

Inflation Expectations and Choices of Households*

Nathanael Vellekoop[†]

University of Toronto and SAFE

Mirko Wiederholt[‡]

Sciences Po and CEPR

April 2019

Abstract

Do household inflation expectations affect consumption-savings decisions? We link survey data on quantitative inflation expectations to administrative data on income and wealth. We document that households with higher inflation expectations save less. Estimating panel data models with year and household fixed effects, we find that a one percentage point increase in a household's inflation expectation over time is associated with a 250-400 euro reduction in the household's change in net worth per year on average. We also document that households with higher inflation expectations are more likely to acquire a car and acquire higher-value cars. In addition, we provide a quantitative model of household-level inflation expectations.

*We thank our discussants Philippe Andrade, Francesco Bianchi, Jane Ryngaert, and Johannes Wohlfart and seminar and conference participants at AEA 2019, Bundesbank, Cambridge, CEBRA 2018, CESifo, ECB, EEA 2018, Halle, Kiel, NYU, Oxford, SED 2017, Yale, and Zürich for helpful comments. We gratefully acknowledge research support from the Research Center SAFE, funded by the State of Hessen initiative for research LOEWE. We thank CentERdata for making the survey data of the Dutch Household Survey available. Satyajit Dutt and Bopjun Gwak provided excellent research assistantship. Results include our own calculations based on microdata made available by Statistics Netherlands.

[†]nvellekoop@hotmail.com

[‡]mirko.wiederholt@sciencespo.fr

1 Introduction

How do individuals form inflation expectations? The answer to this question is of central importance for policy-makers. Inflation expectations are viewed as a key determinant of inflation (Bernanke, 2007, Yellen, 2015). In addition, academics and policy-makers have argued that, in the current context of low nominal interest rates, deflationary expectations could cause large falls in output, while an increase in inflation expectations could stimulate spending (Krugman, 1998). Unsurprisingly, research on expectation formation has been an important input into policy-making. The rational expectations revolution has had a large impact on institutional design such as central bank independence. Learning models have affected how central banks think about disinflations. More recent work on limited information-processing capacity of agents is likely to have an effect on the design of policy communication. We contribute to this literature by studying how households update inflation expectations at the micro level.

A question that seems equally important is: Do inflation expectations affect choices? The main motivation for studying the formation of inflation expectations is the presumption that they affect choices. There exists a vast literature (encompassing the entire literature on New Keynesian models) that assumes a tight link between inflation expectations and consumption-savings decisions of households, and this tight link is often central for conclusions regarding the propagation of shocks and the effectiveness of policy, but empirically it is still an open question whether such a tight link exists. We therefore view the main contribution of this paper as providing empirical evidence on the relationship between inflation expectations and consumption-savings decisions of households by linking panel survey data on inflation expectations to panel administrative data on income, wealth, and car purchases at the household level.

The survey has the unique feature that the inflation expectations of individual households are elicited over several years. We use the Dutch Household Survey. The survey aims to be representative for the Dutch population. Every year households are asked to forecast prices for the next year. Households participate for several years and inflation changes significantly over periods of several years. Since one can track the answers of individual households over time, one can study how individual households update inflation expectations over time. Most papers studying how individuals update inflation expectations over time study data from a survey of professional forecasters or data from the Michigan Survey of Consumers, but professional forecasters

do not appear prominently in Macroeconomic models and households are surveyed at most twice in the Michigan Survey of Consumers. We therefore start the paper in Section 2 by studying how households update inflation expectations at the micro level. We find that a noisy information model with two modifications (heterogeneous intercepts, and a small probability that households answer 10% instead of reporting their true inflation expectation) matches well the micro data.

We then link the survey data on inflation expectations to administrative data on income and wealth. Dutch households are subject to both a wealth tax and an income tax. Wealth and income data are directly provided to Statistics Netherlands from the tax authorities and banks. Our preferred measure of savings is change in net worth during the year. Net worth at the end of the year is the difference between the value of all assets at the end of the year and the value of all liabilities at the end of the year. The change in net worth during the year is the first difference of net worth at the end of the year.

Our main dependent variable is this measure of savings during the year. Our explanatory variable of interest is the inflation expectation reported during the year. We include year and household fixed effects and control for income, wealth, and a large number of observable household characteristics taken mainly from the administrative data. We find that the coefficient on inflation expectations is negative and statistically significant at the 1% level. A one percentage point increase in a household's inflation expectation over time is associated with a 250 euro reduction in the household's change in net worth per year on average. We also include other household expectations and obtain similar results.

We also link the survey data on inflation expectations to administrative data on car purchases. Statistic Netherlands provides registry data containing ownership of cars in each month for each person. The source of the data is the government agency that registers car ownership. Car registration is mandatory for purposes of liability insurance. For all households in the matched admin-survey data we construct a variable whether the household acquires a car in the 12 months after inflation expectations are elicited. We impute the value of the purchased car from the import value of the car and the depreciation schedule that the tax authorities use. We find that households with higher inflation expectations are more likely to acquire a car and acquire higher-value cars.

Krugman (1998) argued that in a liquidity trap, where nominal interest rates are at zero, the central bank can still stimulate consumer spending by raising inflation expectations. This argument

has had a large impact on the literature on forward guidance (e.g., Eggertsson and Woodford, 2004). This argument has also had a large impact on the literature on government spending multipliers, since government spending multipliers can be large when government spending increases inflation, inflation expectations, and thereby consumer spending (Christiano, Eichenbaum and Rebelo, 2011). To the best of our knowledge, this paper is the first paper that provides a direct estimate of the partial equilibrium effect of an increase in inflation expectations on the level of consumer spending. We find that the effect has the predicted sign—an increase in expected inflation stimulates consumer spending—and is large—a one percentage point increase in a household’s inflation expectation over time is associated with a 250 euro reduction in annual savings on average.

It is important to emphasize that we provide an estimate of the partial equilibrium effect of an increase in inflation expectations on consumer spending, since we focus on idiosyncratic variation in inflation expectations, while aggregate variation in inflation expectations is captured through time fixed effects. General equilibrium effects can be computed with a general equilibrium model and will depend on the specificities of the country, for example, through the degree of price stickiness and the conduct of monetary policy. The vast literature on New Keynesian models has identified key determinants of sign and size of these general equilibrium effects, taking as given that there exists a link between inflation expectations and consumption-savings decisions at the household level. We provide evidence in favor of such a link. Given our results in Section 2 of this paper, it also seems important that a general equilibrium model features realistic assumptions about the expectation formation of households.

There exist only a few papers that study the relationship between inflation expectations and choices of households using microdata. Most papers in this literature examine the relationship between quantitative inflation expectations and answers to *qualitative* questions on “readiness to spend” (Bachmann et al., 2015, D’Acunto et al., 2018, Andrade et al., 2017, Arioli et al., 2017). These papers use the Michigan Survey of Consumers (MSC) or similar surveys for other countries. Another group of papers exploit recent innovations in the Federal Reserve Bank of New York’s Survey of Consumer Expectations (SCE): Armantier et al. (2015) find that inflation expectations co-move in a meaningful way with investment choices in a financially incentivized field experiment, while Crump et al. (2018) estimate the elasticity of intertemporal substitution, exploiting the fact that the SCE elicits quantitative measures of both inflation and spending growth expectations.

Malmendier and Nagel (2016) investigate the relationship between their model-implied inflation expectations and financial decisions reported in the Survey of Consumer Finances (SCF) at the cohort level. We take a different approach to the existing papers in this literature by linking panel survey data on inflation expectations to panel administrative data on income, wealth, and car purchases at the household level. To the best of our knowledge, our paper is the first paper that links survey data on expectations to administrative data on income and wealth.^{1,2}

The rest of the paper is organized as follows. Section 2 presents the survey data on inflation expectations and our results on how individual households update inflation expectations over time. Section 3 introduces the administrative data on income, wealth, and car purchases and studies the empirical relationship between inflation expectations and choices. Section 4 concludes.

2 Inflation Expectations of Households

In this section, we investigate how households update inflation expectations over time. The main goal of this section is to present the survey data on inflation expectations and to make the point that a standard model of the dynamics of the *average* inflation expectation also matches fairly well the dynamics of *individual-level* inflation expectations.

2.1 Data

The inflation expectations microdata is from the DNB Household Survey, conducted annually since 1993 and administered by CentERdata at Tilburg University. The survey aims to be representative for the Dutch population. Households participate for several years.³ Since one can track individual households over time, one can study how individual agents update inflation expectations over time.

¹By contrast, D’Acunto et al. (2019) link administrative data on cognitive-ability (IQ) tests with survey data on inflation expectations, planned spending, planned saving, and planned borrowing. In particular, respondents are asked about their inflation expectations, whether now is the right moment to spend on durables, and whether they plan to save or borrow.

²Another distinguishing feature of our work is that we include household fixed effects, which is rare in the literature on inflation expectations and choices. A notable exception is Burke and Ozdagli (2013) who study data on expectations and choices from a series of survey modules appended to the RAND’s American Life Panel over a four-year period.

³For comparison, households are surveyed at most twice in the Michigan Survey of Consumers. Households participate in the Federal Reserve Bank of New York’s Survey of Consumer Expectations for at most 12 months.

The data are collected through the Internetpanel of CentERdata. Households without a computer and/or access to the Internet are provided a basic computer and an Internet connection. The DNB Household Survey consists of six questionnaires. The questionnaire “Health and Income” includes several questions about inflation expectations.

Beginning with the 2008 wave, the main quantitative question on inflation expectations is:

“What is the most likely (consumer) prices increase over the next twelve months, do you think?”

Possible answers are:

1%, 2%, 3%, . . . , 10%

Respondents are then asked four questions regarding their subjective CDF.

In the years 1993-2002, households were only asked for a point prediction. The questions read:

“Do you expect prices in general to rise, to remain the same, or to go down, in the next 12 months?”

“If the answer is rise: By what percentage do you expect prices in general to rise in the next 12 months?”

In the years 2003-2007, households were only asked for their subjective CDF:

“We now would like to learn what you expect will happen to the prices in the next twelve months. What will be the minimum percentage prices could increase over the next twelve months, do you think? If you think prices will decrease, you can fill in a negative percentage by using a minus in front of the number.”

“What is the maximum percentage prices will increase over the next twelve months, do you think?”

Calling the answers MIN and MAX, the respondents were asked 4 questions, with $i \in \{2, 4, 6, 8\}$:

“How likely do you think that it is that the increase in prices in the next twelve months will be less than $\frac{i(MAX-MIN)}{10} + MIN$?”

For the years 2003-2007, where households were only asked for their subjective CDF, we estimate the mean of the subjective CDF using a piece-wise linear interpolation over the probability density function.

Figure 1 shows the cross-sectional distribution of point predictions made in the year 2012. Recall that households were asked: “What is the most likely (consumer) prices increase over the next twelve months, do you think?” Possible answers were: 1%, 2%,..., 10%. As one can see, there is large cross-sectional heterogeneity in the answers. Some households answered 1%, while other households picked 10%. The large majority of respondents made a very good forecast that year: two thirds of households answered 2% or 3% in 2012, and CPI inflation turned out to be 2.5% in 2013.

Figure 2 illustrates the dynamics of the cross-sectional distribution of inflation expectations. The cross-sectional distribution of inflation expectations for year t is described by the 10th percentile (dots), the 90th percentile (small dashes), and the mean (large dashes). To allow comparison between expectations and realizations, we also plot the time series for realized CPI inflation (solid line). The four numbers reported for year t refer to the distribution of forecasts made in year $t-1$ for year t and the realization in year t . The vertical lines mark changes in the survey questions. Cross-sectional heterogeneity in inflation expectations is large in all years and the cross-sectional mean of inflation expectations has properties which are familiar from the literature on predictability of average forecast errors (Coibion and Gorodnichenko, 2012 and 2015). In the case of a difference between realized inflation in year t and the average forecast made in year $t-1$ for year t , the average inflation expectation moves in the direction of realized inflation but the average forecast error is highly persistent. A persistent average forecast error can arise because: (i) inflation expectations and realized inflation move in opposite directions (see the average inflation expectation for year 2002 and realized inflation in year 2002, where euro coins and banknotes were introduced), and (ii) inflation expectations adjust sluggishly (see the period since November 2013, where the main policy rate of the European Central Bank has been 25 basis points or less). Finally, the changes in the survey questions did not coincide with unusual changes in the cross-sectional mean of reported inflation expectations, but they may have coincided with small changes in the cross-sectional heterogeneity in reported inflation expectations.

The panel component of the survey data allows us to track the answers of individual households

over time. Several thousand households participated in the survey. It is difficult to visualize several thousand paths for the reported inflation expectation. We therefore present transition matrices.

In Table 1, we study the answers of all households with an observation in the year after the first observation. The entries are conditional probabilities. The first row contains the relative frequency of answers in year two given that the answer in year one was 1%, the second row contains the relative frequency of answers in year two given that the answer in year one was 2%, and so on. The diagonal entries are 0.47, 0.42, 0.33, 0.36, and 0.30. On average more than one third of households gave the same answer in year two as in year one. The fraction is higher for initially low answers and lower for initially high answers.

Table 2 repeats the exercise for all households who provided an answer in the two years after providing the first answer. The first panel reports transition probabilities comparing answers in years one and two. The second panel shows transition probabilities comparing answers in years two and three. The third panel reports transition probabilities comparing *directly* answers in year one and year three. The two one-year transition matrices in Table 2 are similar to the one-year transition matrix in Table 1, suggesting that there is nothing special about the first year of being in the survey. The last transition matrix comparing directly answers in years one and three reveals a striking feature of the expectations data. The (1,1) entry of panel three (“1 to 3”) is much larger than the product of row one of panel one (“1 to 2”) and column one of panel two (“2 to 3”). This means that answers do not follow a Markov process with common transition probabilities. A household has a higher probability of going from an answer of 2% in year *two* to an answer of 1% in year *three* if the household already said 1% in year *one*. The same observation applies to the other diagonal entries of the last panel. Households tend to return to answers they have given in the past.

Table 3 confirms this finding. In Table 3, we repeat the exercise for all households with three observations in the three years after the first observation. The different panels in the table are the three one-year transition matrices and the transition matrix comparing directly the answers in years one and four. The diagonal entries of the last transition matrix (“1 to 4”) are again much larger than the probabilities implied by a Markov process with common transition probabilities and the one-year transition matrices reported in the first three panels. Households tend to return to answers they have given in the past. To identify this feature of the data, one needs at least three

observations per household.

The panel data also reveal that extreme answers come to a large extent from households who report a fairly normal inflation expectation before or after. The first row of any one-year transition matrix shows that households who say 1% in a given year have a 3-4 percent probability of saying 6% or more in the next year. Moreover, the last row of any one-year transition matrix shows that households who say 6% or more in a given year have a high probability of providing a non-extreme answer in the next year.

In sum, there is large cross-sectional heterogeneity in reported inflation expectations, the average forecast error is highly persistent, households tend to return to answers they have given in the past, and extreme answers come to a large extent from households who provide fairly normal answers before and after.

This completes the description of the survey data on inflation expectations. In the following two subsections, we ask whether a noisy information model of expectation formation can simultaneously match the average inflation expectation and the transition matrices for individual inflation expectations.

2.2 Model

In this subsection, we propose a model of household-level inflation expectations.⁴ In the model, households receive noisy signals on inflation. The noise has an aggregate component and an idiosyncratic component. The aggregate component of noise is interpreted as noise in official inflation statistics or noise in media reports on inflation. The idiosyncratic component of noise is interpreted as coming from randomness in perception due to limited attention.

Noisy signal models of belief formation, where the noise is interpreted as coming from limited attention (Sims, 2003, Woodford, 2003, Maćkowiak and Wiederholt, 2009), have gained popularity in the empirical literature on expectation formation (Patton and Timmermann, 2010, Coibion and Gorodnichenko, 2012 and 2015), since the noise alone generates many features of survey data on inflation expectations: (i) sluggish adjustment of the average expectation due to Bayesian updating in the presence of noisy signals, (ii) cross-sectional dispersion and idiosyncratic transitions in

⁴Malmendier and Nagel (2016) propose a model of cohort-level inflation expectations, but across cohort variation in inflation expectations is small in our sample.

expectations due to the idiosyncratic component of noise, and (iii) unwarranted movements in the average expectation due to the aggregate component of noise (e.g., the sharpest increase in the average inflation expectation in our sample took place in forecasts in 2001 for 2002 at a time where the news media warned about the inflationary effects of the introduction of euro coins and notes).

We modify such a standard noisy signal model of belief formation by introducing an assumption that generates individual-specific intercepts in the equation for household-level inflation expectations, because we documented that households tend to return to answers given in the past (see Section 2.1). While the presence of those intercepts will improve the model's ability to match household-level expectations data, the precise origin of those intercepts will not be important and we will discuss two alternative ways of generating these intercepts at the end of this subsection.

Households' perceived law of motion for inflation is

$$\pi_t = (1 - \rho) c + \rho \pi_{t-1} + u_t, \quad (1)$$

where π_t is the inflation rate in year t , $\rho \in (-1, 1]$ is the autocorrelation coefficient, $c \in \mathbb{R}$ is a constant, and $u_t \sim i.i.d.N(0, \sigma_u^2)$ is the inflation innovation in year t . In every period, each household i receives a noisy signal on inflation. According to the household, the signal is generated as follows

$$s_{it} = \pi_t + \varepsilon_{it}, \quad (2)$$

where $\varepsilon_{it} \sim i.i.d.N(\mu_i, \sigma_\varepsilon^2)$ is the noise in the signal. The noise has an aggregate component and an idiosyncratic component, $\varepsilon_{it} = \bar{\varepsilon}_t + \hat{\varepsilon}_{it}$. The aggregate component of noise, $\bar{\varepsilon}_t$, is interpreted as coming from noisy inflation statistics or noisy media reports on inflation. The idiosyncratic component of noise, $\hat{\varepsilon}_{it}$, is interpreted as coming from randomness in perception due to limited attention. The new assumption is that household i 's subjective mean of the aggregate component of noise, denoted $\mu_i \in \mathbb{R}$, may be non-zero, which captures the idea that the household may believe that official inflation statistics or media reports on inflation are biased. Households remember all signals received in the past and use the steady-state Kalman filter to compute conditional expectations of future inflation.

The standard Kalman filter equations imply that the nowcast for inflation is given by

$$E[\pi_t | \mathcal{I}_{i,t}] = E[\pi_t | \mathcal{I}_{i,t-1}] + K (s_{it} - \mu_i - E[\pi_t | \mathcal{I}_{i,t-1}]).$$

The nowcast for inflation of household i , $E[\pi_t|\mathcal{I}_{i,t}]$, is a linear combination of the household's prior mean, $E[\pi_t|\mathcal{I}_{i,t-1}]$, and the product of the Kalman gain K and the difference between the signal realization and the expected signal realization – after the household has deducted the perceived bias μ_i from the signal to transform the signal into an unbiased signal on current inflation. The perceived law of motion for inflation implies that the forecast for inflation is

$$E[\pi_{t+1}|\mathcal{I}_{i,t}] = (1 - \rho)c + \rho E[\pi_t|\mathcal{I}_{i,t}].$$

Combining the last two equations yields

$$E[\pi_{t+1}|\mathcal{I}_{i,t}] = (1 - \rho)c - \rho K\mu_i + \rho(1 - K)E[\pi_t|\mathcal{I}_{i,t-1}] + \rho K s_{it}.$$

If the signal s_{it} is indeed generated according to equation (2), we arrive at

$$E[\pi_{t+1}|\mathcal{I}_{i,t}] = (1 - \rho)c - \rho K\mu_i + \rho(1 - K)E[\pi_t|\mathcal{I}_{i,t-1}] + \rho K\pi_t + \rho K\varepsilon_{it}. \quad (3)$$

If the household is right about the bias in official inflation statistics, then the term $-\rho K\mu_i$ simply corrects the non-zero mean of $\rho K\varepsilon_{it}$. However, if the household's perceived bias of official inflation statistics or media reports differs from the actual bias of official inflation statistics or media reports, the household's inflation expectation is shifted up or down.

The last equation can be written more concisely as

$$\pi_{t+1|t,i} = \beta_i + \beta_1\pi_{t|t-1,i} + \beta_2\pi_t + \bar{\nu}_t + \hat{\nu}_{it}, \quad (4)$$

where $\pi_{t+1|t,i} \equiv E[\pi_{t+1}|\mathcal{I}_{i,t}]$ is the current period forecast, $\pi_{t|t-1,i} \equiv E[\pi_t|\mathcal{I}_{i,t-1}]$ is the previous period forecast, and π_t is the inflation realization. The noise term in the forecast has an aggregate component, $\bar{\nu}_t \equiv \rho K\bar{\varepsilon}_t$, and an idiosyncratic component, $\hat{\nu}_{it} \equiv \rho K\hat{\varepsilon}_{it}$, because the noise in the household's signal has these two components. Here $\beta_i = (1 - \rho)c - \rho K\mu_i$, $\beta_1 = \rho(1 - K)$, and $\beta_2 = \rho K$. Note that the ratio $\beta_2/\beta_1 = K/(1 - K)$ is a function only of the Kalman gain and the sum $\beta_1 + \beta_2 = \rho$ equals the household's perceived persistence. The last equation shows how noisy signal models generate three features of survey data on expectations: (i) underreaction of the average inflation expectation to inflation innovations since $\beta_2 < \rho$, (ii) cross-sectional dispersion and idiosyncratic transitions in inflation expectations due to the idiosyncratic component of noise, $\hat{\nu}_{it}$, and (iii) unwarranted movements in the average inflation expectation due to the aggregate

component of noise, $\bar{\nu}_t$. The new feature of the model is that heterogeneity in the perceived biases μ_i generates individual-specific intercepts in the equation for household-level inflation expectations.⁵

Finally, from the last equation for a household's inflation expectation one can derive an equation for the average inflation expectation. Averaging across i on both sides of equation (4) and assuming that the idiosyncratic component of noise washes out in the aggregate yields the following equation for the average inflation expectation

$$\bar{\pi}_{t+1|t} = \bar{\beta} + \beta_1 \bar{\pi}_{t|t-1} + \beta_2 \pi_t + \bar{\nu}_t, \quad (5)$$

where $\bar{\pi}_{t+1|t}$ denotes the average inflation expectation in year t and $\bar{\beta}$ denotes the average of β_i .

The individual-specific intercepts in household-level inflation expectations can be derived differently. For example, households may shrink a data-based forecast towards some other view. Patton and Timmermann (2010) propose a model in which professional forecasters shrink a data-based forecast towards some other individual-specific view. According to their model, the forecast of agent i in year t is a weighted average of the conditional expectation, $E[\pi_{t+1}|\mathcal{I}_{i,t}]$, and some other individual-specific view, ξ_i ,

$$\pi_{t+1|t,i} = \omega \xi_i + (1 - \omega) E[\pi_{t+1}|\mathcal{I}_{i,t}].$$

If the conditional expectation is given by equation (3), then the forecast is given by equation (4).⁶

Alternatively, households may misunderstand the survey question on inflation expectations and submit forecasts for *household-level* inflation rates. Kaplan and Schulhofer-Wohl (2017) use scanner data to estimate inflation rates at the household level. One of their main findings is that the household-specific component of household-level inflation is almost uncorrelated over time.⁷

⁵In equation (1) we assume that all households have the same perceived law of motion for inflation. One can generalize the model by assuming that households have individual-specific constants, c_i . The only change in equation (3) is that c is replaced by c_i , implying that β_i in equation (4) becomes $\beta_i = (1 - \rho)c_i - \rho K \mu_i$. Furthermore, if $c_i = c - \mu_i$, the time series average of $E[\pi_{t+1}|\mathcal{I}_{i,t}]$ implied by equation (3) equals the unconditional expectation of inflation implied by equation (1). That is, the perceived law of motion for inflation and the average inflation expectation are consistent. Moreover, the agent can no longer detect mistakes in c_i or μ_i based on observations of the signal alone.

⁶In this case, and assuming $\mu_i = 0$ for ease of exposition, we have $\beta_i = \omega[1 - \rho(1 - K)]\xi_i + (1 - \omega)(1 - \rho)c$, $\beta_1 = \rho(1 - K)$, $\beta_2 = (1 - \omega)\rho K$, and $\nu_{it} = (1 - \omega)\rho K \varepsilon_{it}$.

⁷They find that the cross-sectional correlation between a household's inflation rate in year t and its inflation rate in year $t+1$ is approximately -0.1.

In this case, forecasting household-level inflation is essentially equivalent to forecasting aggregate inflation. But households may not be aware of the fact that the household-specific component of inflation is almost uncorrelated over time. For example, if the *perceived* law of motion for household-level inflation is aggregate inflation plus a constant ($\pi_{it} = \pi_t + \delta_i$), the perceived law of motion for aggregate inflation is given by equation (1), and households pay limited attention to current household-level inflation to forecast future household-level inflation ($s_{it} = \pi_{it} + \varepsilon_{it}$), then the conditional expectation of future household-level inflation is given by

$$E[\pi_{i,t+1}|\mathcal{I}_{i,t}] = (1 - \rho)(c + \delta_i) + \rho(1 - K)E[\pi_{it}|\mathcal{I}_{i,t-1}] + \rho K s_{it}.$$

Finally, if the *actual* law of motion for household-level inflation has the form $\pi_{it} = \pi_t + \zeta_{it}$ where ζ_{it} is uncorrelated over time (i.e., the household-specific component of household-level inflation is uncorrelated over time, as found by Kaplan and Schulhofer-Wohl, 2017), the forecast of household i in year t is again given by equation (4) with $\beta_i = (1 - \rho)(c + \delta_i)$, $\bar{\nu}_t = 0$, and $\hat{\nu}_{it} = \rho K(\zeta_{it} + \varepsilon_{it})$. In the following, the presence – not the origin – of the individual-specific intercepts will matter.

2.3 Comparison of model and data

The noisy signal model presented in the previous subsection (without the individual-specific intercepts) has been frequently used to match *average* inflation expectations in the literature building on Coibion and Gorodnichenko (2012, 2015). We now ask whether this model also matches the empirical transition matrices for individual-level inflation expectations reported in Section 2.1.

To make this point, we proceed in three steps. We assume that β_i has a log-normal distribution. We estimate equation (5) for the average inflation expectation with the time series for inflation and for the average inflation expectation plotted in Figure 2. This yields estimates of $\bar{\beta}$, β_1 , β_2 , and $\sigma_{\bar{\nu}}^2$. If we were only interested in the average inflation expectation, this would complete the exercise. Furthermore, we set the three parameters that appear in equation (4) but not in equation (5) as follows. We set the two parameters of the log-normal distribution for β_i and the variance $\sigma_{\bar{\nu}}^2$ so as to match the estimate of $\bar{\beta}$, the cross-sectional standard deviation of inflation expectations reported in Figure 1, and the (2,2) entry of panel one of Table 3. The rationale is simple. Both a positive variance of β_i and a positive variance of $\hat{\nu}_{it}$ generate cross-sectional dispersion in beliefs, but the intercepts β_i generate permanent differences in views, while the noise terms $\hat{\nu}_{it}$ generate temporary

differences in views. Therefore, the parameters are picked so as to match the overall dispersion in beliefs and the stability of views at the micro level, as summarized by the (2,2) entry of panel one of Table 3.

Table 4 shows the results for estimation of equation (5). The point estimates of $\bar{\beta}$, β_1 , β_2 , and σ_v^2 are 0.40, 0.61, 0.44, and 0.2860 (first column of Table 4).⁸ The estimates of β_1 and β_2 imply the following structural parameters. The estimates of β_1 and β_2 approximately sum to one; and dividing the estimate of β_1 by the estimate of β_2 yields a value around 1.5. According to the model this means households' perceived law of motion for inflation is a random walk ($\beta_1 + \beta_2 = \rho$) and the Kalman gain is about 0.4 ($\beta_1/\beta_2 = (1 - K)/K$). Finally, setting the remaining parameters so as to match the overall dispersion in beliefs and the stability of views at the micro level yields a cross-sectional standard deviation of β_i of 0.6412 and a variance of \hat{v}_{it} of 0.2059. This completes our choice of parameters.

With equation (4) and an actual law of motion for inflation we compute model-implied transition matrices for household-level inflation expectations.⁹ The actual law of motion for inflation is obtained by fitting a first-order autoregressive process to the official annual inflation series for the Netherlands. The estimates are reported in Table 5. The constant equals 0.72, the autocorrelation coefficient equals 0.59, and the variance of the innovation equals 0.7873. We simulate data with equation (4) and the actual law of motion for inflation. Each simulation is based on 2,000 years and 10,000 individuals.¹⁰ The following transition matrices are averages over 50 simulations.

Tables 6-8 show the model-implied transition matrices for individual-level inflation expectations. The (2,2) entry of Table 6 is the targeted moment in these transition matrices. The other moments reported in Tables 6-8 are not targeted. The model-implied one-year transition probabilities are close to the empirical one-year transition probabilities for moderate inflation expectations (1-3%). In particular, the upper-left (3x3) submatrix in Table 6 is similar to the upper-left (3x3) submatrix in panel one of Table 3. The model does much less well matching probabilities of transitions that

⁸The first column uses inflation expectations as reported by households, where the survey question is asked in three different ways as reported in Section 2.1. The second column adjusts inflation expectations to the same format of the most recent question. Results are similar. In both columns, the inflation data are the official annual inflation data published by Statistics Netherlands.

⁹Note that inflation is a variable that enters equation (4).

¹⁰In the simulation, we assume that households with inflation expectations below 1.5% say 1%, households with inflation expectations in the interval [1.5%, 2.5%) say 2%, and so on.

involve high inflation expectations. First, in the model there are essentially no direct transitions from low inflation expectations of 1-2% to very high inflation expectations of 6-10% (see the (1,5) and (2,5) entry of Table 6) and no direct transitions from very high inflation expectations to low inflation expectations (see the (5,1) and (5,2) entry of Table 6); as one would expect from a model with Bayesian agents tracking an inflation process with a small variance of the innovation. By contrast, in the data such transitions are quite frequent (see the (1,5) and (2,5) entry as well as the (5,1) and (5,2) entry of panel one of Table 3). Second, in the model households remain in the brackets 4-5% and 6-10% with a high probability because these are wide brackets (see the last two diagonal entries of Table 6). By contrast, in the data households remain in these brackets with a much lower probability (see the last two diagonal entries of the first panel of Table 3). In sum, there seems to be something special about the high answers. Households transit to them with a fairly high probability but do not stay at these high answers. Moreover, most of the transitions from moderate inflation expectations of 1-3% to very high inflation expectations of 6-10% are direct transitions to the answer category 10%. For this reason, we introduce a specific form of measurement error in the model. We assume that with a probability of 3% a household temporarily answers 10% instead of reporting the true inflation expectation.

Tables 9-11 show the transition matrices for the extended model. The values of $\bar{\beta}$, β_1 , β_2 , and σ_v^2 have remained unchanged. The new feature is that with a probability of 3% a household answers 10% instead of reporting the true inflation expectation. The cross-sectional standard deviation of β_i and the variance σ_v^2 are again chosen so as to match the cross-sectional standard deviation of inflation expectations reported in Figure 1 and the (2,2) entry of panel one of Table 3. The new values are 0.4684 and 0.1373, respectively. The model now matches the last column and the last row of the empirical transition matrices much better.

To conclude, we find that a noisy information model with two modifications (heterogeneous intercepts, and a small probability that households answer 10% instead of reporting their true inflation expectation) that is estimated with the average inflation expectation also matches fairly well transition matrices for individual-level inflation expectations. Based on the model, we interpret the temporary transitions from moderate inflation expectations to 10% inflation expectations and back as a specific form of measurement error, and we exclude these answers from the subsequent analysis of inflation expectations and choices of households.

3 Choices of households

In this section, we report novel results on the relationship between inflation expectations and financial decisions of households.

3.1 Data

We exploit the fact that the same survey with household inflation expectations (the DNB Household Survey) can be linked to administrative data containing household wealth. Statistics Netherlands makes the administrative data available, and provides a working environment where we can merge the survey with the administrative data at the household level.¹¹ This allows us, for example, to link the inflation expectation reported by household j in year t to the wealth held in checking and savings accounts by household j in year t , as reported by banks. Households are asked in 2011 to 2014 whether they give consent to be matched. Of the households appearing in our DHS sample 88% agreed to be matched.¹² Even though consent is asked in 2011, we can uncover observations going back to the year 2008. The last year in our matched sample is 2016. The following two paragraphs provide more information on the income and wealth measures based on administrative data.

We use disposable household income, which is the sum of labor income, business income, and interest income, plus transfers and alimony, minus taxes and mandatory health insurance premiums for all members in the household. Statistics Netherlands imputes a rental value for home-owners' use of the own home, which is part of the income measure. Not measured (or imputed) are income transfers between households, income transfers from abroad, black market income, and alimony paid for children.

Wealth is measured from several administrative sources, directly provided to Statistics Netherlands from the tax authorities and banks. The Netherlands has a wealth tax, which is calculated as a fixed rate on the average holdings of cash, checking and savings accounts, stocks and bonds, real estate not being the primary residence, minus debts (including study loans, but excluding mortgages for the primary residence). Since there is a threshold of 20,000 euro of wealth for the

¹¹Statistics Netherlands checks all output to avoid disclosure of individual households.

¹²On average there are no differences in inflation expectations between households who consented and those who did not.

wealth tax (double the amount for couples), from tax records alone only higher wealth levels would be observed. For households not reporting wealth, Statistics Netherlands imputes wealth holdings based on dividend and interest income. The Netherlands also has a property tax. Local governments provide home-owners with a market-based estimate of the value of their home, typically every year. This estimate of the house value is the tax base for several local taxes and the income tax. Furthermore, for all households, banks report wealth held in checking and savings accounts.

Our preferred measure of savings is change in net worth during the year. Net worth at the end of a year is the difference between the value of all assets at the end of the year and the value of all liabilities at the end of the year. Our main dependent variable is the first difference of net worth. Assets include the wealth held in checking and savings accounts, the value of stocks and bonds (measured as the market value at the beginning of January of the following year), the value of residential property, stock ownership in substantial holdings, and business equity. On the liability side the mortgage value of the own home and the sum of other loans (including study loans) are reported, but not consumer loans or credit line facilities. We use an unbalanced household panel for the nine years 2008-2016.

We estimate empirical models of the following form

$$Y_{i,t} = \beta_0 + \beta_1 \pi_{t|t-1,i} + \mathbf{x}'_{i,t} \delta + \lambda_t + \nu_{i,t}. \quad (6)$$

The dependent variable $Y_{i,t}$ is the change in net worth of household i in year t . We will also consider the level of net worth at the end of year t . Furthermore, we will consider sub-categories of net worth, by looking at total assets, total liabilities, and total deposits (defined as wealth held in checking and savings accounts). The main variable of interest is $\pi_{t|t-1,i}$. In terms of timing, the dependent variable, taken from the administrative data, is measured at the end of year t , while the expectation of inflation over the next 12 months, taken from the survey, is measured at the beginning of year t . Therefore, we denote the inflation expectation by $\pi_{t|t-1,i}$ in the last equation. We control for relevant background characteristics $x_{i,t}$ and year fixed effects λ_t . Standard errors ($\nu_{i,t}$) are clustered at the level of the household. In some specifications we exploit the panel dimension and estimate models with household fixed effects in order to capture unobserved heterogeneity. Since first differences of wealth can result in large swings, we trim the upper and lower ten percent of the distribution in a year. For all regressions we construct for each household the coefficient of variation of net worth (the standard deviation over the mean), and trim the upper and lower one percent. Regressions with

household fixed effects are sensitive to large changes of an outcome variable within a household, and this procedure takes care of that. Trimming of the distribution is done separately for each regression, therefore the number of observations may differ between outcomes. For the years of the matched admin-survey data (2008-2016), respondents to the inflation expectation question were asked to choose one number out of [1%, 2%, ..., 10%], see Section 2.2. We drop all observations with inflation expectations of 10%, for the reasons explained in Section 2.3.

There is no consumption data available in the survey, but we link the matched admin-survey data to administrative data on car ownership. Statistics Netherlands provides registry data containing ownership of cars in each month for each person, for the years 2010-2016. The source of the data is the government agency that registers car ownership. Car registration is mandatory for purposes of liability insurance. In the data we can distinguish between new and used (second-hand) cars. For many cars the import value of a car is registered (almost all cars in The Netherlands are imported), which is the value without dealer modifications or discounts. We impute a current value by using the depreciation schedule that the tax authorities use, which is close to exponential. For all households in the matched admin-survey data we construct a variable whether the household acquires a car in the 12 months after inflation expectations are elicited.

Table 12 shows the summary statistics for the regression sample. Expected inflation is measured in brackets [1%, 2%, ..., 9%],¹³ see Section 2.2 for the wording of the question. Net worth averages around 106,000 euro, with on average 171,000 euro in assets and 65,000 euro in liabilities. Home-ownership rates in the sample are slightly higher than in the population (73% in the sample and around 67% in the population). Households hold on average 27,000 euro as deposits (defined as wealth in checking and savings accounts). After-tax household income is around 23,500 euro. Older households are overrepresented in our sample, which has the advantage that household portfolios are more mature. Overall we have 10,808 household-year observations for 2,131 unique households, with an average panel dimension of close to five years.

¹³We drop all observations with a 10% answer. Dropping all households with at least one 10% answer yields similar results.

3.2 Results on Net Worth

In this subsection, we report results for the regressions without household fixed effects. In Section 1.4, we include household fixed effects. Results are similar.

Table 13 reports the relationship between inflation expectations¹⁴ during the year and the level of net worth at the end of the year. Households with higher inflation expectations have lower net worth. This relationship is statistically significant at 1%, and economically meaningful. A one percentage point increase in inflation expectations is associated with a 4,500 euro decrease in net worth. Net worth consists of assets minus liabilities. In columns (2) and (3) we find that higher inflation expectations are associated with both lower levels of assets and lower levels of liabilities.¹⁵ In the last column we find the same relationship with deposits, though only significant at 10%. An increase of one percentage point in inflation expectations is associated with a decrease of 824 euro in deposits. All regressions control for the disposable household income, education, a number of other household characteristics, regional characteristics and a set of year fixed effects. In sum, households with higher inflation expectations have lower net worth, assets, liabilities, and deposits.

The average relationship between inflation expectations and net worth is not necessarily linear. We run the same regressions reported in Table 13, but with inflation expectations in 9 categories, 1% being the baseline. In Figure 3 we plot the regression coefficients and the 95% confidence intervals. Looking at the top-left panel, we find that a linear relationship is a good approximation for the relationship between inflation expectations and net worth. Confidence intervals at certain brackets are wide, especially at 9%, this is due to the low numbers of observations in that bracket.

In Table 14 we study the relationship between inflation expectations and our preferred measure of savings—the change in net worth during the year. In addition to the same controls as in Table 13, we add net worth, transformed using the inverse hyperbolic sine.¹⁶ Our main finding in this subsection is that households with higher inflation expectations save less. A one percentage point increase in expected inflation is associated with a 300 euro decrease in the change in net worth on

¹⁴Whenever possible we use the inflation expectation of the household head. In the case of a missing value, we use the inflation expectation of the spouse. The variables “female”, “retired”, “college education”, and “age” all refer to the person providing the inflation expectations. “Children in the house” is a dummy variable.

¹⁵Note that the liabilities observed in the administrative data are mainly mortgages and student loans, but not consumer credit.

¹⁶The inverse hyperbolic sine transformation is similar to the log, but allows negative and zero values.

average. For comparison, the unconditional mean of the change in net worth is -380 euro in our sample. The coefficient on inflation expectations is also negative and statistically significant in the regression for change in assets and in the regression for change in deposits (a component of assets). A one percentage point increase in inflation expectations is associated with a 220 euro decrease in the change in assets and a 70 euro decrease in the change in deposits on average.

In Figure 4 we repeat the regressions but with inflation expectations in 9 brackets. Here evidence for a linear relationship is more mixed. For low values of inflation expectations, the linear approximation works quite well, though not for expectations larger than 6%. Given the wide 95% confidence intervals for inflation expectations of 7-9%, zero cannot be excluded.

Taken together, we find that households with higher inflation expectations have smaller net worth and a smaller change in net worth (i.e., they save less).

3.3 Results on Car Acquisitions

To address the relationship between expectations and some direct measure of spending, we turn to the data on car acquisitions. In Table 15 we report marginal effects after Probit regressions in columns (1)-(3), and marginal effects after a Tobit regression in column (4).¹⁷ All regressions control for net worth, year fixed effects and the number of months since the last car acquisition by the household.¹⁸ Overall, we find a positive and sizable effect of inflation expectations on the probability to acquire a car. In column (1), a one percentage point increase in inflation expectations increases the probability of acquiring a car in the next 12 months with 4.3 percentage points. This is a sizable effect compared to an average of 10 percent of households acquiring a car. When we distinguish between used and new cars, we find positive relationships for both, though only statistically significant for used cars (at 1%). In the last column we estimate a relationship between inflation expectations and the log of the value of the car, as imputed using a depreciation schedule. Missing car values are replaced with zero. We find that households with higher inflation expectations acquire higher-value cars. A one percentage point increase in inflation expectations is associated with a 4% increase in car value upon acquisition. Taken together, this shows that households with

¹⁷Results are robust when estimated with linear probability models. The Tobit model is estimated to account for the zero values for those households who do not acquire a car, or for acquisitions for which we do not observe a car value.

¹⁸For first-time car owners a value of zero is used.

higher inflation expectations are more likely to acquire a car, as well as acquire higher-value cars.

3.4 Results on Unobserved Heterogeneity and Including Other Expectations

Turning back to our main result—the negative relationship between inflation expectations and change in net worth—there are two potential concerns. The first concern is unobserved heterogeneity. The second concern is that inflation expectations are confounded with other household expectations. The data allow us to explore both.

First we address unobserved heterogeneity. Using the panel dimension allows us to control for unobserved household traits that are stable over time. In Table 16 we combine the set-up of Tables 13 and 14, but include household fixed effects.¹⁹ We again find a negative and statistically significant relationship between inflation expectations and the change in net worth, see column (5), confirming our main finding. A one percentage point increase in a household’s inflation expectation over time is associated with a 250 euro reduction in the household’s change in net worth per year on average.²⁰

A second potential concern is that inflation expectations are confounded with other household expectations, for example, some general optimism or pessimism about the economy. To the extent that optimism is a trait stable over time, this would be captured by the household fixed effects. To the extent that optimism is time-varying, this might be correlated with other household expectations on macroeconomic variables. We exploit the fact that the survey also contains a module with two other macroeconomic expectations: a point forecast on mortgage interest rates, and a point forecast on house prices in general. Both questions are asked to both renters and home-owners. The questions are asked in a different module than the question on inflation expectations (asked a few months apart). A consequence of the questions being in separate modules is that respondents can miss a survey: the number of observations with expectations on inflation, mortgage interest rates, and house prices drops by around 50%, and the number of households by approximately 40%.

Table 17 shows the results of the same specification as in Tables 13, 14 and 16, but with mortgage interest rate and house price expectations included. Dependent variables in Panel A are in levels (as in Table 13), dependent variables in Panel B are in changes (as in Table 14), and

¹⁹In columns (1)-(4) of Table 16 the dependent variable is the level, while in columns (5)-(8) the dependent variable is the first difference.

²⁰The coefficient on inflation expectations is also negative in all other columns, but not statistically significant.

regressions in even columns include household fixed effects (as in Table 16). Focusing on inflation expectations and changes in net worth, results remain virtually the same: in columns (1)-(2) of Panel B coefficients are negative and statistically significant. Other expectations matter to some extent for changes in net worth, but coefficients are smaller compared to inflation expectations and not always statistically significant. Furthermore, including other expectations increases the size of the coefficient on inflation expectations. For example, in column (2) of Panel B, a one percentage point increase in inflation expectations is associated with a 400 euro decrease in the change in net worth, which is larger than the 250 euro decrease reported in the corresponding specification in Table 16.

In sum, we obtain the same negative relationship between inflation expectations and changes in net worth after including household fixed effects and controlling for other economic expectations.

4 Conclusion

We study panel data on household inflation expectations to understand how households update inflation expectations at the micro level. We find that a noisy information model with two modifications (heterogeneous intercepts, and a small probability that households answer 10% instead of reporting their true inflation expectation) matches well the micro data. We then link the survey data on inflation expectations to administrative data on income, wealth, and car purchases. We find that households with higher inflation expectations save less (i.e., they have a smaller change in net worth). A one percentage point increase in expected inflation is associated with a 300 euro decrease in the change in net worth on average. We obtain a very similar estimate after including household fixed effects. A one percentage point increase in a household's inflation expectation over time is associated with a 250 euro decrease in the household's savings per year on average. Including other expectations increases the estimate to 400 euro. We also find that households with higher inflation expectations are more likely to acquire a car and acquire higher-value cars.

All these findings are consistent with a model where households form inflation expectations based on noisy signals on inflation and increases in inflation expectations stimulate spending. Viewed through the lens of such a model, we provide an estimate of the partial equilibrium effect of an increase in inflation expectations on spending, since we focus on idiosyncratic variation

in inflation expectations, while aggregate variation in inflation expectations is captured through time fixed effects. Providing an estimate of this partial equilibrium effect seems useful for two reasons. First, there is a vast literature (encompassing the entire literature on New Keynesian models) that *assumes* a tight link between inflation expectations and consumption-savings decisions at the household level. Moreover, this tight link is often central for conclusions regarding the effectiveness of monetary and fiscal policy. We find it reassuring that our linked survey-admin data supports the idea of a tight relationship between inflation expectations and consumption-savings decisions at the household level. Second, general equilibrium effects depend on the degree of price stickiness and the conduct of monetary policy and are therefore likely to vary across countries. General equilibrium effects can be computed, for example, with a general equilibrium model and estimates of the degree of price stickiness and monetary policy parameters. Given our results in Section 2 of this paper, it would also seem important that such a general equilibrium model features realistic assumptions about expectation formation of households.²¹

A fascinating next step would be to create exogenous variation in household inflation expectations (e.g., through an information-provision experiment where a set of randomly selected households is provided with inflation information) and study the effect on choices in administrative data. We hope this next step is feasible at some point.

Another fascinating next step would be to build dynamic stochastic general equilibrium models that are consistent with the estimates presented in Sections 2-3.

²¹See Wiederholt (2015) for a New Keynesian model with household inflation expectations that are heterogeneous and adjust sluggishly.

References

- [1] Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon (2017). “Forward Guidance and Heterogeneous Beliefs.” *American Economic Journal: Macroeconomics*, forthcoming.
- [2] Armantier, Olivier, Waendi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar (2015). “Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?” *International Economic Review*, 56(2): 505-536.
- [3] Arioli, Rodolfo, Colm Bates, Heinz Dieden, Ioana Duca, Roberta Friz, Christian Gayer, Geoff Kenny, Aidan Meyler, and Iskra Pavlova (2017): “EU consumers’ quantitative inflation perceptions and expectations: an evaluation.” *European Central Bank Occasional Paper Series*, No 186.
- [4] Bachmann, Ruediger, Tim Berg, and Eric Sims (2015). “Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence.” *American Economic Journal: Economic Policy*, 7(1): 1-35.
- [5] Bernanke, Ben S. (2007). “Inflation Expectations and Inflation Forecasting.” Speech at NBER Monetary Economics Workshop, 2007, available at <http://www.federalreserve.gov/newsevents/speech/Bernanke20070710a.htm>.
- [6] Burke, Mary A. and Ali K. Ozdagli (2013). “Household Inflation Expectations and Consumer Spending: Evidence from Panel Data.” *Federal Reserve Bank of Boston Research Department Working Papers*, No 13-25.
- [7] Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo (2011). “When is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119(1): 78-121.
- [8] Coibion, Olivier and Yuriy Gorodnichenko (2012). “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy*, 120(1): 116-159.
- [9] Coibion, Olivier, and Yuriy Gorodnichenko (2015). “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts.” *American Economic Review*, 105(8): 2644-2678.

- [10] Crump, Richard, Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa (2018). “Subjective Intertemporal Substitution.” *Federal Reserve Bank of New York Staff Reports*, No 734.
- [11] D’Acunto, Francesco, Daniel Hoang, and Michael Weber (2018). “The Effect of Unconventional Fiscal Policy on Consumption Expenditure.” Discussion Paper, University of Maryland, Karlsruhe Institute of Technology, and University of Chicago.
- [12] D’Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber (2019). “IQ, Expectations, and Choice.” NBER Working Paper No. 25496, January 2019.
- [13] Eggertsson, Gauti, and Michael Woodford (2003). “The Zero Bound on Interest Rates and Optimal Monetary Policy.” *Brookings Papers on Economic Activity*, 2003(1): 139-211.
- [14] Kaplan, Greg, and Sam Schulhofer-Wohl (2017). “Inflation at the Household Level.” *Journal of Monetary Economics*, 91: 19-38.
- [15] Krugman, Paul (1998). “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap.” *Brookings Papers on Economic Activity*, 1998(2): 137-205.
- [16] Maćkowiak, Bartosz, and Mirko Wiederholt (2009). “Optimal Sticky Prices under Rational Inattention.” *American Economic Review*, 99(3): 769-803.
- [17] Malmendier, Ulrike, and Stefan Nagel (2016). “Learning from Inflation Experiences.” *The Quarterly Journal of Economics*, 131(1): 53-87.
- [18] Patton, Andrew J., and Allan Timmermann (2010). “Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-sectional Dispersion.” *Journal of Monetary Economics*, 57(7): 803-820.
- [19] Sims, Christopher A. (2003). “Implications of Rational Inattention.” *Journal of Monetary Economics*, 50(3): 665-690.
- [20] Wiederholt, Mirko (2015). “Empirical Properties of Inflation Expectations and the Zero Lower Bound.” Discussion Paper, Goethe University Frankfurt.

- [21] Woodford, Michael (2003). “Imperfect Common Knowledge and the Effects of Monetary Policy.” In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps*, edited by Philippe Aghion et al. Princeton, NJ: Princeton University Press.
- [22] Yellen, Janet L. (2015). “Inflation Dynamics and Monetary Policy.” Speech at the Philip Gamble Memorial Lecture, University of Massachusetts, Amherst, available at <http://www.federalreserve.gov/newsevents/speech/yellen20150924a.htm>.