and Mankiw et al. (2004) find considerable disagreement in inflation perceptions and expectations respectively, indicating that households receive noisy signals about current inflation. Other studies measure how agents update beliefs in response to information treatments (e.g Armantier et al., 2015; Coibion et al., 2022), or use laboratory experiments to test models of rational inattention (Caplin and Dean, 2015; Dean and Neligh, 2019). As in our data, the extent of noise is frequently found to be very heterogeneous across agents (Link et al., 2021). We complement this literature by directly measuring the variance in perceptions, which is an important component of expectation formation in noisy information models when the tracked variable is persistent (Coibion and Gorodnichenko, 2015a; Bordalo et al., 2020). To the best of our knowledge, there are no existing quantitative measures of this variable.[5]

Alongside the results of Angeletos et al. (2020), overshooting has also been documented in laboratory experiments (Beshears et al., 2013; Afrouzi et al., 2020). Within surveys, Armona et al. (2019) use an information treatment to assess how households revise house-price expectations when perceptions of past developments change, and Laudenbach et al. (2021) and Beutel and Weber (2021) use similar methods to study the perceived persistence of stock returns. Relative to this, our method allows us to capture the perceived inflation persistence even of households with very certain perceptions, who may not react to information treatments.[6] Most importantly, we jointly measure perceived persistence and the relative uncertainty over perceived and expected inflation. The correlations between these components matter for the transmission of shocks.

Our findings also reveal how aspects of expectation formation relate to personal characteristics, particularly income and wealth. Galati et al. (2021) and Vellekoop and Wiederholt (2019) document that expectations vary systematically with income and wealth respectively. Our contribution is to go a layer deeper, and show how the perceived persistence and uncertainty in perceptions involved in expectation formation vary along those dimensions.

Finally, we contribute to the literature on the role of expectations in business cycles. Expectations are crucial in any forward-looking environment, so frictions in expectation formation can have important effects on business cycles (e.g. Gorodnichenko and Sergeyev, 2021). Like us, Branch and Evans (2006) and Hommes and Lustenhouwer (2019) find that heterogeneity in the expectation formation process can substantially alter macroeconomic outcomes. We extend this literature by directly measuring the rel-

[5] Armona et al. (2019) ask households to rate their uncertainty over past house price growth on a 1-5 scale, but do not relate this quantitatively to the variance in beliefs.

[6] An alternative method infers perceived inflation persistence from the ratio of expectations at different horizons (Reis, 2020; Andre et al., 2021). However, this assumes that all agents share the same long-run expected inflation.

evant heterogeneity, and showing that the resulting distribution of expectation processes amplifies the aggregate consumption response to inflation shocks by an order of magnitude.

Section 2 outlines a model framework for expectations. Section 3 describes our data and novel survey questions, and Section 4 presents the empirical results. Section 5 calibrates the model to the empirical results. Section 6 concludes.

2 Expectations Framework

In this section we outline a simple theoretical framework for expectations, based on agents receiving noisy signals about inflation, which they believe follows an AR(1) process. We demonstrate how the parameters measured in our survey questions allow us to identify all components of the equations governing the evolution of expectations at the individual level. Different combinations of these components imply different responses of expected inflation to shocks, and cross-sectional heterogeneity in the components has substantial effects on the behaviour of aggregate expectations.

2.1 The Agent

As above, suppose that each agent i believes that inflation follows an AR(1) process:

$$\begin{aligned} \pi_t &= \hat{\rho}_i \pi_{t-1} + \varepsilon_t \\ \varepsilon_t &\sim N(0, \tilde{\sigma}^2_{\varepsilon,i}) \end{aligned} \tag{3}$$

This is their subjective law of motion for inflation, which may not coincide with the true data generating process.

Each period, the agent receives a noisy signal $s_{i,t}$ about current inflation:

$$\begin{aligned} s_{i,t} &= \pi_t + q_{i,t} \\ q_{i,t} &\sim N(0, \sigma^2_q) \end{aligned} \tag{4}$$

To allow for over- or under-confidence in private information (Daniel et al., 1998; Broer and Kohlhas, 2019), suppose that agents perceive the distribution of $q_{i,t}$ to be:

$$q_{i,t} \sim N(0, \tilde{\sigma}^2_{q,i}) \tag{5}$$

where $\tilde{\sigma}^2_{q,i}$ is not necessarily equal to σ^2_q.

After observing the signal, agents update their perception of inflation using the steady state Kalman filter.[7] We show in Appendix A that the evolution of the agent's one-period ahead inflation forecast is given by:

$$\tilde{E}_{i,t}\pi_{t+1} = (1-\chi_i)\hat{\rho}_i\tilde{E}_{i,t-1}\pi_t + \hat{\rho}_i\chi_i(\pi_t + q_{i,t}) \tag{6}$$

$$\chi_i = 1 - \frac{V_i^p}{V_i^f} \tag{7}$$

where χ_i is agent i's Kalman gain. V_i^p denotes the steady-state subjective variance of perceived inflation ($\tilde{Var}_{i,t}(\pi_t)$). Similarly, V_f^i denotes the steady state subjective variance of the one period ahead forecast ($\tilde{Var}_{i,t}(\pi_{t+1})$). These are such that $V_i^f \geq V_i^p$, so the uncertainty reduction about π_t when the agent receives $s_{i,t}$ is weakly positive.

This formula for χ_i is intuitive: the lower is V_i^p relative to V_i^f, the more informative is the signal received, and hence the larger the update the agent optimally makes to their inflation perceptions and forecasts. This formula also highlights that the forecast variance in a given period, as measured in several household surveys, does not place any restrictions on the Kalman gain. To the best of our knowledge our questions are the first attempts to directly identify V_i^p alongside V_i^f in survey data. Note that one cannot identify V_i^p from the cross-sectional dispersion of inflation perceptions without assuming that agents all share the same $\tilde{\sigma}^2_{q,i}$, contrary to the substantial evidence of information heterogeneity among households in particular (e.g. Link et al., 2021).

One can then identify $\tilde{\sigma}^2_{\varepsilon,i}$ using the law of motion for inflation:

$$\tilde{\sigma}^2_{\varepsilon,i} = V_i^f - \hat{\rho}_i^2 V_i^p \tag{8}$$

In addition, the perceived noise in the signal is given by (derivation in Appendix A):

$$\tilde{\sigma}^2_{q,i} = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{9}$$

Our novel survey questions that measure V_i^p and $\hat{\rho}_i$ therefore allow us to infer all parameters of the law of motion for household inflation expectations (equation 6), and the variances of both fundamental shocks and signal noise.

[7] The assumption of steady state filtering is required for identification of the agent's Kalman gain in the absence of panel data. The same assumption is commonly made in the rational inattention literature for analytical tractability (Maćkowiak and Wiederholt, 2009).

2.2 Expectations Impulse Responses

We now consider how inflation expectations respond to a shock to inflation. For simplicity, assume that inflation is an exogenous AR(1) process, with autocorrelation ρ:

$$\pi_t = \rho\pi_{t-1} + \varepsilon_t \tag{10}$$

We consider the impulse response to a one percentage point shock to inflation in time $t = 0$, with inflation and inflation expectations at steady state (zero) before the shock. Abstracting from the effect of realised signal noise $q_{i,t}$,[8] the one-period ahead inflation forecast of agent i, t periods after the shock, is:

$$\tilde{E}_{i,t}\pi_{t+1} = \hat{\rho}_i\chi_i \frac{\rho^{t+1} - (1-\chi_i)^{t+1}\hat{\rho}_i^{t+1}}{\rho - (1-\chi_i)\hat{\rho}_i} \tag{11}$$

See Appendix A for the derivation.

Different combinations of χ_i and $\hat{\rho}_i$ therefore imply very different impulse responses of expectations, even for the same V_i^f. The initial response of expectations to the shock is increasing in χ_i and $\hat{\rho}_i$. The persistence of the response of expectations also increases in $\hat{\rho}_i$; if $\hat{\rho}_i$ is low, the agent believes that inflation is not very persistent, and so largely neglects signals received before time t when forming their time t perception and forecast. Conversely, the persistence decreases in χ_i. If $\hat{\rho}_i$ is sufficiently large, and χ_i sufficiently small, then expectations display the hump-shaped impulse responses observed in previous empirical work (e.g. Angeletos et al., 2020). Equally, as in Angeletos et al. (2020), if $\hat{\rho}_i > \rho$ then expectations overshoot, rising above realized inflation some periods after the shock. See Appendix A for formal derivations of these results.

Figure 1 illustrates these points graphically, displaying impulse responses to the same inflation shock with four different processes for expectations. Panel (a), our baseline, features overextrapolation ($\hat{\rho}_i = 0.9$, $\rho = 0.7$). We set $V_i^f = 1$ and $V_i^p = 0.5$, giving $\chi_i = 0.5$. This parameterization gives a hump-shaped impulse response. In panel (b), we use the same parameters but set $V_i^p = 0$, so that $\chi_i = 1$. Although we still have overextrapolation, perceptions respond instantly to shocks, so there is no hump-shaped response. In panel (c), we set $V_i^p = 0.9$, giving $\chi_i = 0.1$. As in (a), we obtain hump-shaped impulse responses, but the peak response is much smaller. In panel (d), we use the baseline parameters but set a much lower $\hat{\rho}_i$, below the true ρ. As per the results above, we no longer see overshooting or hump-shaped IRFs.

[8] This is as if we are measuring the average expectation across many agents who share the same subjective law of motion and Kalman gain.

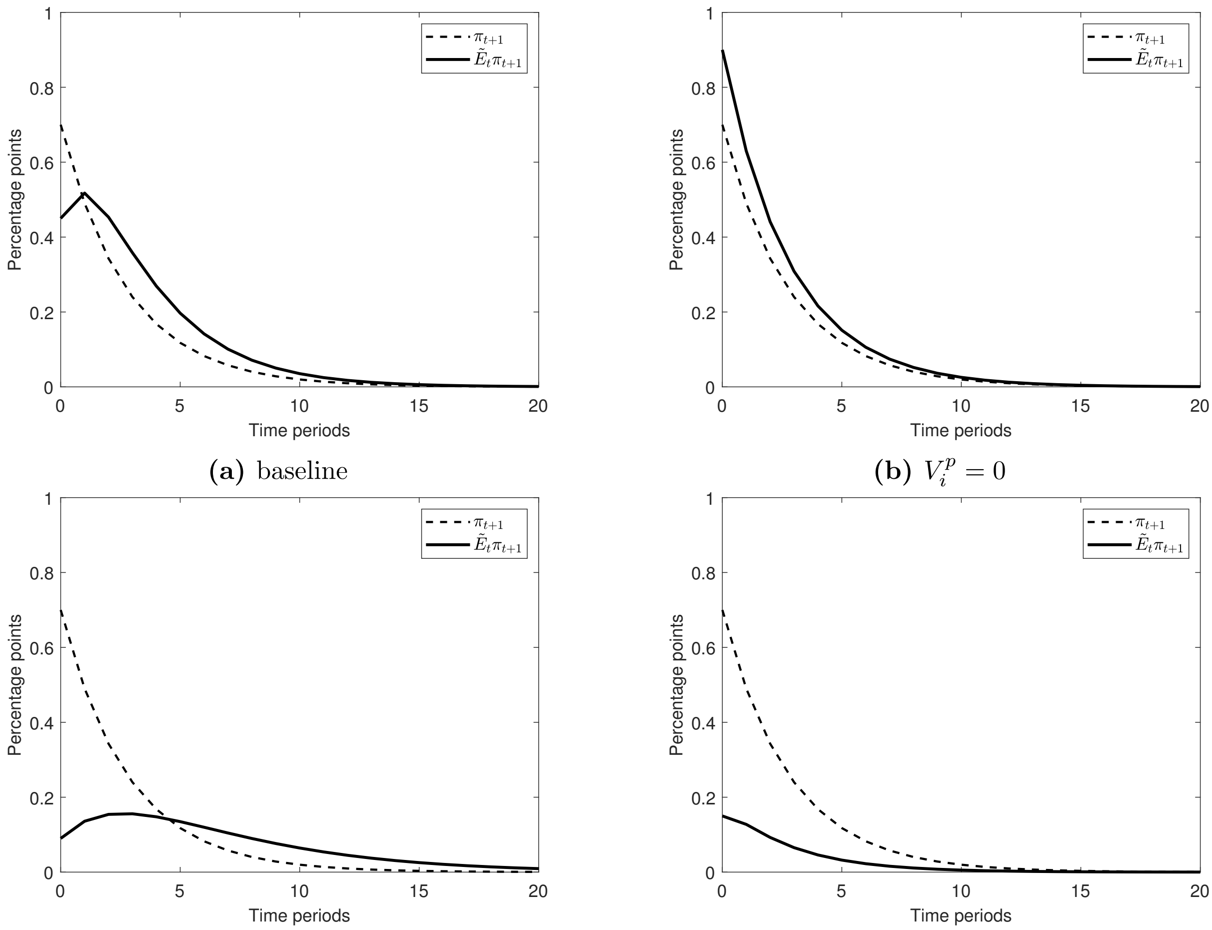

(a) baseline

(b) $V_i^p = 0$

(c) $V_i^p = 0.9$

(d) $\hat{\rho}_i = 0.3$

Figure 1: Impulse responses of one-period ahead inflation and one-period ahead forecasts to a 1 percentage point inflation shock in $t = 0$. Baseline calibration: $V_i^f = 1$, $V_i^p = 0.5$, $\hat{\rho}_i = 0.9$, $\rho = 0.7$.

2.3 The Role of Heterogeneity

If all agents used the same process to form expectations, equation 11 would also describe the impulse response of aggregate inflation expectations to the inflation shock. With heterogeneity in $\hat{\rho}_i$ and χ_i, however, the first period response of aggregate inflation expectations to the shock becomes:

$$\tilde{E}_0\pi_1 = E[\hat{\rho}_i]E[\chi_i] + Cov(\hat{\rho}_i, \chi_i) \quad (12)$$

A positive correlation between $\hat{\rho}_i$ and χ_i therefore amplifies the initial effect of the shock on inflation expectations. Intuitively, in this case those who extrapolate the most from perceived to expected future inflation also update their perceptions the most in the period of the shock. A negative correlation, conversely, dampens the initial response.

Heterogeneity continues to affect aggregate expectations in the periods after the shock. In the special case of $\rho = 0$, we have:

$$\tilde{E}_t\pi_{t+1} = E[\hat{\rho}_i^{t+1}]E[\chi_i(1-\chi_i)^t] + Cov(\hat{\rho}_i^{t+1}, \chi_i(1-\chi_i)^t) \quad (13)$$

To see the implications of heterogeneity here, consider the simple case where $\hat{\rho}_i$ is monotonically increasing in χ_i. The covariance term is then positive for small t, but negative for large t, since $\chi_i(1-\chi_i)^t$ is increasing in χ_i for $t = 0$, but decreasing in χ_i for sufficiently large t. As such, a positive correlation between $\hat{\rho}_i$ and χ_i tends to result in lower persistence of the response in inflation expectations, despite a larger initial response.

The joint distribution of $\hat{\rho}_i$ and χ_i therefore has rich implications for how aggregate expectations respond to shocks. Our empirical exercise allows us to measure this distribution, and so document the properties of these effects.

3 Data

We use data from the November 2021 wave of the Bundesbank Online Panel - Households (BOP-HH) survey, which is administered online to a representative sample of the German population. Our questions were asked to 4110 households.

The core modules of the survey contain several quantitative questions about perceived and expected inflation which are common to other household surveys. Households are asked for a point estimate of the inflation rate over the past 12 months, and for both point and density forecasts of inflation over the next 12 months, with the density forecast obtained by having households fill out the probabilities of inflation falling within a variety of ranges.[9] In addition, a range of household characteristics are collected. Summary statistics for a variety of these characteristics can be found in Appendix B.1.

To these core parts of the survey we add two novel questions for the November 2021 wave. Table 1 shows English translations of these questions (see Appendix B.2 for the German versions seen by respondents).

Table 1: Novel questions added to the BOP-HH survey in November 2020

Question	Text	Sample
1	Now we would like to know how certain you are about your information on the inflation rate or deflation rate over the past 12 months ([Value of point estimate])%. In your opinion, how likely is it that the inflation rate has been between [Low inflation level]% and [High inflation level]% over the past twelve months?	All respondents
Respondents randomly shown one of three scenarios before Question 2		
General	Imagine the following hypothetical situation: Due to an unexpected economic event, the inflation rate increased by one percentage point in the past year.	Group A
Supply	Imagine the following hypothetical situation: Due to unexpected problems with local production technology in the Middle East, the price of crude oil rose in the past year, causing the inflation rate to rise by one percentage point.	Group B
Demand	Imagine the following hypothetical situation: Due to increased defense spending, government spending rose unexpectedly more than usual in the past year, causing the inflation rate to rise by one percentage point. The change is temporary and occurs even though the government's assessment of national security or economic conditions has not changed. In addition, taxes do not change in response to the spending program.	Group C
2	In this situation, would you adjust your inflation expectations for the next 12 months as stated in the first part of the questionnaire? If so, to what extent?	All respondents

The first question is designed to uncover the uncertainty in the household's perceptions of current inflation.[10] The high and low inflation values seen by the respondent are determined by taking their point estimate of current inflation from the core survey, and

[9]These questions are very similar to, among others, the Survey of Consumer Expectations from the Federal Reserve Bank of New York (Armantier et al., 2017), and the ECB's Consumer Expectations Survey survey (Coibion et al., 2021).

[10]We do not use a question with multiple bins, as is done for expected future inflation, as these

adding/subtracting 1 percentage point. If the respondent's point perception of inflation is 5% or more, this range is widened to ±2 percentage points, as uncertainty in expected future inflation is known to rise with point estimates (Ben-David et al., 2018). Answers are in percent, and must be between 0 and 100. Respondents also see a note giving further explanation of the question (see Appendix B.2).

To go from this to the variance of perceived inflation that appears in equation 7, we fit a symmetric triangular distribution using their answer and their point estimate. This is similar to the approach taken by Coibion et al. (2021).[11] The variance of expectations in the Survey of Consumer Expectations is also found by fitting a triangular distribution if respondents only report positive probabilities in two bins of the density forecast questions (Armantier et al., 2017).

The second question elicits the respondent's perceived persistence of the inflation process, denoted $\hat{\rho}_i$ in the equations above. Following Andre et al. (2022), respondents are given a hypothetical situation describing an exogenous shock, and asked how they would expect that to affect inflation. In contrast to Andre et al. (2022), in all of the hypothetical scenarios we tell the respondents that the shock caused inflation to increase by 1 percentage point, so that their answers on how that would change their expected future inflation reflect their estimates of inflation persistence, not their predictions of the immediate impact of the shock, which Andre et al. (2022) find to be heterogeneous across their sample.

In Section 2 we did not distinguish between different kinds of shocks to inflation, and our main empirical analysis will do the same. However, in reality it may be the case that households associate different types of shock with different levels of persistence. To investigate this we randomly split the respondents into three groups. The first group are not told the nature of the inflationary shock, the second see a description of a hypothetical supply shock (oil price), and the third see a description of a demand shock (government spending). The specific hypothetical shocks are adapted from Andre et al. (2022).

To answer, respondents see the following:

1) Yes, from [Value of point estimate]% to __%
2) No

and either input a number in the first line or select 'No'.

Using these questions we obtain $\hat{\rho}_i$ and V_i^p for each respondent. Within the framework in Section 2, we can then back out the implied χ_i for each individual using equation 7.

questions are cognitively demanding. They can result in households dropping out of the survey if there are too many of them.

[11] Coibion et al. (2021) ask households for their most optimistic and pessimistic forecasts for GDP growth, then use those as the end points of an asymmetric triangular distribution, from which they extract the variance.

Using equations 8 and 9 we obtain $\tilde{\sigma}^2_{\varepsilon,i}$ and $\tilde{\sigma}^2_{q,i}$ for each respondent. Full details of the variable construction are in Appendix C.

4 Empirical results

In this section we report the marginal distributions of each component of household expectation formation, along with how they correlate with each other and with household characteristics. There is considerable heterogeneity in both $\hat{\rho}_i$ and χ_i. The correlations between elements of expectation formation, and with household characteristics, are broadly in line with models of rational inattention (Sims, 2003).

Figure 2a plots the CDF of $\hat{\rho}_i$, truncated to remove the approximately 1% of responses that fall outside the $[-5, 5]$ interval. Of the remaining responses, 11% give responses outside the $[0, 1]$ interval, with a substantial fraction giving answers greater than unity. 68% do not revise their expectations at all, indicating an important role for the extensive margin. This distribution gives our first main result:

Result 1 *Conditional on perceiving persistence in* $[0, 1]$*, the average perceived persistence is broadly in line with the data. However, the cross-sectional heterogeneity is large.*

Of those who give a response in $[0, 1]$, the mean $\hat{\rho}_i$ is 0.18, which is close to the "correct" answer of 0.2 based on German data for the past twenty years.[12] Including all responses in the interval $[-5, 5]$, the mean $\hat{\rho}_i$ is somewhat higher at 0.30.

Behind this average, the heterogeneity is substantial. Restricting to responses in [0,1], the maximum possible standard deviation would be 0.5: the standard deviation in our sample is 0.36.

If we consider only those who saw the first inflation scenario, in which the source of the shock was not specified, the mean $\hat{\rho}_i$ for those within the $[0, 1]$ interval is 0.16. The corresponding figures are 0.22 and 0.16 for those shown the supply and demand shocks respectively. Supply shocks are therefore perceived to be slightly more persistent than demand shocks, in line with evidence that supply shocks are of particular importance for consumer inflation expectations (Coibion and Gorodnichenko, 2015b).[13]

[12] We obtain annual CPI data from *www.destatis.de*, calculate the inflation rate, and run a linear projection of this on its lagged value using data from 2002-2021. The projection of year-on-year CPI inflation on the previous year's inflation rate yields a coefficient of 0.21.

[13] Part of this finding could stem from the fact that the demand shock scenario contains a stronger reference to the temporary nature of the shock. However, the unspecified scenario is closer to the supply shock in this respect, with neither stating explicitly that the shock is temporary. The fact that average $\hat{\rho}_i$ is higher in the supply shock scenario than both alternatives supports the interpretation that supply shocks are perceived to be more persistent.

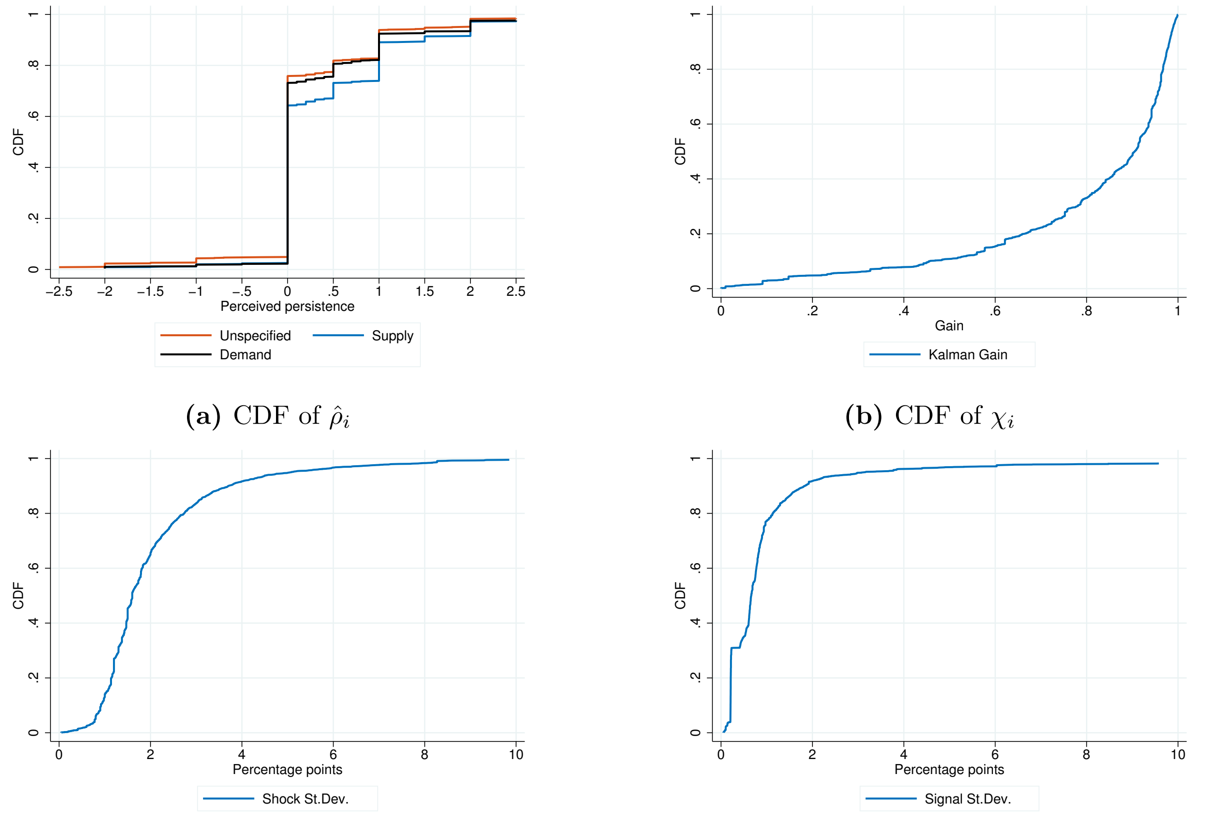

(a) CDF of $\hat{\rho}_i$

(b) CDF of χ_i

(c) CDF of $\tilde{\sigma}_{\varepsilon,i}$

(d) CDF of $\tilde{\sigma}_{q,i}$

Figure 2: CDFs of key parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

Once potential concern here is that the large numbers of households who do not adjust expectations following our scenarios do so because of a tendency to round to the nearest percentage point when reporting inflation expectations (Binder, 2017). To check whether this is driving the results, we restrict the sample to respondents whose answers to the initial inflation expectations question did not end in .0 or .5. We find that *c*.60% of respondents still do not revise their inflation expectations, very similar to the overall sample. Of those who do revise their expectations, we find that there are no longer obvious mass points close to 0.5 and 1.0. We report the CDF in Appendix D. We continue with the full sample here, as rounded expectations may still be important for consumption decisions.

The CDF of χ_i is shown in Figure 2b. The CDFs for V_i^p and V_i^f, which are used to calculate the Kalman gain, can be found in Appendix D.[14] This distribution gives our

[14]Note that for some consumers, we obtain a range of possible values for χ_i, because V_i^p is not pinned down exactly for households who are completely certain that current inflation is within the range shown in our question (see Appendix C for details). Figure 2b shows the mid-point in these cases, for all respondents with a range of width ≤ 0.2. The full CDFs of the upper and lower bounds on χ_i are shown in Appendix D. Figures 2c and 2d similarly show midpoints of implied ranges in these cases.

next main result:

Result 2 *The average implied Kalman gain is high, at 0.8. There is considerable cross-sectional heterogeneity.*

The high average Kalman gain stems from most consumers being considerably more certain about their inflation perceptions than their expectations. Approximately 40% of respondents reported that they were 100% certain that inflation lay within the range that they were presented with, and nearly four in five were at least 80% sure. There is, however, a long tail of very uncertain households, with very low implied Kalman gains, leading to a substantial variance in χ_i.

The average χ_i is substantially above the values obtained from regressions on average forecast errors, which are typically close to 0.5 (e.g. Coibion and Gorodnichenko, 2015a). The discrepancy is unsurprising, since Ryngaert (2018) shows that such regressions will yield biased estimates if agents hold inaccurate beliefs about the persistence of the inflation process. In addition, the survey was conducted in November 2021, when the rate of CPI inflation in Germany exceeded 5%. It is therefore plausible that consumers were particularly well informed about inflation at this point because of media coverage. Conversely, the fact that inflation was rapidly changing could have made consumers less certain in their precise estimate of the inflation rate.

The CDF of $\tilde{\sigma}_{\varepsilon,i}$ is shown in Figure 2c. Note that if households are receiving signals about future shocks to inflation (as in Goldstein and Gorodnichenko, 2022), then the calculations in this figure give a lower bound on the standard deviation of shocks. There is very considerable heterogeneity; a tail of households considerably overestimate the variability of future inflation. This finding again suggests that a large fraction of the uncertainty in inflation expectations relates to future developments, rather than uncertainty about the current inflation rate.

Figure 2d shows the CDF of $\tilde{\sigma}_{q,i}$. In keeping with the high average χ_i, most households have very little noise in their signals, though there is a long tail with very imprecise information.

4.1 Relationships between components of expectation formation

Different models of expectation formation give different implications for how the parameters we identify relate to each other, and to household characteristics. Table 2 shows our next main result:

Result 3 *Households who are more uncertain about current inflation are also more uncertain about future inflation, believe that the inflation process is noisier, and have lower implied Kalman gains.*

This is in line with models of noisy information, in which the households who are more uncertain about current inflation are uncertain because they process less information each period. However, this is only part of the reason for these households' high uncertainty about future inflation, as they also believe that there is more noise in the inflation process than less uncertain households.

Table 2: Pair-wise correlations of subjective law of motion elements.

	$SD_i(\pi_{t+1})$	$SD_i(\pi_t)$	$\hat{\rho}_i$	$SD_i(\varepsilon_{t+1})$	χ_i
$SD_i(\pi_{t+1})$	1.000				
$SD_i(\pi_t)$	0.473***	1.000			
$\hat{\rho}_i$	0.036*	-0.042**	1.000		
$SD_i(\varepsilon_{t+1})$	0.988***	0.440***	-0.025	1.000	
χ_i	0.305***	-0.402***	0.078***	0.327***	1.000

Note: Bundesbank-Online-Panel-Households, November 2021 wave. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range> 0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\hat{\rho}_i$ outside $[-5, 5]$ (c.1% of observations). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Households with high $\hat{\rho}_i$ are on average less uncertain about current inflation, and have higher Kalman gains. However, the results on $\hat{\rho}_i$ are potentially complicated by two factors. First, households may perceive that different shocks have different levels of persistence. Second, households who believe that inflation is non-stationary ($|\hat{\rho}_i| \geq 1$) may behave in systematically different ways to households with $|\hat{\rho}_i| < 1$. We explore these points in more detail in Appendix D. In general, the variation across shock types is small, but there are noticeable differences between those with stationary and non-stationary subjective laws of motion. Within households who believe that inflation is persistent and stationary, greater perceived persistence is associated with less uncertainty about current and future inflation, less perceived noise in the inflation process, and a greater implied Kalman gain. That is consistent with models of endogenous information acquisition; if inflation is more persistent, then information about the current inflation rate is more valuable.

4.2 Correlations of expectation components with household characteristics

Table 3 shows results of regressing each component of expectation laws of motion on a variety of household characteristics. As our application in Section 5 is to consumption, we focus on characteristics known in existing literature to relate to Marginal Propensities to Consume (MPCs). To account for the fact that households with $\hat{\rho}_i = 0$ may be systematically different from those with $\hat{\rho}_i \neq 0$, the final column shows the linear component of an estimated hurdle model, as in Cragg (1971). It therefore shows the estimated associations conditional on the household choosing to update their expectations in light of a small shock. Results from the selection stage, along with other ways of splitting the $\hat{\rho}_i$ responses, can be found in Appendix D.

The key variables displayed are liquid wealth (bank deposits plus securities), illiquid wealth (property plus firm ownership), other wealth, debt, and household income. There is also an indicator for if the household is hand-to-mouth, defined here as having liquid wealth of less than €1250, as this is known to be strongly positively correlated with a household's MPC (e.g. Kaplan et al., 2014). The units of the wealth and debt variables are €1000s.

The first row of coefficients shows our next main result:

Result 4 *Hand-to-mouth households are more uncertain about current inflation, but no more uncertain about future inflation, than other households. They believe inflation is noisier and more persistent, and have higher implied Kalman gains.*

The associations of the components of expectation laws of motion with hand-to-mouth status are large and strongly significant, with the exception of future inflation uncertainty. Hand-to-mouth households have a 14% higher $SD_i(\pi_t)$ than other households. More uncertainty about current, but not future inflation implies that they have a 9.5% larger Kalman gain on average. In addition, their perceived inflation persistence parameter is higher by an average of 0.4, which coupled with the higher current uncertainty implies that they believe the inflation process is substantially noisier than other households. On average hand-to-mouth households believe that $SD_i(\varepsilon_{t+1})$ is 15% higher than non hand-to-mouth households.

Angeletos et al. (2020) find that aggregate inflation expectations display delayed overshooting, as if households perceive inflation to be more persistent than implied by the true data generating process. This result suggests that a substantial portion of that overshooting may be driven by the expectations of hand-to-mouth households. While they have very large MPCs, they are also less able to respond to expectations of future inflation by adjusting consumption and saving decisions, so their expectations may not have as much impact as others' on the transmission of aggregate shocks. This reinforces that aggregate expectations do not fully uncover the dynamics that matter for business cycles. However, this result is not consistent with simple models of rational inattention, since in those models hand-to-mouth households who can't respond to expected inflation would be predicted to process less information as a result. One possibility is that the high Kalman gains for hand-to-mouth households are driven by those who are close to leaving their borrowing constraints, who have highly non-linear policy functions and value information highly as a result (Broer et al., 2021).

Table 3: Regressions of components of subjective laws of motion on household characteristics.

	(1) $\log(SD_i(\pi_{t+1}))$	(2) $\log(SD_i(\pi_t))$	(3) $\log(SD_i(\varepsilon_{t+1}))$	(4) $\log(\chi_i)$	(5) $\hat{\rho}_i$
Hand-to-mouth	0.0200	0.1374**	0.1477***	0.0951**	0.4051**
	(0.0360)	(0.0541)	(0.0570)	(0.0397)	(0.1999)
Liquid wealth	0.0000	-0.0002**	0.0002**	0.0002***	-0.0014***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0005)
Illiquid wealth	0.0000	-0.0000	-0.0000	0.0000	0.0002
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0002)
Other wealth	-0.0001	-0.0001	0.0003	-0.0000	-0.0010
	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0007)
Debt	0.0000	0.0001	0.0001	0.0002*	-0.0009
	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0006)
log(income)	-0.0775***	-0.1247***	-0.1736***	-0.0019	0.0756
	(0.0235)	(0.0348)	(0.0371)	(0.0332)	(0.1777)
HH Controls	Yes	Yes	Yes	Yes	Yes
Hurdle model	No	No	No	No	Yes
Observations	4382	3161	2292	2024	3194

Note: Bundesbank-Online-panel-Households, November 2021 wave. The household controls are age, age^2, gender, region, education, occupation category, and employment status. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Also consistent with this finding is the result that higher liquid wealth above the hand-to-mouth threshold is associated with less uncertainty in current inflation and less perceived inflation persistence, though these effects are quantitatively small in comparison. Each further €1000 is associated with a 0.02% reduction in $SD_i(\pi_t)$ and a 0.001 fall in $\hat{\rho}_i$. The lower current uncertainty relative to uncertainty over the future should imply that higher liquid wealth is associated with a lower Kalman gain, but in fact the reverse is true. This is the result of some households giving answers to Question 1 such that we cannot precisely impute their χ_i, so the samples are not equal across these regressions. However, although statistically significant, the association of χ_i with wealth beyond hand-to-mouth status is small in magnitude. In Appendix D we show that the positive coefficient on liquid wealth in the χ_i regression is driven by a small number of very wealthy households.

The final row of coefficients leads to our next main result:

Result 5 *Higher income is associated with less uncertainty about current and future inflation, and less perceived noise in the inflation process, but is not associated with differences in perceived inflation persistence or implied Kalman gains.*

As has been documented in other contexts (e.g. Ben-David et al., 2018), households with higher income are less uncertain about future inflation: a 10% rise in household income is associated with a 0.8% fall in $SD_i(\pi_{t+1})$. If this is due to high-income households acquiring more precise information about inflation, the relationship of income to uncertainty about current inflation should be even stronger (equation 7). There is evidence of that, as the same 10% rise in income is associated with a 1.2% fall in $SD_i(\pi_t)$. However, this difference across current and future uncertainty is not large, so there is no significant relationship between income and χ_i. Rather, the bulk of the lower uncertainty for high-income households is explained by them believing that the inflation process is itself less volatile. A 10% rise in income is associated with a 1.7% reduction in $SD_i(\varepsilon_{t+1})$. There is no significant correlation between income and perceived persistence.

5 Implications for Aggregate Consumption

In this section we show that the heterogeneity observed in the data has large effects on the aggregate consumption response to inflation shocks in an otherwise standard partial equilibrium consumption-saving model. Calibrating the model using the full distribution of expectations processes identified in our data yields aggregate consumption responses to inflation shocks which are an order of magnitude larger, and substantially more persistent, than those obtained in a representative-agent version of the model.

This amplification arises chiefly because the consumption function is very convex in perceived inflation persistence, which implies that heterogeneity in $\hat{\rho}_i$, and its covariance with other aspects of expectation formation, have large effects on aggregate shock responses. These conclusions therefore rely on the properties of standard consumption-saving models.[15] Dräger and Nghiem (2021) find that German households in a different

[15] A similar approach is taken by Roth et al. (2021), who also calibrate a standard partial equilibrium

survey behave in line with a consumption Euler equation of the form used here, suggesting this is reasonable. Similarly, Hanspal et al. (2021) find evidence that perceived persistence is important in determining household behaviour in the context of Covid-19 expectations.

5.1 Consumption-Saving Model

Unconstrained households have an infinite horizon and face no borrowing constraint. They choose consumption each period to maximise the expected discounted sum of CRRA utility over consumption, and invest any unspent exogenous income $y_{i,t}$ in risk-free one-period bonds with a gross nominal interest rate of i_t. The household's log-linearised consumption function is:[16]

$$\hat{c}_{i,t} = \sum_{h=0}^{\infty} \beta^h \Big((1-\beta)\tilde{E}_{i,t} y_{i,t+h} - \beta\gamma^{-1}\tilde{E}_{i,t} i_{t+h} + \beta\gamma^{-1}\tilde{E}_{i,t}\pi_{t+h+1} \Big) \tag{14}$$

where β is the consumer's discount factor and γ is the coefficient of relative risk aversion. To isolate the effect of a shock to expected inflation, we hold expected $y_{i,t+h}$ and i_{t+h} constant in all exercises here. With this assumption, and using the perceived law of motion for inflation (equation 3), consumption can be written as:[17]

$$\hat{c}_{i,t} = \frac{\beta\gamma^{-1}}{1-\beta\hat{\rho}_i}\tilde{E}_{i,t}\pi_{t+1} \tag{15}$$

This highlights that as well as increasing the size and persistence of the response of inflation expectations to the shock, a higher $\hat{\rho}_i$ also increases the responsiveness of consumption to those expectations, as it implies larger changes in longer-horizon expectations. Importantly, the effect of $\hat{\rho}_i$ on the scale of the response of consumption to inflation expectations is potentially very large in the case of overextrapolation, because $\beta/(1-\beta\hat{\rho}_i)$ is convex in $\hat{\rho}_i$.

Using the process for inflation expectations from Section 2 and aggregating across households, the aggregate consumption response to a one percentage point inflation shock is (details in Appendix A):

$$\hat{c}_t = E\left[\frac{\beta\gamma^{-1}}{1-\beta\hat{\rho}_i}\hat{\rho}_i\chi_i\frac{\rho^{t+1}-(1-\chi_i)^{t+1}\hat{\rho}_i^{t+1}}{\rho-(1-\chi_i)\hat{\rho}_i}\right] \tag{16}$$

consumption model using insights from survey data. While there are questions on consumption intentions in the BOP-HH survey, the sample answering both those questions and ours is very small.

[16] See Proposition 29 in the Online Appendix to Gabaix (2020).

[17] This result assumes that $|\beta\hat{\rho}_i| < 1$. In all calibrations of the model to the survey data we drop the minority of households for whom this is not the case.

In $t = 0$, this is equal to:

$$\hat{c}_0 = \beta\gamma^{-1}\left(E[\chi_i]E\left[\frac{\hat{\rho}_i}{1-\beta\hat{\rho}_i}\right] + Cov\left(\chi_i, \frac{\hat{\rho}_i}{1-\beta\hat{\rho}_i}\right)\right) \tag{17}$$

Heterogeneity in expectation laws of motion therefore affects aggregate consumption in two ways. First, heterogeneity in $\hat{\rho}_i$ substantially amplifies the initial aggregate consumption responses to inflation, because $\hat{\rho}_i/(1-\beta\hat{\rho}_i)$ is highly convex in $\hat{\rho}_i$. If even a small fraction of agents believe that inflation is close to a unit root, they respond very strongly to current inflation, generating large aggregate consumption responses. Note that heterogeneity in $\hat{\rho}_i$ increases the size of the consumption response *relative* to the response in inflation expectations. Similar aggregate impulse responses in inflation expectations may therefore correspond to very different effects of those expectations on consumption decisions.

Second, any correlation between $\hat{\rho}_i$ and χ_i will further distort the aggregate consumption response away from the representative-agent case. Intuitively, the response of aggregate consumption is amplified if the households who obtain precise information about the shock are also the ones who respond most strongly to that information. This is an example of the 'narrative heterogeneity channel' discussed in Macaulay (2022).

In all impulse responses below we also account for the fact that some households are hand-to-mouth, and so do not respond to expectations. For them, we set $\hat{c}_t = 0$ as we abstract from indirect effects of nominal shocks through incomes.

5.2 Model Implied Impulse Responses

We now generate IRFs of aggregate one-year ahead inflation expectations and consumption in three cases. First, we consider full information rational expectations (FIRE): all households know that the annual autocorrelation of inflation ρ is 0.21, and observe π_t precisely. Second, we maintain homogeneity, and calibrate the model using the population averages for χ_i and $\hat{\rho}_i$ in the survey data. Finally, we allow for heterogeneity, taking into account the joint distribution of χ_i and $\hat{\rho}_i$ from the data. In all cases we exclude observations where the perceived persistence lies outside [0,1].[18]

The implied IRFs for aggregate one-year ahead inflation expectations are shown in Figure 3a. The rational expectation responds by 0.21 on impact, because $\rho = 0.21$. Expectations respond less on impact in both of the cases using the survey data, because $E[\chi_i] < 1$ and $E[\hat{\rho}_i] \approx \rho$. On impact, expectations in these two cases are approximately the same. Aggregate inflation expectations are however somewhat more persistent un-

[18]This is done to exclude outliers, but note that Result 4 suggests the excluded households are disproportionately hand-to-mouth, so they would not respond to expectations anyway.

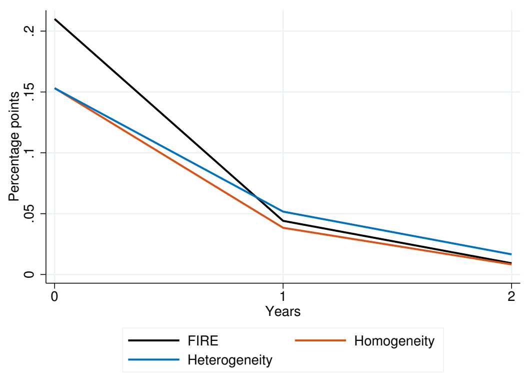

(a) Expectations

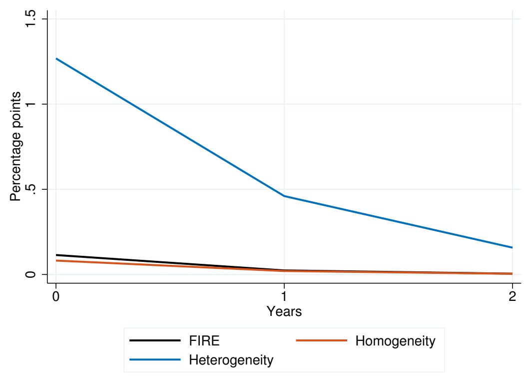

(b) Consumption

Figure 3: Implied IRFs of one-period ahead inflation expectations and consumption. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

der heterogeneity: a year after the shock, average $\tilde{E}_{i,t}\pi_{t+1}$ is c.35% greater than under homogeneity.

The IRFs for aggregate consumption are shown in Figure 3b. They differ a great deal between cases, which gives our final main result:

Result 6 *The model-implied consumption response under heterogeneity is 15.5x greater on impact than under homogeneity. The persistence of the consumption response under heterogeneity (that is, $\frac{\hat{c}_1}{\hat{c}_0}$) is c.45% greater than under homogeneity.*

On impact, the rational expectations response is small, at 0.11%. It has persistence of 0.2, and so the deviation from steady state in $t = 1$ is negligible. The homogeneous case has an even smaller initial response of consumption of 0.08%, and marginally greater persistence.

The heterogeneous case, however, has a vastly larger initial consumption response of 1.27%. This principally reflects the heterogeneity in the perceived persistence, which generates a great deal of amplification for any given initial response in inflation expectations. The response is then also more persistent, with a persistence of 0.35 between $t = 0$ and $t = 1$. Aggregate consumption therefore remains substantially above steady state in the two years following the shock.

Although the consumption function in equation 14 is standard in DSGE modeling, it does make strong assumptions about the way that households translate their expectations into consumption decisions, including that households have an infinite planning horizon, and are rational when making consumption decisions. Relaxing these assumptions would likely somewhat reduce the size of the consumption response to an inflation shock. We have also abstracted from the fact that consumers may expect changes in their real incomes or nominal interest rates in responses to inflation shocks. Our qualitative point, however, remains: if consumers are forward-looking in their consumption decisions, as is generally assumed in DSGE models, then assuming away heterogeneity in expectation formation is likely to be a poor assumption. Neglecting the considerable heterogeneity in the data can change the modeled response to shocks not by a few percentage points, but by an order of magnitude.

Figures 3a and 3b also provide a stark demonstration of the challenges involved in inferring how expectations affect the response of macroeconomic variables to shocks using only aggregate data. Comparing the homogeneity and heterogeneity cases, we see that similar IRFs in aggregate inflation expectations can relate to entirely different IRFs in aggregate consumption, depending on the degree of heterogeneity.

In Appendix E, we repeat this exercise separately for the distributions of expectation laws of motion observed for each separate shock scenario posed in Question 2 of the survey. Amplification is large in all cases, though the greater perceived persistence of supply shocks implies the largest amplification in that case. We also repeat the exercise excluding those whose answers are rounded to a multiple of 0.5, as in Section 4. Aggregate consumption still responds by 4.4× more on impact under heterogeneity even when excluding these households.

6 Conclusion

Inflation expectations form a crucial component of aggregate dynamics in most theories of the business cycle. However, we demonstrate that the quantities measured by existing surveys of expectations are consistent with multiple laws of motion for expectations, with very different aggregate implications. To separate these models, we use novel survey data to elicit (i) consumers' certainty in their inflation perceptions, and (ii) how persistent they perceive inflation to be, at the individual level.

We find that, on average, consumers are relatively confident in their inflation perceptions and perceive little persistence in inflation. Averages alone consequently suggest that inflation expectations have little role in generating amplification and persistence in the consumption responses to shocks. However, underlying these averages is a large degree of heterogeneity, which has major implications for shock transmission through expectations. Accounting for the observed heterogeneity in expectation laws of motion increases the aggregate consumption response to an inflation shock by an order of magnitude in an otherwise standard partial-equilibrium model. The persistence of consumption responses to shocks also increases substantially relative to the representative-agent case.

The key reason for this is that individual consumption functions are highly non-linear in the components of expectation formation, so heterogeneity in those parameters, and correlations between them, can lead to radically different model predictions. This nonlinearity in consumption functions also implies that the same impulse response for aggregate expectations can imply very different consumption responses, depending on the distribution of expectation laws of motion. The importance of heterogeneity is further underlined by our finding that the components of expectation formation are systematically correlated with whether the household is hand-to-mouth, and with household income, both of which are known to relate to consumption behaviour (Kaplan et al., 2014; Kueng, 2018). Exploring the distribution of these components for expectations of other variables, and how the distributions change over time and states of the world, could be a fruitful avenue for future research, with further implications for quantitative macroeconomic models.

References

Adam, K. (2007). Experimental Evidence on the Persistence of Output and Inflation. *The Economic Journal*, 117(520):603–636.

Afrouzi, H., Kwon, S. Y., Landier, A., Ma, Y., and Thesmar, D. (2020). Overreaction and working memory.

Andre, P., Haaland, I., Roth, C., and Wohlfart, J. (2021). Inflation Narratives. *ECONtribute Discussion Papers*, (127).

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample. *Review of Economic Studies*, (forthcoming).

Angeletos, G.-M., Huo, Z., and Sastry, K. A. (2020). Imperfect Macroeconomic Expectations: Evidence and Theory. *NBER Macroeconomics Annual*, 35(1):2–86.

Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W., and Zafar, B. (2015). Inflation expectations and behavior: do survey respondents act on their beliefs? *International Economic Review*, 56(2):505–536.

Armantier, O., Topa, G., van der Klaauw, W., and Zafar, B. (2017). An overview of the Survey of Consumer Expectations. *Federal Reserve Bank of New York Economic Policy Review*, 23(2).

Armona, L., Fuster, A., and Zafar, B. (2019). Home price expectations and behaviour: Evidence from a randomized information experiment. *Review of Economic Studies*, 86(4):1371–1410.

Badarinza, C. and Buchmann, M. (2009). Inflation perceptions and expectations in the euro area: the role of news. *ECB Working Paper Series*, (1088).

Ben-David, I., Fermand, E., Kuhnen, C. M., and Li, G. (2018). Expectations Uncertainty and Household Economic Behavior. *NBER Working Paper Series*, (25336).

Beshears, J., Choi, J. J., Fuster, A., Laibson, D., and Madrian, B. C. (2013). What goes up must come down? experimental evidence on intuitive forecasting. *American Economic Review*, 103(3):570–74.

Beutel, J. and Weber, M. (2021). Beliefs and Portfolios: Causal Evidence.

Bhandari, A., Borovicka, J., and Ho, P. (2019). Survey Data and Subjective Beliefs in Business Cycle Models. *Federal Reserve Bank of Richmond Working Papers*, 19(14):1–60.

Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90:1–12.

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

Branch, W. A. and Evans, G. W. (2006). Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory*, 127(1):264–295.

Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2021). Information and Wealth Heterogeneity in the Macroeconomy.

Broer, T. and Kohlhas, A. N. (2019). Forecaster (Mis-)Behavior.

Caplin, A. and Dean, M. (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, 105(7):2183–2203.

Coibion, O., Georgarakos, D., Gorodnichenko, Y., Kenny, G., and Weber, M. (2021). The Effect of Macroeconomic Uncertainty on Household Spending. *NBER Working Paper Series*, 28625.

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

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

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary Policy Communications and their Effects on Household Inflation Expectations. *Journal of Political Economy*, (forthcoming).

Cragg, J. G. (1971). Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods. *Econometrica*, 39(5):829.

Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (1998). Investor Psychology and Security Market Under- and Overreactions. *The Journal of Finance*, 53(6):1839–1885.

Dean, M. and Neligh, N. (2019). Experimental Tests of Rational Inattention.

Dominitz, J. and Manski, C. F. (2011). Measuring and Interpreting Expectations of Equity Returns. *Journal of Applied Econometrics*, 370(November 2010):352–370.

Dräger, L. and Nghiem, G. (2021). Are consumers' spending decisions in line with a euler equation? *Review of Economics and Statistics*, 103(3):580–596.

Fuster, A., Laibson, D., and Mendel, B. (2010). Natural Expectations and Macroeconomic Fluctuations. *Journal of Economic Perspectives*, 24(4):67–84.

Gabaix, X. (2020). A behavioral new Keynesian model. *American Economic Review*, 110(8):2271–2327.

Galati, G., Moessner, R., and van Rooij, M. (2021). Anchoring of consumers' long-term euro area inflation expectations during the pandemic.

Goldstein, N. and Gorodnichenko, Y. (2022). Expectations Formation and Forward Information. *NBER Working Paper Series*, (29711).

Gorodnichenko, Y. and Sergeyev, D. (2021). Zero Lower Bound on Inflation Expectations. *NBER Working Paper Series*, (29496).

Hajdini, I. (2020). Misspecified Forecasts and Myopia in an Estimated New Keynesian Model.

Hanspal, T., Weber, A., and Wohlfart, J. (2021). Exposure to the covid-19 stock market crash and its effect on household expectations. *Review of Economics and Statistics*, 103(5):994–1010.

Hommes, C. and Lustenhouwer, J. (2019). Inflation targeting and liquidity traps under endogenous credibility. *Journal of Monetary Economics*, 107:48–62.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity*, 45(1 (Spring)):77–153.

Kueng, L. (2018). Excess Sensitivity of High-Income Consumers. *The Quarterly Journal of Economics*, 133(4):1693–1751.

Laudenbach, C., Weber, A., and Wohlfart, J. (2021). Beliefs About the Stock Market and Investment Choices: Evidence from a Field Experiment. *CEBI Working Paper Series*, (17/21).

Link, S., Peichl, A., Roth, C., and Wohlfart, J. (2021). Information Frictions Among Firms and Households. *CESifo Working Paper*, 8969.

Macaulay, A. (2022). Heterogeneous information, subjective model beliefs, and the time-varying transmission of shocks.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2020). Rational Inattention: A Review. *CEPR Discussion Papers*, no. 15408.

Maćkowiak, B. and Wiederholt, M. (2009). Optimal Sticky Prices under Rational Inattention. *American Economic Review*, 99(3):769–803.

Mankiw, N. G., Reis, R., and Wolfers, J. (2004). Disagreement about Inflation Expectations. *NBER Macroeconomics Annual*, 18:209–270.

Reis, R. (2020). Comment on "Imperfect Expectations: Theory and Evidence". *NBER Macroeconomics Annual*, 35.

Roth, C., Wiederholt, M., and Wohlfart, J. (2021). The Effects of Forward Guidance: Theory with Measured Expectations. *ECONtribute Discussion Papers*, (126).

Ryngaert, J. (2018). What Do (and Don't) Forecasters Know about U.S. Inflation?

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

Vellekoop, N. and Wiederholt, M. (2019). Inflation Expectations and Choices of Households. *SAFE Working Paper*, (250).

A Proofs

Steady-State Kalman Filter

Before receiving the signal at time t, the agent's subjective joint distribution for π_t and the signal received is:

$$\begin{pmatrix} \pi_t \\ \pi_t + q_{i,t} \end{pmatrix} \sim N\left(\begin{pmatrix} \hat{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} \\ \hat{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} \end{pmatrix}, \begin{pmatrix} V^f_{i,t-1} & V^f_{i,t-1} \\ V^f_{i,t-1} & V^f_{i,t-1} + \tilde{\sigma}^2_{q,i} \end{pmatrix} \right) \tag{18}$$

The conditional mean of π_t given the signal is then:

$$\begin{aligned} \tilde{E}_{i,t}\pi_t &= \hat{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} + \frac{V^f_{i,t-1}}{V^f_{i,t-1} + \tilde{\sigma}^2_{q,i}}(\pi_t + q_{i,t}) \\ &= (1 - \chi_{i,t})\hat{\rho}_i \tilde{E}_{i,t-1}\pi_{t-1} + \chi_{i,t}(\pi_t + q_{i,t}) \end{aligned} \tag{19}$$

where:

$$\chi_{i,t} = \frac{V^f_{i,t-1}}{V^f_{i,t-1} + \tilde{\sigma}^2_{q,i}} \tag{20}$$

The conditional variance of π_t

$$\begin{aligned} V^p_{i,t} &= V^f_{i,t-1}\left(1 - \frac{(V^f_{i,t-1})^2}{V^f_{i,t-1}(V^f_{i,t-1} + \tilde{\sigma}^2_{q,i})}\right) \\ &= \frac{V^f_{i,t-1}\tilde{\sigma}^2_{q,i}}{V^f_{i,t-1} + \tilde{\sigma}^2_{q,i}} \end{aligned} \tag{21}$$

In steady state, the variance of the perception is:

$$V^p_i = \frac{V^f_i \tilde{\sigma}^2_{q,i}}{V^f_i + \tilde{\sigma}^2_{q,i}} \tag{22}$$

The steady state Kalman gain is then:

$$\begin{aligned} \chi_i &= \frac{V^f_i}{V^f_i + \tilde{\sigma}^2_{q,i}} \\ &= 1 - \frac{V^p_i}{V^f_i} \end{aligned} \tag{23}$$

Response of Inflation Expectations to Shocks

The individual inflation perception is given by equation 19. Iterating backwards to time 0, we obtain:

$$\tilde{E}_{i,t}\pi_t = \chi_i \sum_{s=0}^{t} ((1-\chi_i)\hat{\rho}_i)^s(\pi_{t-s} + q_{i,t-s}) \tag{24}$$

The h period ahead forecast is then:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i \hat{\rho}_i^h \sum_{s=0}^{t} ((1-\chi_i)\hat{\rho}_i)^s(\pi_{t-s} + q_{i,t-s}) \tag{25}$$

Abstracting from $q_{i,t}$, the forecast is simply:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i \hat{\rho}_i^h \sum_{s=0}^{t} ((1-\chi_i)\hat{\rho}_i)^s \pi_{t-s} \tag{26}$$

Then noting that $\pi_{t-s} = \rho^{t-s}\varepsilon_0$, we obtain:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i \hat{\rho}_i^h \rho^t \sum_{s=0}^{t} ((1-\chi_i)\hat{\rho}_i \rho^{-1})^s \varepsilon_0 \tag{27}$$

Provided that $(1-\chi_i)\hat{\rho}_i\rho^{-1} \neq 1$, then evaluating the summation and rearranging yields the result:

$$\tilde{E}_{i,t}\pi_{t+h} = \chi_i \hat{\rho}_i^h \frac{\rho^{t+1} - ((1-\chi_i)\hat{\rho}_i)^{t+1}}{\rho - (1-\chi_i)\hat{\rho}_i} \varepsilon_0 \tag{28}$$

Setting $h = 1, \varepsilon_0 = 1$ yields equation 11.

Persistence of Expectations

Using equation 28, note that:

$$\frac{\tilde{E}_{i,1}\pi_2}{\tilde{E}_{i,0}\pi_1} = \rho + (1-\chi_i)\hat{\rho}_i$$

As such, the persistence of the response of expectations to the shocks is increasing in ρ and $\hat{\rho}_i$, and decreasing in χ_i (assuming that $\hat{\rho}_i > 0$). If $\rho + (1-\chi_i)\hat{\rho}_i > 1$, then the expectation rises between the period that the shock hits and the period after, giving a hump-shaped response. This condition is both necessary and sufficient for a hump-shaped impulse response.

Consumption responses

Consider an unconstrained agent facing an infinite-horizon consumption savings problem. As in Gabaix (2020), take income as given for simplicity. The consumption function is given by:

$$\hat{c}_{i,t} = \sum_{h \geq 0} \beta^h ((1-\beta)\tilde{E}_{i,t}\hat{y}_{i,t+h} - \beta\gamma^{-1}\tilde{E}_{i,t}i_{t+h} + \beta\gamma^{-1}\tilde{E}_{i,t}\pi_{t+h+1}) \tag{29}$$

Since we hold expected income and nominal interest rates at steady state, we have $\tilde{E}_{i,t}\hat{y}_{i,t+h} = 0$ and $\tilde{E}_{i,t}i_{t+h} = 0$ for all t and h. The consumption function then reduces to:

$$\hat{c}_{i,t} = \beta\gamma^{-1} \sum_{h \geq 0} \beta^h \tilde{E}_{i,t}\pi_{t+h+1} = \beta\gamma^{-1} \frac{1}{1-\beta\hat{\rho}_i} \tilde{E}_{i,t}\pi_{t+1} \tag{30}$$

To proceed, substitute in for the one period ahead expectation in time t using equation 28:

$$\hat{c}_{i,t} = \beta\gamma^{-1}\frac{1}{1-\beta\hat{\rho}_i}\hat{\rho}_i\chi_i\frac{\rho^{t+1}-(1-\chi_i)^{t+1}\hat{\rho}_i^{t+1}}{\rho-(1-\chi_i)\hat{\rho}_i} + d_{i,t} \tag{31}$$

Here, $d_{i,t}$ is an idiosyncratic noise term, which is a linear function of $q_{i,t}$, $q_{i,t-1}$, ...,$q_{i,0}$, and so has mean zero. Averaging across agents, one obtains:

$$\hat{c}_t = \beta\gamma^{-1}E\left[\frac{1}{1-\beta\hat{\rho}_i}\hat{\rho}_i\chi_i\frac{\rho^{t+1}-(1-\chi_i)^{t+1}\hat{\rho}_i^{t+1}}{\rho-(1-\chi_i)\hat{\rho}_i}\right] \tag{32}$$

Note that in $t=0$, this becomes:

$$\hat{c}_0 = \beta\gamma^{-1}E\left[\frac{\hat{\rho}_i}{1-\beta\hat{\rho}_i}\chi_i\right] \tag{33}$$

Applying the definition of a covariance then leads to equation 17.

B Survey details

B.1 Summary statistics

Table 4 shows summary statistics for the key variables used in our analysis, and several other household characteristics. The construction of $SD_i(\pi_t), SD_i(\pi_{t+1}), SD_i(\varepsilon_{t+1}), \chi_i, \hat{\rho}_i$ is described in Appendix C.

Income and wealth (of different types) is reported in bins in the survey. We take the mid-point of each bin. We code the lowest bin for income as if the lower bound is zero, and again take the midpoint (all wealth variables have a separate bin for zero). The top bin is coded as if it had the same width as the second-highest bin. Liquid wealth is defined as the sum of bank deposits and securities. Illiquid wealth is defined as the sum of property and firm ownership. Debt is defined as the sum of secured and unsecured debt. A respondent is classified as hand-to-mouth if their liquid wealth is below €1250.

B.2 Novel questions for November 2020

The novel questions as they appeared to the respondents in German are recorded below.

Question 1

Nun möchten wir wissen, wie sicher Sie sich über Ihre Angabe zur Inflationsrate oder Deflationsrate in den letzten 12 Monaten sind ([Value of point estimate])%.

Table 4: Summary statistics for expectations point estimates and components, respondent characteristics, and income/wealth.

	Mean	Std Dev.	Min	Max
Panel A: Expectations				
$\tilde{E}_{i,t}\pi_t$	4.10	2.61	0	30
$\tilde{E}_{i,t}\pi_{t+1}$	5.06	5.14	1	60
$SD_i(\pi_t)$	1.75	4.70	0.41	40.41
$SD_i(\pi_{t+1})$	1.71	1.35	0.30	8.80
$SD_i(\varepsilon_{t+1})$	1.87	1.51	0.04	12.12
χ_i	0.80	0.23	0	1
$\hat{\rho}_i$	0.30	0.84	-5	5
Panel B: Demographics				
Age	56.87	14.66	16	80
Female	0.37	0.48	0	1
Higher Education	0.39	0.49	0	1
Is Working	0.55	0.50	0	1
Panel C: Income and Wealth				
Income	3.95	1.97	0.25	11
Liquid Wealth	90.49	154.90	0	1250
Illiquid Wealth	315.38	383.64	0	2375
Other Wealth	12.41	48.75	0	625
Debt	47.89	109.34	0	955
Owns Securities	0.62	0.48	0	1
Hand-to-mouth	0.14	0.34	0	1

Note: Bundesbank-Online-Panel-Households, November 2021 wave. For cases where χ_i is set-identified, respondents are excluded if the parameters are estimated very imprecisely (range$>$ 0.2). For all remaining set-identified parameters, the mid-point of the range is used. Observations of $\tilde{E}_{i,t}\pi_t, \tilde{E}_{i,t}\pi_{t+1}, SD_i(\pi_{t+1})$, and $SD_i(\pi_t)$ below the 1st or above the 99th percentile of that variable's distribution are also excluded as outliers, as are observations of $\hat{\rho}_i$ outside $[-5, 5]$ (c.1% of observations). All income and wealth variables are in €1000s, and income refers to monthly net income of the household. Higher Education is an indicator for if the respondent has a bachelor's degree or higher, not including vocational training.

Frage: Wie wahrscheinlich ist es Ihrer Meinung nach, dass die Inflationsrate in den letzten zwölf Monaten zwischen [Low inflation level]% und [High inflation level]% lag?
Hinweis: Bei dieser Frage geht es darum, wie Sie die Wahrscheinlichkeit einschätzen, dass die von Ihnen angegebene Inflationsrate oder Deflationsrate in den letzten 12 Monaten tatsächlich ungefähr diesen Wert angenommen hat. Ihre Antworten können zwischen 0 und 100 liegen, wobei 100 bedeutet, dass Sie absolut sicher sind. Kleinere Zahlen bedeuten, dass Sie sich weniger sicher sind.
Input field Prozent

Question 2

Input filter: if Group A
Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund eines unerwarteten wirtschaftlichen Ereignisses hat sich die Inflationsrate im vergangenen Jahr um einen Prozentpunkt erhöht.
Input filter: if Group B
Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund von uner-warteten Problemen mit der lokalen Produktionstechnologie im Nahen Osten ist der Rohölpreis im vergangenen Jahr gestiegen, was zu einem Anstieg der Inflationsrate um einen Prozentpunkt geführt hat.
Input filter: if Group C
Stellen Sie sich die folgende hypothetische Situation vor: Aufgrund gestiegener Verteidigungsausgaben sind die Staatsausgaben im vergangenen Jahr unerwartet stärker als üblich gestiegen, was zu einem Anstieg der Inflationsrate um einen Prozentpunkt geführt hat.
Die Änderung ist vorübergehend und tritt ein, obwohl sich die Einschätzung der Regierung zur nationalen Sicherheit oder den wirtschaftlichen Bedingungen nicht geändert hat. Darüber hinaus ändern sich die Steuern nicht als Reaktion auf das Ausgabenprogramm.
Frage: Würden Sie in dieser Situation Ihre im vorderen Teil des Fragebogens genannten Inflationserwartungen für die nächsten 12 Monate anpassen? Wenn ja, inwiefern?

1) Ja, von [Value of point estimate]Prozent auf __Prozent
2) Nein

C Variable construction

To obtain V_i^p, we take answers to Question 1 and apply the following formula, which takes the variance of a symmetric triangular distribution fitted to household i's answers.

$$V_i^p = \begin{cases} \frac{1}{6}\left(1 - \sqrt{1 - \frac{x_{1i}}{100}}\right)^{-2} & \text{if } \mathbb{E}_i(\pi_t) \in (-5, 5) \\ \frac{2}{3}\left(1 - \sqrt{1 - \frac{x_{1i}}{100}}\right)^{-2} & \text{if } \mathbb{E}_i(\pi_t) \notin (-5, 5) \end{cases} \tag{34}$$

where x_{1i} is respondent i's response to Question 1. Note that for households who report $x_i = 0$, this method provides an upper bound on their $Var_i(\pi_t)$.

To obtain $\hat{\rho}_i$, we set $\hat{\rho}_i = 0$ for all households who select 'No' in answer to Question 2. For all other households, we set:

$$\hat{\rho}_i = x_{2i} - \mathbb{E}_i(\pi_{t+1}) \tag{35}$$

where x_{2i} is respondent i's response to Question 2.

We then calculate V_i^f. For those agents who are certain that future inflation will lie within one specific bin, we calculate an upper bound on the variance using the symmetric triangular distribution, just as for the perception. The lower bound on V_i^f is given by zero.

For the remaining agents, we calculate V_i^f by taking the midpoints of each of the bins in the probability distribution. Denote these midpoints as z_j for the bins $j = 1, ..., n$. Denote the probability assigned to each bin as p_j. We then calculate the mean:

$$\bar{z}_i = \sum_{j=1}^{n} p_{i,j} z_j \tag{36}$$

The variance is then:

$$V_i^f = \sum_{j=1}^{n} p_{i,j}(z_j - \bar{z}_i) \tag{37}$$

The calculation of the Kalman gain is complicated by the fact that that for some respondents we have ranges of possible V_i^p or V_i^f, in which case we can only find ranges for χ_i and the other key parameters. We now describe how we calculate these parameters for each such case.

Case (i): V_i^p and V_i^f both point-identified

Calculate χ_i using:

$$\chi_i = 1 - \frac{V_i^p}{V_f^i} \tag{38}$$

Back out $\tilde{\sigma}^2_{\varepsilon,i}$ using:

$$\tilde{\sigma}^2_{\varepsilon,i} = V_i^f - \hat{\rho}_i^2 V_i^p \tag{39}$$

Finally, calculate $\tilde{\sigma}^2_{q,i}$ using:

$$\tilde{\sigma}^2_{q,i} = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{40}$$

Datapoints are inconsistent with Kalman filtering (and so are dropped) if $\chi_i < 0$ or $\tilde{\sigma}^2_{\varepsilon,i} < 0$.

Case (ii): V_i^f point-identified, V_i^p set-identified

This case occurs if the respondent is certain that π_t lies within the specified interval, but places strictly positive probability in multiple intervals in the expectation question. V_i^p is then bounded below by zero, and the upper bound is calculated using the symmetric triangular distribution as above.

Denote the upper bound on V_i^p by a, so that $V_i^p \in [0, a]$. Under steady state Kalman filtering, it must be that $V_i^p \leq V_i^f$ and $V_i^p \leq \hat{\rho}_i^{-2} V_i^f$. The latter is more restrictive if $|\hat{\rho}_i| > 1$. This may shrink the upper bound on V_i^p, and hence raise the lower bound on the Kalman filter. As such, $V_i^p \in [0, \tilde{a}]$, where $\tilde{a}$ is given by:

$$\tilde{a} = \min(V_i^f, \hat{\rho}_i^{-2} V_i^f) \tag{41}$$

Then we have the following ranges for the key parameters:

$$\chi_i \in \left[1 - \frac{\tilde{a}}{V_i^f}, 1\right] \tag{42}$$

$$\tilde{\sigma}^2_{q,i} \in \left[0, \frac{V_i^f \tilde{a}}{V_i^f - \tilde{a}}\right] \tag{43}$$

$$\tilde{\sigma}^2_{\varepsilon,i} \in \left[V_i^f - \hat{\rho}_i^2 \tilde{a}, V_i^f\right] \tag{44}$$

Case (iii): V_i^f set-identified, V_i^p point-identified

In this case, the consumer is not certain that current inflation lies within the specified interval, but is certain that future inflation lies within one specific interval. As such, V_i^p is known, but $V_i^f \in [0, b]$, where b is given by the symmetric triangular distribution.

Under steady state Kalman filtering, it must be the case that $V_i^f \geq V_i^p$ and $V_i^f \geq \hat{\rho}_i^2 V_i^p$. Hence $V_i^f \in [\tilde{b}, b]$, where:

$$\tilde{b} = \max(V_i^p, \hat{\rho}_i^2 V_i^p) \tag{45}$$

Note that if $\tilde{b} > b$, then the observations must be dropped as they are inconsistent with steady state Kalman filtering. Using the equation for the Kalman gain, we then have:

$$\chi_i \in \left[1 - \frac{V_i^p}{\tilde{b}}, 1 - \frac{V_i^p}{b}\right] \tag{46}$$

The variance of the signal then lies in the interval:

$$\tilde{\sigma}_{q,i}^2 \in \left[\frac{bV_i^p}{b - V_i^p}, \frac{\tilde{b}V_i^p}{\tilde{b} - V_i^p}\right] \tag{47}$$

Note that if $\tilde{b} = V_i^p$, then the upper end of this interval is infinite, implying the signal may be infinitely noisy (i.e. contains no information).

Finally, the perceived variance of the shock lies in the range:

$$\tilde{\sigma}_{\varepsilon,i}^2 \in [\tilde{b} - \hat{\rho}_i^2 V_i^p, b - \hat{\rho}_i^2 V_i^p] \tag{48}$$

Case (iv): V_i^f and V_i^p both set-identified

In this case, the consumer is certain that current inflation lies within the specified interval, and certain that future inflation will lie within one specific interval. Hence, we have $V_i^p \in [0, a]$ and $V_i^f \in [0, b]$. If $|\hat{\rho}_i| \leq 1$, then χ_i is unrestricted within the interval $[0, 1]$. If $|\hat{\rho}_i| > 1$, then χ_i is bounded below as described above. Hence $\chi_i \in [0, 1]$ if $|\hat{\rho}_i| \leq 1$, and $\chi_i \in [1 - \hat{\rho}_i^{-2}, 1]$ if $|\hat{\rho}_i| > 1$.

We then know that:

$$\tilde{\sigma}_{q,i}^2 = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \tag{49}$$

If $|\hat{\rho}_i| < 1$, this can take any value. It could be infinite large if $V_i^p = V_i^f$, and could be zero if $V_i^p = 0$ but $V_i^f > 0$. If $|\hat{\rho}_i| > 1$, then $V_i^f \geq \hat{\rho}_i^2 V_i^p$. In that case, $\tilde{\sigma}_{q,i}^2$ could still be zero, but the maximum value it can now take is:

$$\tilde{\sigma}^2_{q,i} = \frac{V_i^f V_i^p}{V_i^f - V_i^p} \leq \frac{V_i^f V_i^p}{\hat{\rho}_i^2 V_i^p - V_i^p} \tag{50}$$

$$= \frac{V_i^f}{\hat{\rho}_i^2 - 1} \leq \frac{b}{\hat{\rho}_i^2 - 1} \tag{51}$$

To summarize, then, $\tilde{\sigma}^2_{q,i} \in [0, \infty)$ if $|\hat{\rho}_i| \leq 1$, and $\tilde{\sigma}^2_{q,i} \in [0, \frac{b}{\hat{\rho}_i^2 - 1}]$ if $|\hat{\rho}_i| > 1$.

Turning to $\tilde{\sigma}^2_{\varepsilon,i}$, this could always be zero in this case. The maximum it could be is b if $V_i^f = b$ and $V_i^p = 0$. Hence $\tilde{\sigma}^2_{\varepsilon,i} \in [0, b]$.

D Additional empirical results

D.1 Additional parameter distributions

Figure 4 plots the CDFs of V_i^p and V_i^f. In cases where these are only set-identified, this plots the upper bound from fitting a symmetric triangular distribution. The lower bound in all such cases is 0. On average households are substantially less uncertain about current inflation than about future inflation.

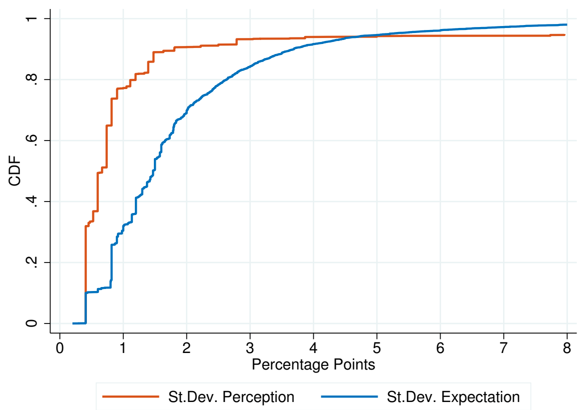

Figure 4: CDFs of V_i^p and V_i^f. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

For those respondents where we can only identify ranges for V_i^p and V_i^f, we can similarly only identify bounds for the inferred parameters $\chi_i, \tilde{\sigma}^2_{q,i}, \tilde{\sigma}^2_{\varepsilon,i}$. Figure 5 shows the distributions of these parameters if we take the upper or lower bounds of the parameter ranges for those households respectively.

Figure 6 plots the distribution of $\hat{\rho}_i$ when we exclude respondents whose response to the initial inflation expectations question ends in .0 or .5. Even excluding these households

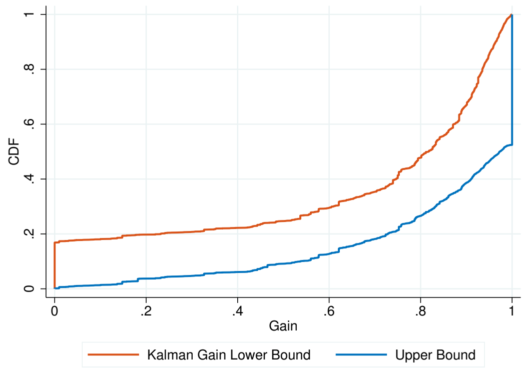

(a) χ_i

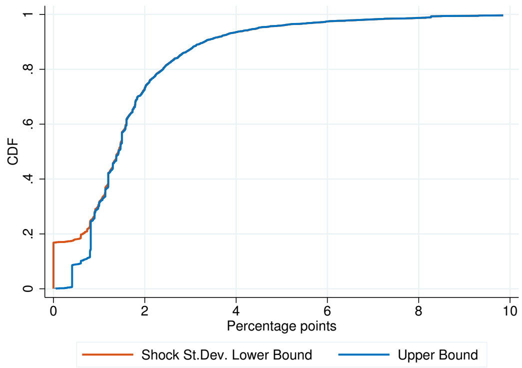

(b) $\tilde{\sigma}_{\varepsilon,i}$

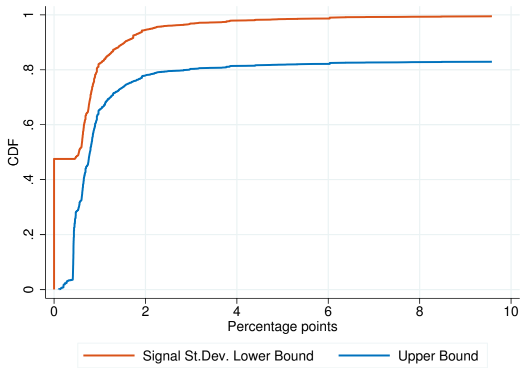

(c) $\tilde{\sigma}_{q,i}$

Figure 5: CDFs of upper and lower bounds for inferred parameters. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

with the strongest tendency to round their answers, there is a large mass with $\hat{\rho}_i = 0$, and substantial heterogeneity.

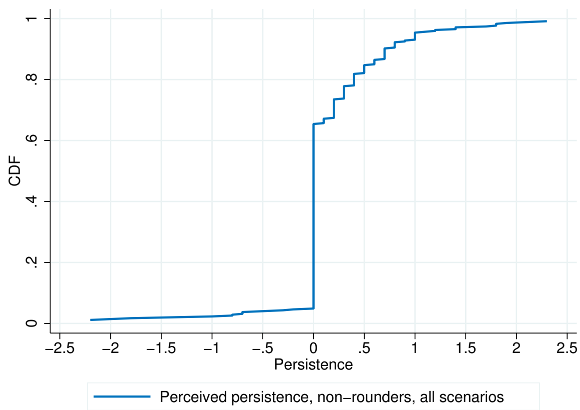

Figure 6: CDF of perceived persistence, only respondents whose response to initial inflation expectations question does not end in .0 or .5. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

Table 5: Means of elements of expectation laws of motion, by inflation perception

	$SD_i(\pi_t)$	$SD_i(\pi_{t+1})$	$SD_i(\varepsilon_{t+1})$	χ_i	$\hat{\rho}_i$
$\tilde{E}_{i,t}\pi_t < 0$	1.07	2.60	2.53	0.67	-0.13
	(0.27)	(0.52)	(0.48)	(0.13)	(0.30)
$\tilde{E}_{i,t}\pi_t \in [0, 2)$	0.76	1.89	1.83	0.71	0.26
	(0.04)	(0.16)	(0.16)	(0.04)	(0.08)
$\tilde{E}_{i,t}\pi_t \in [2, 4)$	0.66	1.83	1.77	0.79	0.26
	(0.01)	(0.03)	(0.03)	(0.01)	(0.02)
$\tilde{E}_{i,t}\pi_t \in [4, 6)$	0.70	2.12	2.06	0.82	0.28
	(0.01)	(0.04)	(0.04)	(0.01)	(0.02)
$\tilde{E}_{i,t}\pi_t \in [6, 8)$	1.44	3.35	3.22	0.74	0.22
	(0.10)	(0.19)	(0.19)	(0.03)	(0.09)
$\tilde{E}_{i,t}\pi_t \in [8, 10)$	1.44	3.95	3.89	0.81	0.17
	(0.17)	(0.44)	(0.41)	(0.07)	(0.17)
$\tilde{E}_{i,t}\pi_t > 10$	2.11	4.83	4.73	0.72	0.04
	(0.18)	(0.29)	(0.30)	(0.04)	(0.07)

Note: Bundesbank-Online-panel-Households, November 2021 wave. Standard errors in parentheses.

D.2 Relationship between components of expectation formation and point estimates

Table 5 shows the means of the elements of the expectation laws of motion, broken down by the respondent's reported inflation perception ($\tilde{E}_{i,t}\pi_t$). Note that this only includes respondents whose responses are consistent with steady-state Kalman filtering. Those with an inflation perception far away from the actual value (which was approximately 5% at the time of the survey) tend to be the least certain in their perceptions. Those with the highest perceptions are also the least certain in their expectations and perceive the noise in the inflation process to be the highest. Those with perceptions that are either very high or very low also tend to have very low perceived persistence on average, and lower Kalman gains.

D.3 Relationship between components of expectation formation: further details

Table 6 breaks down the different elements of the expectation law of motion according to $\hat{\rho}_i$, divided into five categories; those who believe the price level is mean-reverting ($\hat{\rho}_i < 0$), those who believe persistence is zero, those who believe inflation is persistent but stationary ($\hat{\rho}_i \in (0, 1)$), those who believe that inflation follows a unit root ($\hat{\rho}_i = 1$), and those who believe inflation is explosive with positive persistence ($\hat{\rho}_i > 1$). This

Table 6: Means of elements of expectation laws of motion, broken down by persistence type

	$SD_i(\pi_t)$	$SD_i(\pi_{t+1})$	$SD_i(\varepsilon_{t+1})$	χ_i	HTM	Owns Securities
$\hat{\rho}_i < 0$	1.27	1.76	1.68	0.86	0.09	0.66
	(0.26)	(0.11)	(0.17)	(0.02)	(0.03)	(0.05)
$\hat{\rho}_i = 0$	1.80	1.71	1.96	0.79	0.14	0.56
	(0.09)	(0.03)	(0.04)	(0.01)	(0.01)	(0.01)
$\hat{\rho}_i \in (0,1)$	0.85	1.58	1.64	0.80	0.09	0.66
	(0.13)	(0.06)	(0.07)	(0.01)	(0.02)	(0.03)
$\hat{\rho}_i = 1$	1.75	1.56	1.52	0.78	0.10	0.59
	(0.23)	(0.05)	(0.06)	(0.01)	(0.01)	(0.02)
$\hat{\rho}_i > 1$	2.32	2.01	2.05	0.93	0.19	0.46
	(0.35)	(0.09)	(0.14)	(0.01)	(0.02)	(0.03)

Note: Bundesbank-Online-panel-Households, November 2021 wave. Standard errors in parentheses.

follows the classification for stock return beliefs in Dominitz and Manski (2011). Table 6 includes all respondents, with the most extreme 1% of responses for each variable excluded as outliers.[19]

There is a highly non-linear relationship between the standard deviation of the perception and the persistence type. In particular, those who perceive that inflation is persistent but mean reverting are the most confident in their inflation perceptions. Those who believe that $\hat{\rho}_i > 1$ have the lowest confidence in their perceptions. This fits with the notion that those who track inflation most closely are also those who have the best knowledge of its dynamic properties. Those who believe inflation has zero persistence and those who believe it is explosive tend to also believe that the noise in the inflation process is highest.

As noted in Result 4, those who believe that $\hat{\rho}_i > 1$ are more likely to be hand-to-mouth than any of the other persistence types, over twice as likely if one only includes those whose responses are consistent with Kalman filtering. They are also much less likely to own securities.

Finally, note that those who believe that $\hat{\rho}_i > 1$ have the highest χ_i on average. However, this is partly mechanical, since if $|\hat{\rho}_i| > 1$ then that places a lower bound on the values of χ_i that are consistent with steady-state Kalman filtering. Between the groups with $\hat{\rho}_i \in [0,1]$, the average Kalman filter varies little.

Table 7 shows regressions of each component of the expectation laws of motion on $\hat{\rho}_i$, split in two ways. The first panel splits respondents according to which hypothetical scenario they were shown before Question 2, to explore the role of different shock types. That is, each dependent variable is regressed on $\hat{\rho}_i$ interacted with a categorical variable

[19]Note that the first two columns include those for whom $V_i^p > V_f^i$, whose responses are inconsistent with steady-state Kalman filtering.

reflecting which shock scenario the respondent saw.

The second panel splits households into some of the persistence categories outlined above, specifically those who believe the price level is mean-reverting ($\hat{\rho}_i < 0$), those who believe inflation is persistent but stationary ($\hat{\rho}_i \in (0, 1)$), and those who believe inflation is non-stationary with positive persistence ($\hat{\rho}_i \geq 1$). The final panel shows the results of regressing each dependent variable on an indicator equal to 1 if the household does no updating of expectations at all when faced with the hypothetical shock ($\hat{\rho}_i = 0$).

In the first panel, the relationships of $\hat{\rho}_i$ with other elements of expectation laws of motion are generally weaker for respondents who were not shown a specific shock scenario.

While there are some significant differences across shock types, the magnitudes are generally small. For that reason we pool households across shock types for the analysis in Section 4.2.

The differences are much larger, however, across persistence types. Panel 2 shows that within households who believe inflation is persistent and stationary, greater perceived persistence is associated with less uncertainty about current and future inflation, less perceived noise in the inflation process, and a greater implied Kalman gain. This is consistent with models of endogenous information acquisition, as with a more persistent inflation process information about the current rate of inflation is more valuable.

Although there are only weak relationships between $\hat{\rho}_i$ and uncertainty over current and future inflation across the whole sample, the second panel reveals that this is driven by weak relationships among those who believe in inflation processes that are qualitatively different from those seen in the data. Among those who believe that inflation is persistent but stationary, the relationships between $\hat{\rho}_i$ and uncertainty are very strong. Since an AR(1) process estimated on German CPI inflation over the previous 20 years implies a persistence of $\rho = 0.2$, this suggests that the group of households most aware of the time-series properties of inflation behave as predicted by models of rational inattention (e.g. Sims, 2003). However outside of this group, households behave less in line with those predictions.

D.4 Correlations of expectation components with household characteristics: further details

Table 8 column 1 shows the full results of the hurdle model regressing $\hat{\rho}_i$ on household characteristics. The upper panel is as in Table 3 column 4, and the lower panel is the selection step. The remaining columns of Table 8 repeat this exercise, splitting by shock type. The association between $\hat{\rho}_i$ and being hand-to-mouth is weaker for those who saw the demand shock scenario, though this is imprecisely estimated.

Table 9 repeats the analysis in Table 3, with the adjustment that we allow the coefficient on liquid wealth to differ for those with very high liquid wealth (above the 90th percentile in the survey). The positive association between χ_i and liquid wealth found in Table 3 is only present for the high-wealth households.

Table 7: Breakdown of $\hat{\rho}_i$ relationships with other expectation law of motion components by shock type and persistence category.

	(1) $SD_i(\pi_{t+1})$	(2) $SD_i(\pi_t)$	(3) $SD_i(\varepsilon_{t+1})$	(4) χ_i
Panel A: Shock type				
Shock unspecified × $\hat{\rho}_i$	0.0735	0.00338	-0.0474	0.00495
	(0.0745)	(0.0234)	(0.0692)	(0.0118)
Supply × $\hat{\rho}_i$	0.0224	-0.0332*	-0.133*	0.0304***
	(0.0800)	(0.0197)	(0.0764)	(0.00788)
Demand × $\hat{\rho}_i$	0.151	-0.0657***	0.0527	0.0462***
	(0.0993)	(0.0158)	(0.0979)	(0.00853)
Constant	2.068***	0.748***	2.043***	0.790***
	(0.0299)	(0.0114)	(0.0299)	(0.00526)
Panel B: Persistence type				
$\hat{\rho}_i < 0$ × $\hat{\rho}_i$	0.0589	0.0855***	0.229***	-0.0532***
	(0.0840)	(0.0227)	(0.0807)	(0.0107)
$\hat{\rho}_i \in (0, 1)$ × $\hat{\rho}_i$	-0.490***	-0.326***	-0.550***	0.0514*
	(0.157)	(0.0346)	(0.158)	(0.0294)
$\hat{\rho}_i \geq 1$ × $\hat{\rho}_i$	0.100	-0.0464***	-0.0928	0.0439***
	(0.0630)	(0.0145)	(0.0612)	(0.00567)
Constant	2.089***	0.768***	2.087***	0.783***
	(0.0334)	(0.0129)	(0.0334)	(0.00570)
Panel C: Updating indicator				
$\hat{\rho}_i \neq 0$	-0.129**	-0.0995***	-0.314***	0.0293***
	(0.0571)	(0.0189)	(0.0562)	(0.00983)
Constant	2.129***	0.770***	2.128***	0.789***
	(0.0343)	(0.0136)	(0.0343)	(0.00596)
Observations	2383	2383	2383	2383

Note: Bundesbank-Online-panel-Households, November 2021 wave. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Regressions of $\hat{\rho}_i$ on household characteristics, split by shock type.

	(1)	(2)	(3)	(4)
	$\hat{\rho}_i$	$\hat{\rho}_i$	$\hat{\rho}_i$	$\hat{\rho}_i$
$\hat{\rho}_i$				
Hand-to-mouth	0.4051**	0.3869	0.3946	0.0122
	(0.1999)	(0.3446)	(0.2733)	(0.3584)
Liquid wealth	-0.0014***	-0.0006	-0.0017**	-0.0018
	(0.0005)	(0.0005)	(0.0007)	(0.0013)
Illiquid wealth	0.0002	-0.0003	0.0004	0.0007*
	(0.0002)	(0.0002)	(0.0003)	(0.0004)
Other wealth	-0.0010	-0.0002	-0.0015	-0.0032*
	(0.0007)	(0.0009)	(0.0016)	(0.0017)
Debt	-0.0009	-0.0006	-0.0004	-0.0014
	(0.0006)	(0.0008)	(0.0009)	(0.0011)
log(income)	0.0756	0.2334	0.1750	-0.4421
	(0.1777)	(0.2145)	(0.2243)	(0.3435)
selection($\hat{\rho}_i \neq 0$)				
Hand-to-mouth	-0.0939	-0.1104	-0.2293*	0.0921
	(0.0773)	(0.1419)	(0.1346)	(0.1362)
Liquid wealth	0.0000	0.0002	0.0002	-0.0002
	(0.0002)	(0.0003)	(0.0003)	(0.0003)
Illiquid wealth	0.0001	0.0001	0.0001	-0.0000
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Other wealth	0.0003	0.0011	-0.0007	0.0012
	(0.0005)	(0.0009)	(0.0009)	(0.0008)
Debt	0.0000	0.0002	-0.0007	0.0006
	(0.0003)	(0.0004)	(0.0005)	(0.0004)
log(income)	-0.0414	-0.1728	0.0061	0.0070
	(0.0581)	(0.1080)	(0.1059)	(0.1000)
HH Controls	Yes	Yes	Yes	Yes
Hurdle model	Yes	Yes	Yes	Yes
Shock type	All	Unspecified	Supply	Demand
Observations	3194	1053	1057	1084

Note: Bundesbank-Online-panel-Households, November 2021 wave. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Regressions of components of subjective laws of motion on household characteristics, interacting liquid wealth with an indicator for having liquid wealth above the 90th percentile.

	(1) $\log(SD_i(\pi_{t+1}))$	(2) $\log(SD_i(\pi_t))$	(3) $\log(SD_i(\varepsilon_{t+1}))$	(4) $\log(\chi_i)$	(5) $\hat{\rho}_i$
Hand-to-mouth	0.0056	0.1147**	0.1156**	0.0805*	0.4012*
	(0.0372)	(0.0556)	(0.0587)	(0.0421)	(0.2095)
High liquid wealth=0	-0.0004*	-0.0008**	-0.0006*	-0.0001	-0.0015
× Liquid wealth	(0.0002)	(0.0003)	(0.0003)	(0.0003)	(0.0016)
High liquid wealth=1	-0.0000	-0.0002**	0.0002**	0.0002**	-0.0014***
× Liquid wealth	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0005)
Illiquid wealth	0.0000	-0.0000	0.0000	0.0000	0.0002
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0002)
Other wealth	-0.0001	-0.0001	0.0003	-0.0000	-0.0010
	(0.0002)	(0.0002)	(0.0003)	(0.0002)	(0.0007)
Debt	0.0000	0.0001	0.0000	0.0002	-0.0009
	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0006)
log(income)	-0.0748***	-0.1212***	-0.1684***	0.0003	0.0764
	(0.0236)	(0.0349)	(0.0371)	(0.0334)	(0.1786)
HH Controls	Yes	Yes	Yes	Yes	Yes
Hurdle model	No	No	No	No	Yes
Observations	4382	3161	2292	2024	3194

Note: Bundesbank-Online-panel-Households, November 2021 wave. Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

E Further impulse response exercises

Using the survey responses to the different hypothetical scenarios in Question 2, we can further compare the effects of heterogeneous expectation laws of motion for different types of shock. We find somewhat greater amplification and persistence in consumption responses to supply shocks than other types of shock. A comparison of the IRFs between the three cases is shown in Figures 7a and 7b. This result is consistent with the higher average perceived persistence of supply shocks discussed in Section 4.

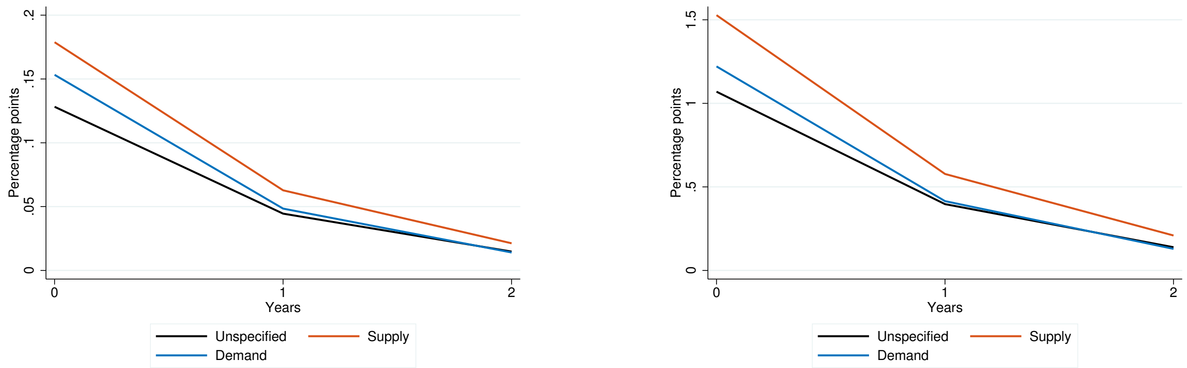

(a) Expectations

(b) Consumption

Figure 7: Implied IRFs of one-period ahead inflation expectations and consumption by shock. Source: Bundesbank-Online-Panel-Households, November 2021 wave.

As in Section 4, we also repeat these exercises with the distributions of subjective models after excluding those whose answers are rounded to a multiple of 0.5. The model with heterogeneity does deliver smaller initial consumption responses in this case, but it is still 4.4× larger than under homogeneity. As such, the result that heterogeneity generates very significant amplification of the transmission of inflation shocks to consumption still holds.