Who Cares about Inflation?

Endogenous Expectation Formation of Heterogeneous Households

Lukas Nord*

European University Institute

This Version: July 14, 2022

Abstract.

This paper builds a joint theory of endogenous inflation expectations and consumption-savings choices of heterogeneous households. We introduce imperfect information about future inflation rates in a consumption-savings model and allow households to exert costly effort to reduce uncertainty about future price changes. High wealth households are more exposed to future inflation due to its effect on real interest rates and hence choose to be better informed. The joint distribution of wealth and inflation expectations generated by the model is consistent with key features of the data. The implied consumption response to news about inflation is hump shaped in wealth: Wealthier households pay closer attention and update their expectations more in response to any signal received, but change their consumption less after any given update in expectations due to the income effect of future inflation. We show this mechanism to reduce the on-impact aggregate consumption response to forward guidance policies by up to 55% compared to an attentive counterfactual.

Keywords: Household Heterogeneity, Expectation Formation, Inflation Expectations.

JEL: D31, D84, E21, E31.

*lukas.nord@eui.eu. I am indebted to my advisors, Árpád Ábrahám and Russell Cooper, for their continuous guidance and support. This work benefited greatly from helpful comments and discussions with Annika Bacher, Tobias Broer, Thomas Crossley, Wouter Den Haan, Dimitris Georgarakos, Nils Grevenbrock, Philipp Grübener, Carl-Wolfram Horn, Kurt Mitman, Ben Moll, Stefan Ried, Fabian Stürmer-Heiber, and Mirko Wiederholt, my discussant Fabio Rummler as well as workshop and conference participants at Bonn, Bundesbank, and EUI. All remaining errors are mine. In this paper use is made of data of the DNB Household Survey administered by CentERdata (Tilburg University, The Netherlands).

1 Introduction

In recent years, central banks have begun to rely more and more on forward guidance – influencing households' behavior through signals about future macroeconomic outcomes such as inflation – as a policy instrument to stimulate current demand. Who adjusts their expectations in response to news about future inflation remains as much an open issue as how heterogeneous households respond to changes in their expectations. Understanding both of these points is essential to evaluate the effectiveness of forward guidance policies.

When central banks want to stimulate current demand by signaling higher inflation for future periods, the response of consumption is determined in two steps: First, households need to update their expectations based on the signal they receive. Second, they respond to their updated expectations by adjusting their consumption behavior. This paper considers both of these steps jointly by introducing endogenous expectation formation in an otherwise standard consumption-savings model with heterogeneous households. We show that wealth is an important determinant of both households' incentives to pay close attention to signals about future inflation rates and their consumption response to any given change in their inflation expectations. Based on this finding we argue that allowing for an endogenous wealth effect on the formation of households' inflation expectations substantially reduces the responsiveness of aggregate consumption to signals about future inflation rates, because those likely to adjust their expectations are less responsive in their consumption. Explicitly accounting for heterogeneity in wealth and its impact on the formation of inflation expectations hence suppresses the effectiveness of forward guidance policies at its origin.

For a discussion of the incentives to pay attention to future inflation it is important to note that heterogeneous households are not exposed to inflation uniformly. Through its effect on real interest rates households are more affected by inflation the more they borrow or save between periods. If expectation formation is in any way costly, we should therefore consider households with higher wealth – and hence larger exposure – to be more willing to face these cost.[1] How heterogeneous households respond to expected inflation likewise depends on their wealth holdings. Any given change in (expected) inflation rates will have different consequences for households with different asset levels.

[1] In this paper, the focus lies on inflation as a risk to the real interest rate that affects all saving and borrowing uniformly. Wages are real and all households face the same effective inflation rate. This assumption is discussed in later sections.

While an increase in expected inflation is good news for debtors since it reduces the real value of their future repayments, it is bad news for savers who would be the recipients of those payments. This heterogeneous income effect makes households' consumption response to any given change in inflation expectations a declining function of their wealth. Taken together, both effects imply a negative correlation between households' updating their expectations and their potential consumption response to changes in expectations, dampening the responsiveness of aggregate consumption to signals about future inflation rates.

This paper formalizes the intuitive arguments above in a theoretical framework, disciplining it with empirical observations on inflation expectations along the wealth distribution. Our approach is novel in considering the joint formation of inflation expectations and consumption-savings decisions of heterogeneous households.

We begin by developing a theory of endogenous inflation expectations. Households are assumed to understand the underlying inflation process and to be uncertain only about future inflation rates. To reduce this uncertainty, they can exert costly effort. The proposed framework is sufficiently tractable to integrate it into a heterogeneous agents model while capturing key features of the data. Analytical results show that this model of expectation formation implies the standard deviation of forecast errors as well as the mean absolute error across a group of households to be decreasing in the effort they exert. We use these results to discipline our model with cross-sectional statistics from the joint distribution of inflation expectations and wealth, making use of the Dutch National Bank's Household Survey. We find the standard deviation of forecast errors and the mean absolute error across households to be decreasing in absolute wealth in the data. Both richer as well as indebted households have more precise and less dispersed expectations compared to those around zero net wealth. Integrating the proposed model of expectation formation into an infinite-horizon consumption-savings problem allows us to study jointly the formation of and response to households' inflation expectations. The calibrated model matches the empirically observed pattern of forecast errors along the wealth distribution. Households with higher net savings or debt endogenously choose to be better informed about future inflation as they are more exposed to inflations' effect on real interest rates.

The theoretical framework allows us to back out the consumption response to any signal about future inflation rates – the marginal propensity to consume on signal (MPCS)

– from households' policy functions. We show that the MPCS depends on two factors: How much a household updates its expectations in response to any signal received, and how it reacts to any given change in its expectations. Households' consumption response to any given change in expectations is decreasing in wealth due to an expected income effect. This income effect arises as higher expected inflation ceteris paribus reduces the expected future value of savings. In contrast, as richer households are endogenously paying closer attention to inflation rates, they update their expectations more in response to any signal received, making their consumption more responsive to news about inflation. Combined, these two forces yield a hump shaped pattern for MPCS' along the wealth distribution.[2]

To highlight the importance of our findings at the aggregate level, we conduct a forward guidance exercise within our framework. Capturing the on-impact effect of forward guidance, we simulate the aggregate consumption response to a one percentage point increase in all signals received by households about next period's inflation. We show that under endogenous expectation formation, forward guidance misses out on up to 55% of the effect it could have if all households choose to be as informed as the most attentive. This result is driven by low wealth households, who are potentially most responsive to any change in their inflation expectations, but fail to update their expectations in response to the signal as they do not pay close attention to news about future price changes. The (richer) households paying attention to the signal and updating their expectations in response perceive higher inflation as a loss in their real income, yielding a relatively lower consumption response.

Most existing models of inflation expectations, summarized in Coibion and Gorodnichenko (2012), study expectation formation in isolation from other household choices and often abstract as well from underlying heterogeneity among agents, Madeira and Zafar (2015) being one of few exceptions. We extend this work to include wealth as a direct determinant of expectation formation when precise expectations are costly. In this regard we are closest in spirit to the literature on rational inattention founded by Sims (2003) and surveyed in Mackowiak et al. (2018). In a heterogeneous agent framework, Carroll et al. (2020) and Auclert et al. (2020) introduce sticky expectation formation but abstract from endogeneity of expectations with respect to households' idiosyncratic state.

[2]The group of indebted households who would be both likely to update their expectations and strongly respond in their consumption is small and therefore quantitatively less important.

We share the analysis of heterogeneous incentives to form precise expectations with Broer et al. (2018). They discuss the endogenous choices of households to use precise laws of motion for aggregate capital in an economy à la Krusell and Smith (1998) and find substantial heterogeneity in the utility loss from not using full information. In their model, choices are based on simulated lifetime utilities and forecasting capital impacts forecasts about both returns and wages, accounting for the difference to our findings. Recent work has found households' consumption responses to income shocks to be an important determinant of Macroeconomic outcomes, leading to a large and growing literature on heterogeneous households' marginal propensity to consume out of transitory income (MPC), sampled e.g. in Kaplan and Violante (2021). In contrast, we focus on households marginal propensity to consume in response to signals about future inflation (MPCS) and show that heterogeneity along this margin is important to consider for policy analysis. Studying forward guidance in a framework with heterogeneous agents, McKay et al. (2016) show how occasionally binding borrowing constraints can reduce the responsiveness of aggregate consumption to interest rate changes in the distant future, alleviating the so called *Forward Guidance Puzzle.* Compared to their framework with full information, we show how an endogenous correlation between expectation updating and consumption responses can dampen the effects of forward guidance also in the short run. Similar to previous theoretical approaches, most empirical work on inflation expectations has as well abstracted from wealth as a potential determinant of expectation formation. Closest to our analysis is Ben-David et al. (2018) who study the relation between uncertainty about macroeconomic variables such as inflation or house prices and socio-economic status of households. Their data does not include households' asset holdings but they find uncertainty about macroeconomic variables to decrease in income and employment – both highly correlated with wealth. Another strand of the empirical literature considers the impact of expectations on households' consumption savings choices. Among others, Armantier et al. (2015), Crump et al. (2015), Dräger and Nghiem (2018), or Vellekoop and Wiederholt (2019) evaluate the consistency of households' choices with their expectations. Coibion et al. (2019) study how expectations and consumption respond to exogenous news about inflation. They find a negative consumption response to higher inflation, driven by high wealth households in line with our model. Also in line with our theory, Lieb and Schuffels (2019) find the likelihood of positive durable consumption expenditure in response to higher inflation expectations to be decreasing in wealth. We contribute to this literature by highlighting the importance of considering jointly heterogeneity in households' incentives to form precise expectations and their potential response to such expectations.

The remainder of the paper is organized as follows: Section 2 introduces a framework for endogenous inflation expectations. Section 3 presents empirical findings on the joint distribution of wealth and inflation expectations. Section 4 incorporates the endogenous expectation framework in a consumption-savings model with heterogeneous households. Section 5 analyses households' consumption responses to news about inflation and discusses aggregate consequences of endogenous expectation formation. Section 6 concludes.

2 Modeling Endogenous Inflation Expectations

To allow for two-way interactions between households' consumption-savings choices and their inflation expectations, we require a model with endogenous expectation formation. This section provides an endogenous expectation framework, which will later be incorporated into a consumption-savings problem. In the interest of a clear exposition and computational tractability, we keep the expectation formation process as simple as possible while at the same time rich enough to account for key features of the data. Households are assumed to understand the underlying process and perfectly observe current inflation but to be uncertain about the shock component to future inflation rates. This uncertainty can be reduced endogenously by households exerting costly effort. The setup yields heterogeneous effort choices if the gains of reducing uncertainty about future inflation rates are distributed unevenly across households. In this section, focus lies on how effort transmits into individual expectations and cross-sectional moments of expectation errors. Section 4 discusses households' effort choice.

Assume inflation follows a first-order autoregressive process

$$\pi_{t+1} = (1-\rho)\mu + \rho\pi_t + e_{t+1} \qquad e_{t+1} \sim \mathcal{N}(0, \sigma_e^2). \tag{1}$$

π_t is inflation in period t, μ is the long run mean of inflation and ρ its persistence across periods. e_t is a shock to inflation, which is i.i.d across time.

Households know that inflation follows (1) and agree about (the true) μ, ρ and σ_e^2. In contrast to most of the literature on expectation formation, households perfectly observe current and all past inflation rates.[3] In period t, π_τ is known for all $\tau \leq t$. This assumption keeps the state space of the household problem small. We believe this is justified, given that information about current and past inflation rates is easily accessible online.[4] Furthermore, when embedding the expectation formation process in a consumption-savings model, it will be important for households to know current prices in order to pin down their budget set in real terms. Therefore, households are assumed to be uncertain only about future inflation rates.

Households form expectations with respect to the shock to future inflation, e_{t+1}. In period t, household i can exert some effort n_t^i to influence the noise in a signal $\hat{e}_{t+1}^i$ he receives about next period's shock. The signal he receives follows

$$\hat{e}_{t+1}^i = e_{t+1} + s_{t+1}^i \qquad s_{t+1}^i \sim \mathcal{N}(0, \sigma_s^2(n_t^i)), \tag{2}$$

where the noise component s_{t+1}^i can be influenced by households' effort choice. We assume its standard deviation to be a decreasing but convex function of effort ($\sigma_s'(n) < 0,\ \sigma_s''(n) > 0$) and s_{t+1}^i to be pure noise, i.e.

$$e_t \perp\!\!\!\perp s_t^i \ \forall i, t \qquad s_t^i \perp\!\!\!\perp s_t^j \ \forall i, j, t \qquad s_t^i \perp\!\!\!\perp s_{t+s}^i \ \forall i, t, s. \tag{3}$$

Households have identical priors about the shock corresponding to the true unconditional distribution $e_{t+1} \sim \mathcal{N}(0, \sigma_e^2)$.[5] Based on the signal received, the household updates his prior belief according to Bayes Rule. Let $\omega_{t+1}^i(n_t^i) = \frac{\sigma_e^2}{\sigma_e^2 + \sigma_s^2(n_t^i)}$ be the weight he attaches to the signal, yielding his posterior belief about the shock as

$$e_{t+1}|_{\hat{e}_{t+1}^i, n_t^i} \sim \mathcal{N}\left(\omega_{t+1}^i(n_t^i)\hat{e}_{t+1}^i, \omega_{t+1}^i(n_t^i)\sigma_s^2(n_t^i)\right). \tag{4}$$

[3]See e.g. Vellekoop and Wiederholt (2019).

[4]One can reinterpret our assumption as the first marginal bit of effort providing full information about present and past inflation rates. With the assumptions imposed below on the cost of effort, the first marginal bit of information is costless and hence always obtained.

[5]Heterogeneity in prior variance, especially if correlated with households wealth, would complicate the analysis substantially. Assuming a common (unbiased) prior about the mean of the shock is without loss of generality. Relaxing this assumption would yield similar results as introducing heterogeneous beliefs about μ, see appendix B.1.

Household i's expected value for inflation is determined by his expectation about the future shock and is given as

$$\mathbb{E}_t[\pi_{t+1}|\hat{e}^i_{t+1}, n^i_t] = (1-\rho)\mu + \rho\pi_t + \omega^i_{t+1}(n^i_t)\hat{e}^i_{t+1} \tag{5}$$

implying an ex-post forecast error of

$$err^i_{t+1} = \mathbb{E}_t[\pi_{t+1}|\hat{e}^i_{t+1}, n^i_t] - \pi_{t+1} = \omega^i_{t+1}(n^i_t)s^i_{t+1} - (1-\omega^i_{t+1}(n^i_t))e_{t+1}. \tag{6}$$

The first term in the error captures households' over-reaction to noise while the second term captures under-reaction to news contained in the signal, as is standard in models with Bayesian updating.

The standard deviation of households' posterior belief about future inflation can be referred to as their *subjective uncertainty* (SU) and is given by

$$SU^i_{t+1} = \sqrt{\omega^i_{t+1}(n^i_t)\sigma^2_s(n^i_t)} = \sqrt{\frac{\sigma^2_e\sigma^2_s(n^i_t)}{\sigma^2_e + \sigma^2_s(n^i_t)}}. \tag{7}$$

Under the assumptions that $\sigma'_s(n) < 0,\ \sigma''_s(n) > 0$, one can show that SU is a decreasing and convex function of households' effort n.

We can also derive theoretical moments for a group g of households with equal choices for $n^i_t = \bar{n}^g_t$. An equal choice for n implies identical weights $\omega^i_{t+1}(n^i_t) = \omega^g_{t+1}(\bar{n}^g_t)$. The model implies that the forecast errors within a group g will be normally distributed with a variance, across households and time, given by

$$\mathrm{Var}^g(err^i_{t+1}) = (\omega^g_{t+1}(\bar{n}^g_t))^2\sigma^2_s(\bar{n}^g_t) + (1-\omega^g_{t+1}(\bar{n}^g_t))^2\sigma^2_e = \frac{\sigma^2_e\sigma^2_s(\bar{n}^g_t)}{\sigma^2_e + \sigma^2_s(\bar{n}^g_t)} = \left(\overline{SU}^g_{t+1}\right)^2, \tag{8}$$

where we make use of the assumption that noise is uncorrelated across households. Hence the within group standard deviation of forecast errors across households exerting effort $\bar{n}^g_t$ can be interpreted as the standard deviation of the posterior belief about future inflation of households in that group $\overline{SU}^g_{t+1}$. With our model, disagreement among households becomes a measure of how noisy a signal about future inflation rates these households chose to receive.

As an additional measure of forecast precision, we can derive the mean absolute error of households by using the fact that s and e are normally distributed and uncorrelated with each other and over time. This implies a normal distribution for the expectation error among a group of households with mean zero and variance given in (8). By the properties of folded normal distributions, the average absolute error is given as

$$\mathbb{E}^g[|err^i|] = \sqrt{\mathrm{Var}^g(err^i_{t+1})\frac{2}{\pi}} = \overline{SU}^g_{t+1}\sqrt{\frac{2}{\pi}}. \tag{9}$$

Therefore, our model predicts a strong co-movement of the standard deviation and mean absolute error for a group of households exerting similar effort, driven by how noisy a signal they chose to receive about future inflation.

An important implication of our theoretical results is that they rationalize the use of cross-sectional moments to learn something about households' expectation formation. They allow us to discipline a model of joint consumption-savings choices and expectation formation with the standard deviation of expectation errors at different points of the wealth distribution. Before we turn to incorporating endogenous expectations into a consumption-savings framework we therefore report in the next section on the joint distribution of expectation errors and wealth in the data.

In order to reduce the state space of the problem and incorporate it into a heterogeneous agent framework, we have kept the expectation formation process as simple as possible, but sufficiently rich to account for key features of the data. In Appendix B.1 we show that the results are robust to additional sources of heterogeneity in expectations, such as fundamental disagreement about the long run mean of inflation.

3 Expectations Along the Wealth Distribution

This section presents empirical observations on the joint distribution of households' wealth levels and inflation expectations, which we will use to discipline our model. As suggested by the results of the previous section, we study the cross-section of households at different points of the wealth distribution and focus on two statistics: The standard deviation of forecast errors and the average absolute forecast error. After outlining the data used and methodology applied we present our baseline findings before concluding with some additional robustness tests.

3.1 Data and Methodology

To gain insight into the joint distribution of inflation expectations and wealth we use data from the Dutch National Bank's Household Survey (DHS). This dataset is unique in providing comprehensive data on both households' wealth and their inflation expectations. We combine observations from the survey waves 2010-2018. The choice of period reflects changes made to the questionnaire on inflation expectations in the 2008 wave and excludes the financial crisis episode. We use data at the individual level, as presented in the DHS, but restrict our sample to household heads to avoid within household correlations. We take heads' answers to be representative for their household.

In the survey, households are simultaneously asked about their current wealth in a variety of asset classes and their expectations of one year ahead inflation. We compute households' net financial wealth as the sum of all assets less liabilities reported in the DHS, excluding houses and related mortgages, business equity and vehicles. Whenever referring to "wealth" in the remainder of this paper, we apply this definition. In the baseline results, we focus on financial wealth as we believe it to capture best the resources out of which the household decides to consume or save in response to changes in inflation rates.[6] Using the described wealth measure, we construct decile groups based on households' position in the wealth distribution in the year of observation. We pool observations across waves that are in the same wealth decile for their wave. Table A.1 in the appendix reports summary statistics for these groups.[7] It also shows that results are robust to pooling all observation across years and defining wealth deciles based on the full sample.

Participants in the DHS are asked to report a point forecast for the inflation rate over the following 12 months, choosing from the set of whole numbers between 1 and 10. Ex-post errors are computed by subtracting the realized inflation rate over the next 12 months from this forecast. As the exact month of the observation is unknown (the survey generally takes place between April and October each year), we subtract June-to-June inflation as an approximation to the forecasted rate. As an example, for

[6] Previous real estate or durable goods purchases are unlikely to be re-considered in response to small fluctuations in expected inflation rates.

[7] The table shows different numbers of missing observations for inflation expectations across wealth deciles, with the highest number of missings in the second decile. To test robustness with respect to differential numbers of missings, we have constructed bounds in the spirit of Lee (2009). All main findings are robust to the number of missing values across deciles. Results are omitted for brevity but are available upon request.

an observation of the 2016 wave inflation is the change in the Dutch CPI between June 2016 and June 2017.[8]

3.2 Empirical Observations

Our focus is on observations at the wealth decile group level. For each wealth decile group, we report the within-group standard deviation of forecast errors and the mean absolute error.[9] Figure 1 presents our baseline empirical results. Both the mean absolute error and the within-group standard deviation increase between the first and second decile and decline as wealth increases further until reaching a stable level in the upper half of the wealth distribution. At its lowest level, both variables are about 0.6 pp. lower than at their peak in the second decile. Both the initial increase and the subsequent decline are statistically significant at the 95% level.

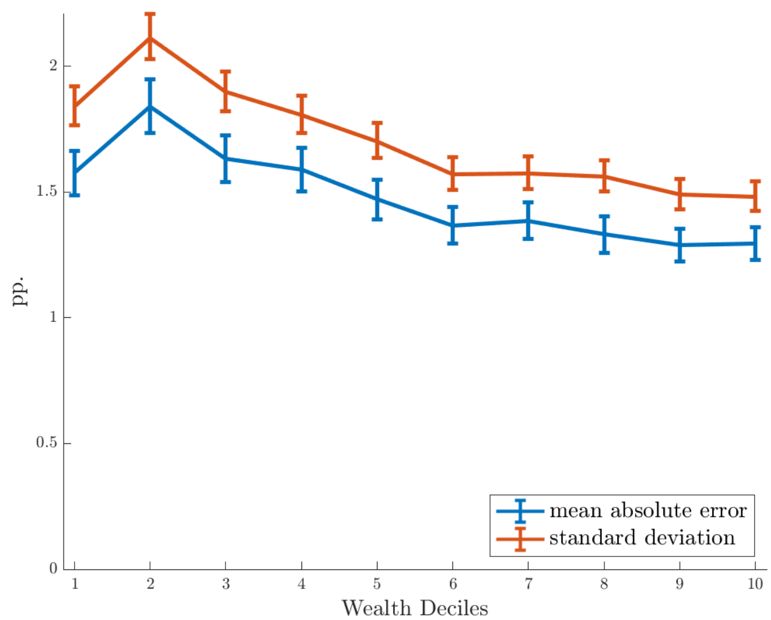

Figure 1: Expectation Errors by Wealth Decile Groups

The figure plots the within-decile group standard deviation of errors and the mean absolute forecast error. Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018.

[8]None of the results is altered qualitatively if June-to-June inflation is replaced by inflation in the current year (annual inflation in 2016 in our example) or the following year (annual inflation in 2017).

[9]Baseline results are unweighted. Using household weights has no significant impact on our findings.

For the interpretation of these findings it is important to note that the second wealth decile group is centered around zero net financial wealth, i.e. net debtors are concentrated in the first decile group. Through the lens of our model of expectation formation, the results suggest that wealthier as well as indebted households choose to exert higher effort in order to form precise expectations about future inflation rates.

To validate our approach to modeling households' expectation formation, Figure A.6 in the appendix plots the histograms of errors by decile. Our theory would suggest that these errors should be normally distributed within decile groups. Despite limitations such as the discreteness and truncation of expectation data due to the sample question in the DHS, the fitted kernel densities align well with their respective normal counterparts.

The role of age and education

It is well established that other demographic characteristics are highly correlated with positions in the wealth distribution.[10] The two most important for our analysis are age and education. An argument can be made that more experienced (as older) people could be better at forming expectations. Similar argument applies for more educated individuals. As education and age correlate positively with wealth this could be driving the finding in Figure 1. As we include neither education nor age in our model, we test for robustness and repeat our analysis controlling for age and education respectively.

Testing for robustness towards age and education, we look at the data on quintile group level to allow for a sufficient number of observations within each age/wealth and education/wealth cell. At the quintile level, debtors are pooled with households around zero wealth. Figure 2 reports the within wealth quintile group standard deviation of errors by age groups and education groups. The general downward trend of disagreement in wealth persists after controlling for either age or education. Age appears to have little explanatory power beyond the impact of wealth, providing an argument against experience as a driving force for expectation formation. College education, however, appears to somewhat decrease disagreement compared to less educated groups. Similar findings hold for the mean absolute expectation error.[11]

[10] See e.g. Cooper and Zhu (2016).

[11] These results are presented in Figure A.4 in the appendix.

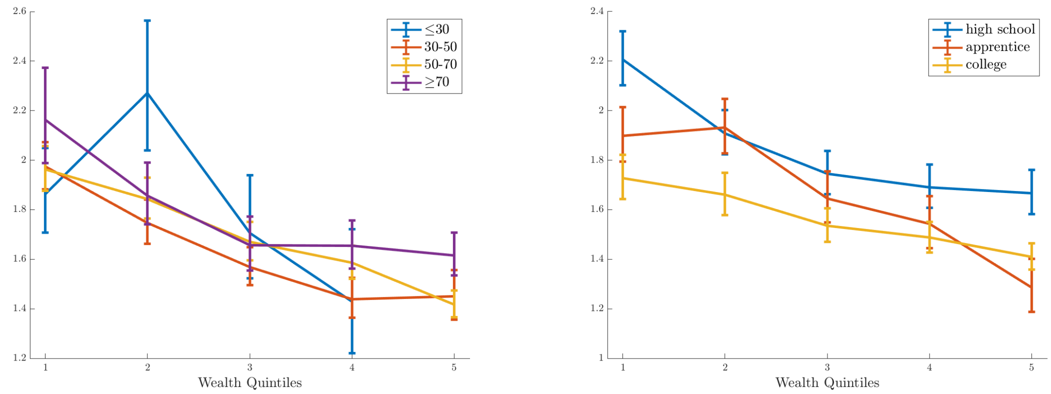

(a) by age

(b) by education

Figure 2: Standard Deviation of Expectation Errors by Wealth Quintiles – Controls

The figure plots the within-quintile group standard deviation of errors by age (a) and education groups (b). Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018. Combination of youngest age and highest wealth quintile omitted due to lack of observations.

Measure of wealth

To test robustness with respect to the considered measure of wealth, we repeat the analysis dividing households into decile groups based on two alternative measures: A first including both housing and associated mortgages as well as a second considering only positive financial assets. Including housing wealth leaves the results qualitatively unchanged, as Figure 3 shows. The peak of both mean absolute error and standard deviation of errors remains in the second decile group (again around zero net wealth). Both decline to either side and the overall decline between peak and trough in both variables is of similar magnitude as before. Different from previous results there is a hump shaped pattern between the 4th and 10th decile especially in the standard deviation of errors. This is perfectly in line with the correlation of financial wealth and housing wealth: Median financial wealth increases up to the 4th decile of wealth including housing, declines again for deciles 5 to 7 before increasing substantially for deciles 8 to 10. We take this as further support for financial wealth as the relevant measure to consider.

Excluding debt from the wealth measure, we find that both mean absolute errors and the standard deviation of errors are declining in asset holdings as Figure 4 shows. This is as expected given the limited amount of financial debt (and hence limited netting

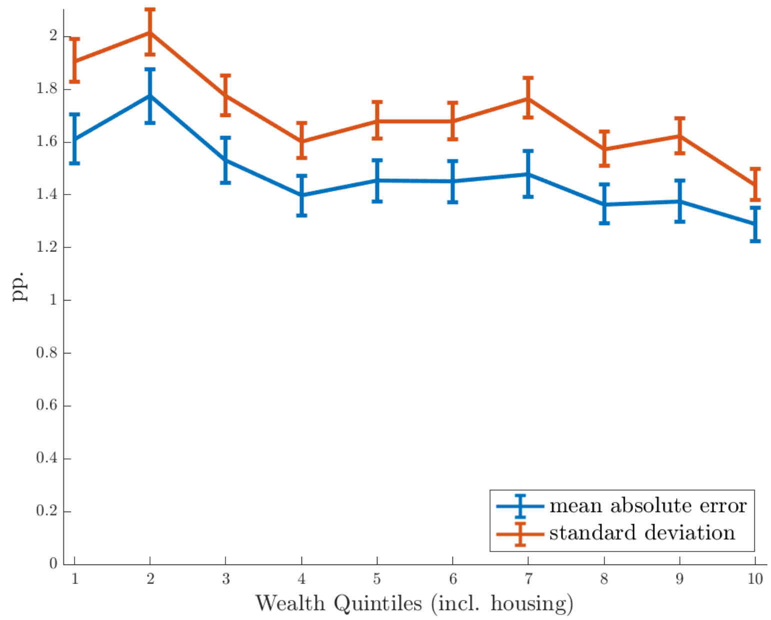

Figure 3: Expectation Errors by Wealth Decile Groups – Housing

The figure plots the within-group standard deviation of errors and mean absolute errors by net financial wealth decile groups, including housing and mortgages in the wealth measure. Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018.

of financial asset positions) in the DHS dataset.[12] Again, the overall decline in both measures along the wealth distribution is of similar magnitude as in the baseline results, and both flatten over the the 6th-10th decile of total financial assets.

Individual Level Analysis

While measuring the standard deviation of errors requires us to pool households into groups, differences in the absolute forecast error can also be tested at the household level. To do so, we regress households' absolute forecast error on indicators for their wealth decile and controls. The specification is given in equation (10).

$$abs(err^i_{t+1}) = \alpha + \sum_{d=2}^{10} \beta^d \mathbb{1}_{dec_{i,t}=d} + \gamma X_{i,t} + \epsilon_{i,t} \tag{10}$$

[12]Median debt is zero for all but the first decile of financial wealth and averages liabilities are below EUR 1,000 for deciles 2 to 10.

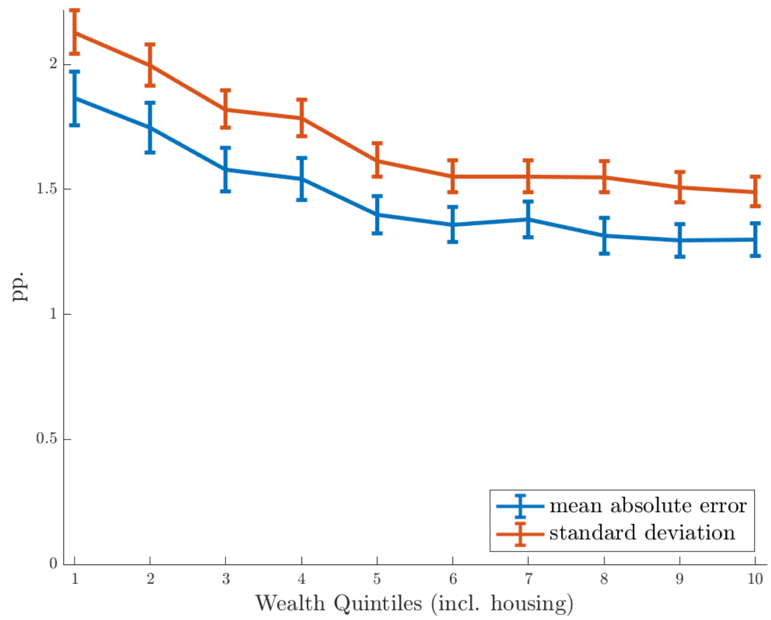

Figure 4: Expectation Errors by Wealth Decile Groups – No Debt

The figure plots the within-group standard deviation of errors and mean absolute errors by total financial assets. Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018.

The controls $X_{i,t}$ include indicators for households age and education group as well as home ownership status.[13] Coefficients β^d must be interpreted as the difference in absolute forecast errors of households in decile d relative to households in the first wealth decile. Figure 5 plots the coefficients β^d along with the corresponding findings from Figure 1 for comparision. The similarity of both lines in Figure 5 suggests that controlling for age, education and homeownership does not alter the findings of absolute forecast errors along the wealth distribution. Full results are reported in Table A.2 in the appendix and show an insignificant effect of age and education (except for college degrees) but significantly lower errors for home owners.

Additional data sources

For further robustness, we repeat our analysis with US data from the Michigan Survey of Consumers (MSC). Compared to the DHS data, the MSC contains substantially less

[13]The estimation of β^d relies on variation in wealth across households. The median household is in the sample for 3 years and the analysis is carried out at annual frequency, making it difficult to obtain sufficient within household variation in wealth to be able to control for household fixed effects.

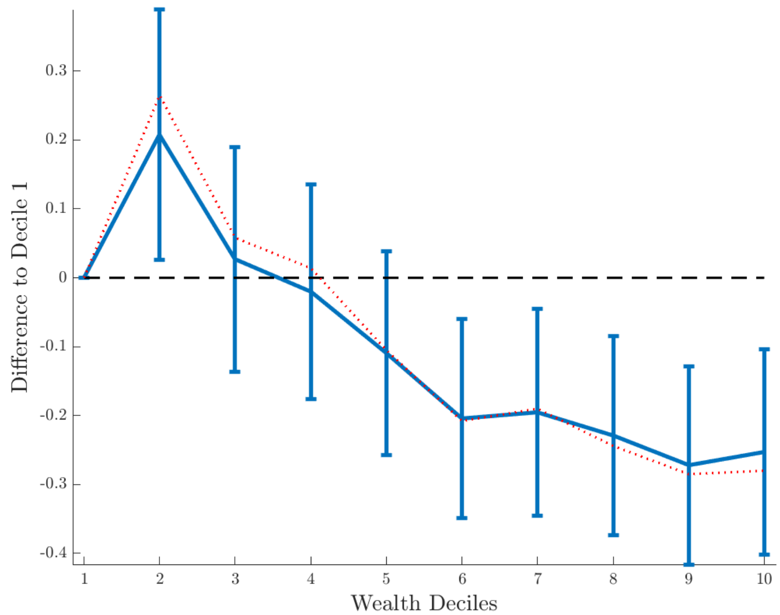

Figure 5: Mean Absolute Error - Individual Level Regression

The figure plots in blue the etimated coefficients β^d from running (10) at the household level. Bars provide confidence bands at the 95% level, based on standard errors clustered at the household level. For comparison, the red dotted line reports the corresponding results from Figure 1, i.e. the difference in mean absolute errors vs. the first decile group.

information on household wealth. Reported stock market investment has to be used as an approximation of net financial wealth and hence there are no households with negative wealth, making it impossible to test for a hump shaped pattern. Detailed findings are provided in Appendix A.1. The patterns reported for the DHS are strongly supported by findings from the Michigan Survey of Consumers, i.e. both the standard deviation of errors and the mean absolute forcast error are strictly decreasing in households stock market investment.

Alternative Mechanisms

While the mechanism under study in this paper relies on wealth levels influencing households' expectation formation, alternative mechanisms might be proposed to explain the reported patterns. The literature often imposes causality to run from the level of expectations to wealth levels, higher inflation expectations implying lower savings, abstracting from any reverse effect (see e.g. Vellekoop and Wiederholt, 2019; Crump et al., 2015). There are two important differences, in timing and in the moments considered, compared to the present paper. First, while this literature speaks to how the level of expectations on inflation between yesterday and today impact today's wealth levels, we focus on how today's wealth impacts dispersion in expectations between today and tomorrow. Second, while previous work has focussed on the mean inflation forecast at the household level, we focus on the dispersion in forecast errors across households.

A recent literature has found heterogeneity in realized inflation rates due to heterogeneity in households' consumption baskets, higher income households experiencing lower inflation (see e.g. Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021). While systematic differences in realized inflation rates could account for the observed dispersion in expectations across groups, it is unlikely to account for the pattern of dispersion in expectations within wealth deciles along the distribution as consumption baskets are strongly correlated with income/wealth.

Another possible driver of the observed patterns could be that more financially literate households are at the same time wealthier and better able to form expectations about inflation. Direct measures of financial literacy are available in the DHS only for a special module in 2005, i.e. not in our sample (Deuflhard et al., 2019). However, patterns are robust by education and age group, both likely correlated with financial literacy.

In our analysis we focus on net financial wealth and abstract from portfolio composition. This is justified by looking at the share of wealth held through checking, savings or deposit accounts, savings certificates and deposit books in the DHS. These assets arguably have predetermined, nominal interest rates and therefore all carry the same one-for-one exposure to inflation. In our sample the share of these assets over total assets is around 80% for most parts of the wealth distribution except for the wealthiest households, where it drops to about 50% for the top decile. Therefore, changes in portfolio composition are unlikely to account for the decline in dispersion of forecast errors between the second and sixth decile of the wealth distribution. In theory, an active portfolio choice would also make incentives to learn about inflation an increasing function of beginning of period wealth. As long as we assume portfolio composition to be adjustable at annual frequency, what matters is not the initial exposure to different assets but overall wealth, as the total size of the portfolio determines the benefits of forming precise expectations and possibly adjusting its composition going forward.[14]

Taking stock

Our empirical findings show that both the standard deviation as well as the mean absolute error across households co-move with households' wealth in a meaningful way: Richer and indebted households exhibit lower dispersion and mean absolute errors in their forecasts of inflation compared to their counterparts around zero net wealth. The findings are robust to covariation with age or education as well as our definition of wealth and can be replicated in US data. These results are in line with our theory of expectation formation if richer and indebted households choose to exert higher effort to learn about future inflation. We use these findings to discipline a consumption-savings model with endogenous expectation formation of heterogeneous households in the following section.

4 Savings Choice and Endogenous Expectations

Building on the results of the previous sections, we are now in a position to incorporate the expectation formation presented in Section 2 into an infinite-horizon consumption-savings model. The model explicitly considers the effect of wealth on households' expectation formation and at the same time allows us to trace out their responses to changes in expectations. The dynamic setting generates a joint distribution of expectations and wealth, which can be validated against the empirical findings in Section 3.

4.1 Household Problem with Endogenous Expectations

At the beginning of each period, a household knows the assets carried over from the previous period a and learns about his real income y as well as the current inflation rate π, which are both stochastic over time. Together, these variables determine the available resources for consumption and saving. Based on this information, the household decides on his effort n. After deciding on n, he receives a signal about the shock to inflation between the current and the next period and updates his belief about future inflation. He will base his choice over consumption today and savings on the updated belief. Households' income is assumed to follow a Markov process with transition matrix Π_y. We assume income y to be real income.[15] Savings and borrowing are subject to

[14] See Peress (2004) for theoretical results along these lines and the related discussion in Section 4.3.

[15] This choice is motivated by the fact that labour income, the largest component of non-financial income, for the Netherlands over our sample period is to a large extend protected from inflation through

a nominal interest rate. We abstract from interest rate risk and assume the nominal interest between any two periods to be constant at r^n. We do so to discuss the effect of inflation risk in isolation. Qualitatively, the findings presented below rely on this assumption only to the extent that nominal interest rates do not move one-for-one with inflation. As long as nominal rates co-move disproportionately, changes in inflation will induce fluctuations in the real interest rate. A constant nominal interest rate together with real income define inflation in our model effectively as a risk only to the real interest rate.

With all other notation as introduced above, households' information choice problem is given as

$$\tilde{V}(a, y, \pi) = \max_{n \in [0, \bar{n}]} \mathop{\mathbb{E}}_{\hat{e}'}[V(a, y, \pi, n, \hat{e}')|n] \tag{11}$$

where we restrict the choice of effort to be positive and impose an upper limit $\bar{n}$ on how much the households can learn about future inflation to rule out perfect foresight.

The subsequent consumption-savings choice, conditional on chosen effort n and received signal $\hat{e}'$, can be described as the solution to

$$\begin{aligned} V(a, y, \pi, n, \hat{e}') = \max_{c, a'} & \left(c^{1-\gamma} + \beta \left(\mathop{\mathbb{E}}_{\pi', y'} [\tilde{V}(a', y', \pi')^{1-\alpha} | \hat{e}', n, \pi, y] \right)^{\frac{1-\gamma}{1-\alpha}} \right)^{\frac{1}{1-\gamma}} \\ s.t. \quad & c + a' = \frac{1 + r^n}{1 + \pi} a + y - \mathcal{F}(n) \\ & a' \geq \bar{a}, \ c \geq 0 \end{aligned} \tag{12}$$

where the budget constraint is written in real terms, a is today's nominal asset level divided by yesterday's prices and $\bar{a}$ is the borrowing limit.[16] Expectations over π' are based on households' updated belief taking into consideration π, $\hat{e}$ and the previous choice for n. The law of motion of inflation and the expectation formation based on the signal are as presented in Section 2. Preferences of the household are recursive as in Epstein and Zin (1989), allowing for independence of risk aversion and intertemporal substitution.

collective bargaining agreements. According to OECD data, collective bargaining coverage in the Netherlands was well above 80% for the period under study.

[16]For details see appendix B.2.

We model the cost of effort as a monetary cost, representing both the opportunity cost of spending time on forming expectations as well as the cost of acquiring information. For the cost of effort and the relationship between effort and noise in the signal, we assume functional forms

$$\sigma_s(n) = \frac{\chi}{1+n} \quad \text{and} \quad \mathcal{F}(n) = (\theta n)^{\phi}. \tag{13}$$

These choices yield convex cost of and convex gains from exerting effort.[17] Note that with these functional forms, χ is the variation in the noise if zero effort is exerted, i.e. the maximum variation possible, and that zero effort implies zero cost.

4.2 Calibration

The calibration of the model aims to replicate the patterns presented in Figure 1. Our calibration strategy is twofold: First, a range of parameters is set exogenously. These include preference parameters $\gamma = 1.5$ and $\alpha = 8$ which we chose in line with previous work.[18] We furthermore assume the cost of information to be quadratic ($\phi = 2$). The inflation process is estimated from Dutch annual inflation rates for the period 1988-2018. This yields a long run mean of about 2% and an annual persistence of about 0.5, similar to the estimates of Vellekoop and Wiederholt (2019). The nominal interest rate r^n is set at 4% for a steady state real rate of 2%. Second, we calibrate β, $\bar{a}$, θ, χ and $\bar{n}$ as well as the process for y jointly for the model to fit the data on households' wealth and their expectation errors along the wealth distribution. Calibration targets include the position of the peak of households' errors in the second decile, the beginning of the flattened part of the standard deviation of errors in the sixth decile, the magnitude of the drop in error standard deviation of 0.57pp[19] as well as the share of wealth held by each decile of the wealth distribution. All parameters (and their interaction) influence a wide range of model statistics. Nevertheless, β and $\bar{a}$ are particularly important to determine the lower end of the wealth distribution while θ, χ and $\bar{n}$ reproduce the slope and level of errors along the wealth distribution. Bounding n generates a flat standard

[17] $\sigma_s'(n) < 0,\ \sigma_s''(n) > 0$ and $\mathcal{F}'(n) > 0,\ \mathcal{F}''(n) \geq 0,$ iff $\phi \geq 1$.

[18] Papers applying Epstein-Zin preferences in a consumption-savings framework with idiosyncratic risk include Cooper and Zhu (2016), Ampudia et al. (2018), Campanale and Sartarelli (2018) and Kaplan and Violante (2014). They agree about the intertemporal elasticity of substitution. We chose the risk aversion from the lower end of the range of their estimates, a conservative choice closer to more standard CRRA preferences.

[19] We target the difference between the standard deviation of errors in the second decile group versus the average over deciles 6-10.

deviation of errors across high wealth groups. Under our calibration, the maximum possible effort $\bar{n}$ reduces the standard deviation of noise to 0.5pp, half the standard deviation of shocks to the inflation rate. Exerting effort $\bar{n}$ comes at a cost of less than 0.1% of average income, speaking to the fact that little is necessary to deter households from acquiring information about future inflation. As in Castañeda et al. (2003), the process for y is calibrated to generate the distribution of wealth. Similar to their results, one high earnings state with lower persistence is necessary to generate a long right tail of the wealth distribution. Table 1 summarizes our parameter choices.

Table 1: Dynamic Model – Calibration

	Parameter	Value	Target
intertemp. substitution	γ	1.5	Literature
risk aversion	α	8.0	Literature
time preference	β	0.9779	fraction of debtors
borrowing limit	$\bar{a}$	-7.5	total debt
income states	y	[0.45 1 8]	wealth distribution
income transition	Π_y	$\begin{bmatrix} 0.975 & 0.025 & 0 \\ 0.057 & 0.931 & 0.012 \\ 0 & 0.15 & 0.85 \end{bmatrix}$	wealth distribution
nominal interest rate	r^n	0.04	2% SS real rate
persistence inflation	ρ	0.5	Dutch data, 1988-2018
long-run mean inflation	μ	0.02	Dutch data, 1988-2018
std. inflation shocks	σ_e	0.01	Dutch data, 1988-2018
curvature cost of effort	ϕ	2	quadratic cost of effort
scale cost of effort	θ	0.0015	range flat std. errors
maximum std. of noise	χ	0.1	peak std. errors
upper bound on effort	$\bar{n}$	17.5	low std. errors

Table 2 presents the fit of our model with respect to the wealth distribution. The model performs well along this dimension. It only struggles to match the strong concentration of wealth at the top as well as the total amount of debt. We argue, that the failure to match the concentration at the top has negligible relevance for our results regarding expectations, since the model performs much better in matching the total fraction of wealth held by the flat part of the expectation distribution (wealth deciles 6-10 jointly). As expectation formation is similar within this range, not matching the correct distribution of wealth within the upper half of the wealth distribution will not have consequences for our results regarding expectation formation.[20] The failure to

[20]To match the top tail of the wealth distribution, the literature often introduces heterogeneity in time preferences (see e.g. (Krusell and Smith, 1998)). Introducing such a positive correlation between

match the total amount of debt arises from the difficulty to match jointly total debt as well as the fraction of debtors, which we share with many similar models.[21] The model does well with respect to the fraction of indebted households, a feature important to match the position of the peak in the expectations distribution. It falls short in fully matching the amount of net liabilities of indebted agents. This imprecision is slightly biasing the standard deviation of expectation errors and mean absolute errors in the first wealth decile upwards, as we will see below.

Table 2: Wealth Distribution

Decile	1	2	3	4	5	6	7	8	9	10
Data	-6.14%	-0.01%	0.35%	1.01%	2.14%	3.82%	6.01%	10.07%	18.81%	63.94%
Model	-2.64%	-0.88%	0.65%	2.24%	4.01%	6.14%	9.00%	13.63%	22.45%	45.39%

Data refers to net financial wealth in the DNB Household Survey (waves 2010-2018). Compared to simulated, model implied wealth distribution.

4.3 Endogenous Expectations along the Wealth Distribution

Figure 6 presents the model implied equivalent to our baseline empirical findings in Figure 1. The model matches well qualitatively and quantitatively the differences in both the standard deviation of errors across households and their mean absolute forecast errors along the wealth distribution: A peak in the second wealth decile, a flattening over wealth deciles 6-10 and the quantitative magnitude of the decline between deciles 2 and 6. The model captures qualitatively the untargeted decline in both the standard deviation and mean absolute error for the first wealth decile vis-à-vis the second. As in the data, the first decile consists of households with negative net wealth. The shortfall in reproducing the quantitative magnitude of this decline is due to the left tail of the wealth distribution in the model being less spread out compared to the left tail of the net wealth distribution in the data. Where the model is off by the largest margin quantitatively is the level of both the standard deviation and the mean absolute error. In the data, both curves are about one percentage point higher than in the model. Note, however, that in order to isolate the effect of the proposed mechanism we abstract entirely from any exogenous dispersion in beliefs such as e.g. fundamental disagreement about the

wealth levels and households' weight on future utility would only strengthen our results further as it would make high wealth households care even more about future inflation.

[21] To match both jointly we would need to introduce additional model features, such as e.g. a wedge between borrowing rates and the return on savings, from which we abstract here to keep the exposition as simple as possible.

long run mean μ or heterogeneous biases in the signal. The level of error dispersion is in line with the fraction attributed to our mechanism in Figure B.1 after controlling for disagreement in long-run means. Exogenously imposing additional sources of dispersion would likely shift the reported measures up and towards the data equivalent.

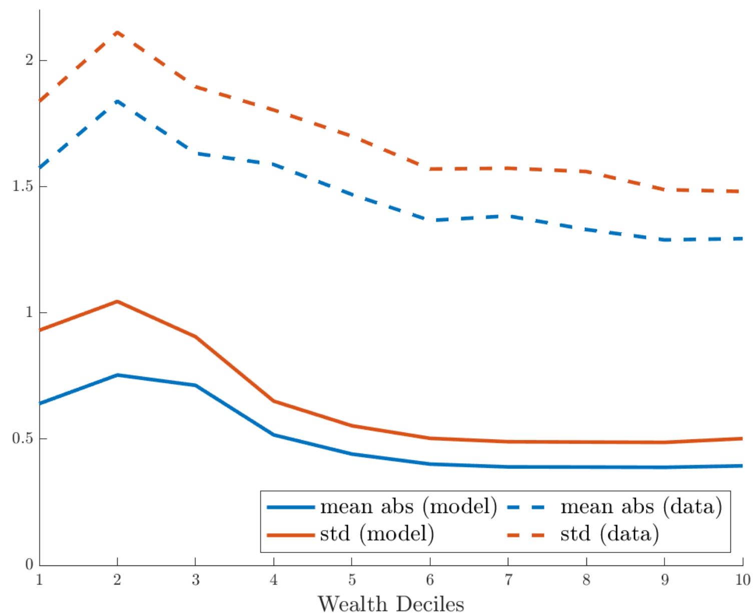

Figure 6: Expectation Errors by Wealth Decile Groups

Simulated, model implied statistics versus targeted data moments from Figure 1.

We have shown in Section 2 that, with our model of expectation formation, the driving force behind changes in mean absolute errors and standard deviation of errors is the noise in signals households receive about future inflation rates and hence the effort they choose to reduce this noise. Our quantitative findings suggest that, indeed, wealthier and indebted households endogenously choose to exert more such effort, enabling the model to replicate the empirical patterns. But why does the choice of effort vary with wealth? When inflation is a risk to the real interest rate, the more an agent wants to save or borrow between periods, the more he is exposed to fluctuations in the inflation rate. As future savings are positively correlated with current wealth, the richer (or the more indebted) a household is today, the more he will expose himself to inflation going forward. This exposure drives the incentives of households to exert effort and reduce the perceived uncertainty about future inflation.[22]

[22]We provide a more detailed discussion of the exposure effect in a two period framework in Appendix C.

Discussion of Assumptions

Before we highlight potential consequences of endogenous expectations along the wealth distribution by studying households' consumption responses to signals about future inflation, we revisit four key assumptions underlying our results:

First is our choice of preferences. The qualitative finding of effort choices increasing in absolute wealth levels does not rely on the assumption of recursive preferences, it pertains also under more standard CRRA utility. A sufficient level of risk aversion is, however, important to quantitatively generate a steep decline in the standard deviation of errors as it leads to a stronger increase of the gains from effort with wealth. Epstein-Zin preferences allow for high risk aversion without marginalizing the intertemporal elasticity of substitution, which is important for our analysis of consumption responses to signals below.

Second is our empirical measure of wealth. What ultimately matters for the formation of households' inflation expectations in the model are beginning of period resources $\frac{1+r^n}{1+\pi}a + y$. These determine the *potential* exposure to inflation until the next period as they pin down the general range of future savings/borrowing. The *actual* exposure will then be given by the realized savings/borrowing choice within this range, but this happens only after the household has formed his expectations and is therefore endogenous to his effort choice. Motivated by the states relevant to households' expectation formation in the model, beginning of period wealth is the model-consistent empirical measure to consider.

Third, we have abstracted from modeling portfolio composition. In this regard, it is important to distinguish our analysis from work on the distributional consequences of surprises in inflation or monetary policy more general as e.g. in Doepke and Schneider (2006), Auclert (2019), and Tzamourani (2019). These papers focus on the *ex-post* distributional consequences of inflationary shocks, while we are concerned with the *ex-ante* anticipation of such shocks. In theory, households' exposure to future inflation is independent of the composition of beginning of period wealth as long as this composition is adjustable going forward. We argue that this is the case for financial wealth at annual frequency, the time horizon at which we have data and to which we calibrate the model. Households' balance sheets going forward are endogenous to their expectation formation. Including a portfolio choice into the model is likely to only strengthen results as the benefits from information in the presence of portfolio choice are increasing in wealth, shown e.g. in Peress (2004). His results suggest that when aggregate risk is distorting the relative returns of different assets, households with larger portfolios can gain more from acquiring information and rebalance their asset holdings optimally. Therefore, if inflation is distorting relative asset returns, again richer households would have higher incentives to form precise expectations.

Fourth, we have also abstracted from any exposure of non-asset (labour) income to inflation risk. Our results rely on this assumption to the extent that the exposure of labour income to inflation has to be sufficiently below the exposure of asset income. "Sufficiency" is determined by the levels of absolute risk aversion along the wealth distribution. What is important for our findings is that the residual absolute exposure to inflation, the exposure households face after controlling for all indexation of wages and asset returns to inflation, increases enough along the wealth distribution to outweigh the decrease in absolute risk aversion.[23] The high collective bargaining coverage along with the low portfolio share of potentially indexed assets or debt in our sample provide evidence for a sufficient difference in residual absolute exposure.

5 Expectations and Consumption Responses

To conclude the analysis, we turn to households' consumption responses to a signal about future inflation and how these depend on their wealth levels. Aggregating the individual responses yields the on-impact response of aggregate consumption to forward guidance policies, which we discuss in the final part of this section.

5.1 The Marginal Propensity to Consume upon Signal

The starting point to trace out aggregate effects of endogenous expectation formation is the relationship between wealth and households' marginal propensity to consume in response to a signal about future inflation rates. We will refer to this metric as *MPCS* and define it as the relative change in a household's consumption policy in response to a change in the signal he receives about tomorrows shock to inflation $\hat{\epsilon}$, when holding all other variables (a,y,π,n) constant. Defined in this way, the MPCS is the semi-elasticity of a household's consumption policy with respect to the signal he receives. This measure

[23] For an extended discussion of the interplay between absolute risk aversion and exposure in a two period example see Appendix C.

has two components,

$$\text{MPCS} = \underbrace{\frac{1}{c}\frac{\partial c}{\partial\,\mathbb{E}[\epsilon]}}_{\text{MPCE}} \times \left.\frac{\partial\,\mathbb{E}[\epsilon]}{\partial\hat{\epsilon}}\right|_n . \tag{14}$$

We will refer to the first term as the marginal propensity to consume in response to a change in expectations (MPCE). It is the percentage change in current consumption in response to a change in expectations about future inflation rates, i.e. the semi-elasticity of consumption with respect to expected inflation. The second term captures the change in expectations in response to a change in the signal, where $\mathbb{E}[\epsilon]$ stands for households' full subjective distribution over the future shock. Since we are concerned here only with a change in its value but not the fact that a household receives a signal and further under assumptions as above, the only moment of the subjective distribution affected by the realization of the signal is the posterior mean.[24] Applying Bayesian updating as before, in our framework the change in the subjective mean is given explicitly as

$$\left.\frac{\partial\bar{\epsilon}}{\partial\hat{\epsilon}}\right|_n = \omega(n)d\hat{\epsilon} = \frac{\sigma_e^2}{\sigma_e^2 + \sigma_s^2(n)}d\hat{\epsilon}. \tag{15}$$

It becomes clear immediately how the response of expectations to a signal depends on effort n: The more effort is exerted, i.e. the less noisy a signal is perceived to be, the more a household will respond to this signal by updating his expected mean of future inflation – a standard result of Bayesian updating.

We analyze households' MPCS' quantitatively and compute the change in current consumption for each household if he receives a signal of $\hat{\epsilon} = 0.01$ instead of $\hat{\epsilon} = 0$. We distinguish four different scenarios defined by how noisy they perceive the signal to be. An *endogenous* scenario follows our baseline model where noise is determined by the endogenous choice of effort and heterogeneous across households. To disentangle households' MPCEs from how their expectations respond to a signal, we compare this benchmark to three scenarios in which noise is equalized across households: An *inattentive* scenario, setting the perceived noise of all households equal to that of the endogenously least informed. An *attentive* scenario, assigning to all agents the noise of

[24]The standard deviation of the posterior distribution responds only to the fact that a signal is received and to the noise attached to such signal but is independent of the value the signal takes. Under our assumptions on how households' form their expectations, the change in mean and standard deviation are sufficient to characterize the response of the entire distribution.

the endogenously most informed households. A *flat* scenario, in which the noise of all agents is chosen in order to match the unconditional standard deviation of errors in our baseline economy. All three have in common that the second term in (14) is constant across households and forces them to update their expectations in response to the signal in the same way, isolating differences in their MPCE. In the endogenous scenario, we also let $\left.\frac{\partial \mathbb{E}[\epsilon]}{\partial \hat{\epsilon}}\right|_n$ vary according to the endogenous effort choices of households.

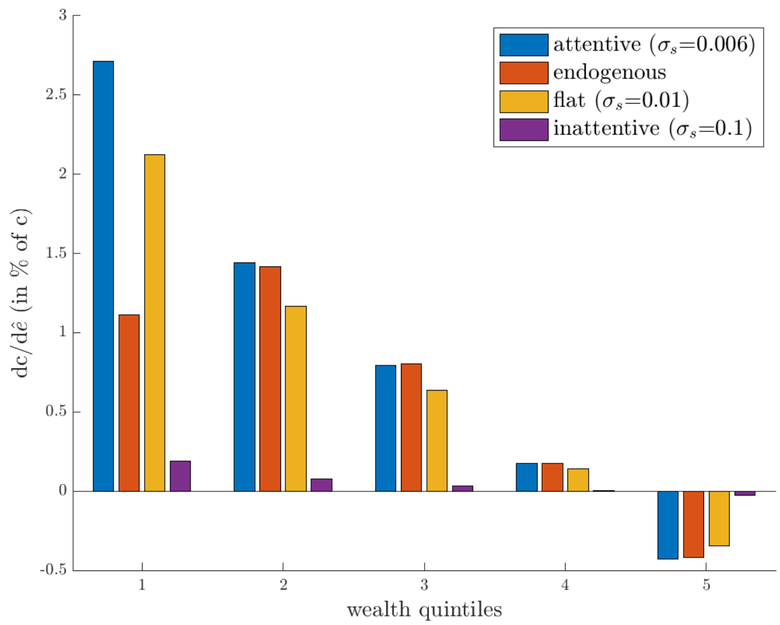

Figure 7: Marginal Propensity to Consume on Signal

Percentage change in consumption (aggregated by wealth quintile) on impact if $\pi = 2$ and $\hat{e}$ changes from 0 to 1pp. Endogenous: Noise as endogenously chosen. Attentive: All HHs σ_s=0.006. Inattentive: All HHs σ_s=0.1. Flat: All HHs σ_s=0.01.

Figure 7 plots the MPCS' aggregated by quintile of the wealth distribution.[25] We begin by looking at the three cases in which we keep $\left.\frac{\partial \mathbb{E}[\epsilon]}{\partial \hat{\epsilon}}\right|_n$ constant across households. The figure shows the MPCS for the inattentive, attentive and flat scenarios to be decreasing in wealth. Remember that from equation (14) MPCS = MPCE $\times \left.\frac{\partial \mathbb{E}[\epsilon]}{\partial \hat{\epsilon}}\right|_n$. Following this decomposition, it has to be the MPCE that is declining in wealth. This is due to the interaction of income and substitution effects in expectation of future inflation rates. For a household who previously would not have held any savings or debt between periods, a change in expected inflation comes down to a change in the expected

[25]To aggregate, we use the stationary wealth distribution of the converged economy if inflation is constant at two percent.

relative price of consumption today versus tomorrow, generating a substitution effect on current consumption. For a household who initially would have held either savings or debt, the substitution effect is accompanied by an income effect as a change in expected inflation implies a change in expected real financial income in the future. This income effect counteracts the substitution effect for saving households while it reinforces the substitution effect for borrowing households. For the case considered in Figure 7, a change in the signal from zero to one percentage point reveals to saving households that they will tomorrow be poorer than previously expected, hence diminishing their consumption response compared to households with little savings or debt. A good predictor for future savings in the model is the current asset level of a household, implying a MPCE declining in wealth.

The difference in the magnitude of MPCS' between the three cases with constant noise across households is driven by how much they update their expectations in response to the signal. The least informed ("inattentive") households choose a standard deviation of noise (σ_s) as high as 0.1 compared to a standard deviation of 0.01 of the actual shock (σ_e). Therefore, they attach little weight to any signal they receive ($\omega^{inatt} \approx 0.01$), do not update their beliefs in response and hence do not change their consumption behavior. This is why the MPCS for inattentive households remains low. For the flat scenario, σ_s decreases to 0.01 and hence $\omega^{flat} \approx 0.5$. For the attentive scenario we assume the standard deviation of the noise to be 0.006. Therefore, they attach more weight to any signal they receive ($\omega^{att} \approx 0.74$) and respond stronger in terms of consumption. The increase in the MPCS is not linear in ω across scenarios since a change in effort also affects household's uncertainty, i.e. the standard deviation of their inflation expectations, and hence their precautionary saving motive.

In the endogenous scenario, both terms in equation (14) interact. From our analysis so far we know the MPCE to be decreasing in wealth. From section 4.2 we know households effort choice and hence $\left.\frac{\partial \mathbb{E}[\epsilon]}{\partial \hat{\epsilon}}\right|_n$ to be increasing in wealth. The interaction between these two forces yields a hump shaped pattern of MPCS' along the wealth distribution. At the lowest wealth levels the increase in effort following an increase of resources outweighs the decline in MPCEs. From the second quintile onwards the decline in MPCEs dominates as effort is almost constant over the upper half of the wealth distribution. The figure shows that at low levels of wealth endogenous effort leads to a substantially lower consumption response compared to the counterfactual attentive scenario. This gap is how the influence of wealth on expectation formation has an impact on macroeconomic aggregates.

5.2 A Forward Guidance Exercise

Campbell et al. (2012) famously coined the terms of *odyssean* and *delphic* forward guidance, the former referring to policy makers commitment to some future policy action and the latter standing in for an attempt to influence expectations about the future path of economic variables. Our model naturally lends itself to a discussion of the channel behind delphic forward guidance, as it provides an understanding into how heterogeneous households respond to signals about future inflation rates. More specifically, we can provide an approximation to how much endogenous expectation formation can decrease the effectiveness of such forward guidance policies. While a full general equilibrium analysis is beyond the scope of our setup, we will be able to capture the initial consumption response to a change in households' inflation expectations. Following Auclert and Rognlie (2020), any demand shock can be decomposed into a partial equilibrium consumption response on impact and a general equilibrium multiplier. Our results should be interpreted as capturing the initial partial equilibrium increase in aggregate demand which is then amplified through a general equilibrium multiplier.

To highlight households' response to delphic forward guidance we conduct a quantitative exercise: Assume the economy to be stationary at $\pi = 0.02$. In this economy we shift the signal of every household by 0.01, such that all signals are drawn from $\mathcal{N}(0.01, \sigma_s^2(n_t^i))$ instead of $\mathcal{N}(0, \sigma_s^2(n_t^i))$. For each household we compute the change in consumption compared to the original signal and obtain an aggregate response using the stationary distribution of households. We do so under two different assumptions about how noisy households perceive their signals to be: Our benchmark scenario, where households' choose their effort endogenously, as well as the counterfactual attentive scenario, where all households' are as informed as the most informed inside the model economy. The attentive scenario provides an upper bound on how effective forward guidance could be as it assumes all households to attach the highest possible weight to any signal received and hence a strong updating of expectations. The difference between the two scenarios provides us an estimate for the potential consumption response that forward guidance misses out on due to some households not paying attention to inflation.

Table 3: Forward Guidance Exercise

calibration	**attentive**	**endogenous**	**missing potential**
baseline	0.20	0.09	0.11 (55%)
adjusted	0.13	0.08	0.05 (42%)

The table reports aggregated MPCS' in pp as defined in (14) if signals are drawn from $\mathcal{N}(0.01, \sigma_s^2(n_t^i))$ instead of $\mathcal{N}(0, \sigma_s^2(n_t^i))$ when the economy is stationary at $\pi = 0.02$. The first row reports results for our baseline calibration, the second row for an alternative calibration with $\bar{n} = 10$.

The first row in table 3 presents results for our baseline calibration. It shows that due to endogenous expectation formation forward guidance loses approximately 55% of its consumption response on impact, a sizable drop in the partial equilibrium response necessary to trigger any general equilibrium effects. As outlined in the previous section and especially Figure 7, the missing potential lies with households around zero net wealth who exert little effort in forming precise expectations, perceive any signal about future inflation as noisy, and hence do not update their expectations despite having the largest potential consumption response if they would do so. Any higher order (general equilibrium) effects that rely on this initial trigger will also be attenuated. Reaching those households' to whom higher future inflation does not imply a decrease in future income from asset holdings could therefore substantially increase the effectiveness of delphic forward guidance policies. Central banks should take this into account when designing the communication of their policies.

It is important to set this result in relation to previous work on the role of frictional expectation formation for the effectiveness of forward guidance policies, such as e.g. Wiederholt (2015) or Angeletos and Lian (2018). While most of this literature has focussed on the overall effect of imperfect expectation formation compared to a full information counterfactual, the result highlighted in this section is driven by differentials in expectation formation across households and how they are correlated with the general responsiveness to the policy announcement. This is also why our counterfactual is not a full information economy but one where we eliminate differences in attention across households.

Our baseline calibration has attributed the entire decline in the standard deviation of errors along the wealth distribution to endogenous factors and, in this regard, provides an upper bound on the effect of differences in expectation formation across households on forward guidance.[26] To test the robustness of our estimate to this assumption, we adjust the calibration to match the decline in subjective uncertainty after controlling for dispersion in beliefs about the long-run mean of inflation μ as presented in Appendix B.1. This provides some lower bound as it assumes any decline in dispersed beliefs about μ to be entirely exogenous, restricting the endogenous gap of attention between high and low wealth households. Instead of a decline of 0.57 between peak and low of the standard deviation of errors, we now target a drop of only 0.34. This target is met by adjusting $\bar{n}$ to 10 and keeping all other parameters as they were before. The second row of Table 3 presents results for this alternative calibration. While in general the response of consumption is weaker than under the baseline calibration due to the reduced attentiveness (and hence reduced updating of expectations upon a signal) of the most informed households, forward guidance still loses about 42% of its initial effect on consumption when moving from maximum attention of all households' to endogenous expectation formation. This is due to the first marginal reduction in noise increasing ω more than the last and the non-linear effects of ω on consumption responses due to precautionary saving behavior.

6 Concluding Remarks

This paper provides a framework to discuss the joint formation of households' inflation expectations and savings choices. We argue that wealth levels are important for both the formation of expectations and households' response to expected inflation. Looking at empirical observations from the DHS dataset, the standard deviation of forecast errors and mean absolute errors are declining in absolute wealth. We exploit changes in these cross-sectional statistics along the wealth distribution to discipline a consumption-savings problem with endogenous expectation formation, where households can exert effort to reduce uncertainty about future inflation rates. The model matches the empirical observations. The mechanism behind this finding works through the heterogeneous exposure to inflation that households at different points in the wealth distribution face. The model allows us to back out marginal propensities to consume in response to signals about future inflation. These MPCS' are hump shaped in wealth, driven by a negative

[26] Along another dimension, the failure of the calibrated model to match the top of the wealth distribution dampens the consequences of the proposed mechanism. With higher inequality in wealth, as observed empirically, the effects of endogenous expectation formation would increase further. More wealth inequality implies larger dispersion in MPCEs and hence even more importance for who pays attention to future inflation rates.

correlation between households' consumption response to expected inflation and the change in their expectations in response to signals. At the aggregate level, small MPCS' of low wealth households (due to a lack of attention to inflation) can substantially reduce the effectiveness of forward guidance policies.

While an empirical analysis of MPCS' lies beyond the scope of this paper, others have conducted related work in the DHS dataset: Lieb and Schuffels (2019) find the likelihood of positive durable consumption expenditure in response to higher inflation expectations to be decreasing in wealth. Similarly, Coibion et al. (2019) report a stronger decline in durable consumption in response to (exogenously) higher inflation expectations for households with higher wealth levels. This can be seen as support for MPCEs declining in wealth due to the interaction of income and substitution effects. Unrelated to inflation but in line with our theory, Fuster et al. (2020) find more exposed participants to be willing to pay a higher cost for information about future house prices in an experiment. More work along these lines is necessary for a full empirical evaluation of our theory, especially with regard to the effect wealth has on how expectations respond to signals.

Our paper also leaves room for further theoretical work on the topic. One possible addition to the analysis presented here can be to include a portfolio choice into our model. As mentioned before, such an extension is unlikely to alter the findings presented in this paper. It might nevertheless yield interesting additional results on the implications of costly inflation expectations for wealth inequality, as suggested by the findings of Peress (2004) and Lei (2019). While we focus on uncertainty and endogenous expectations about the shock to inflation rates, the model can be extended to other sources of heterogeneity in expectations such as learning about the underlying model. Our extension to include fundamental disagreement provides a starting point for work in this direction. More importantly, a computationally demanding but interesting application of the mechanism described in this paper would be to introduce our model of expectation formation into a general equilibrium environment. Recent work by Carroll et al. (2020) and Auclert et al. (2020) has included imperfect expectations in general equilibrium models with heterogeneous households. These papers rely so far on exogenous updating of expectations. It would be important to understand the impact of our findings on MPCS' in their general equilibrium setting. We leave these extensions for future research.

References

Ampudia, Miguel, Russell Cooper, Julia Le Blanc, and Guozhong Zhu, "MPC Heterogeneity in Europe: Sources and Policy Implications," Working Paper 25082, National Bureau of Economic Research September 2018.

Angeletos, George-Marios and Chen Lian, "Forward Guidance without Common Knowledge," *American Economic Review*, 2018, *108* (9), 2477–2512.

Argente, David and Munseob Lee, "Cost of Living Inequality During the Great Recession," *Journal of the European Economic Association*, 2021, *19* (2), 913–952.

Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, "Inflation Expectations and Behavior: Do Survey Respondents act on their Beliefs?," *International Economic Review*, 2015, *56* (2), 505–536.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, *109* (6), 2333–67.

_ and Matthew Rognlie, "Inequality and Aggregate Demand," 2020.

_ , _ , and Ludwig Straub, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," Working Paper 26647, National Bureau of Economic Research January 2020.

Ben-David, Itzhak, Elyas Fermand, Camelia M. Kuhnen, and Geng Li, "Expectations Uncertainty and Household Economic Behavior," Working Paper 25336, National Bureau of Economic Research December 2018.

Broer, Tobias, Alexandre Kohlhas, Kurt Mitman, and Kathrin Schlafmann, "Heterogenous Information Choice in General Equilibrium," 2018.

Campanale, Claudio and Marcello Sartarelli, "Life-cycle Wealth Accumulation and Consumption Insurance," 2018.

Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 2012, pp. 1–80.

Carroll, Christopher D., Edmund Crawley, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "Sticky Expectations and Consumption Dynamics," *American Economic Journal: Macroeconomics*, 2020, *12* (3), 40–76.

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull, "Accounting for the US Earnings and Wealth Inequality," *Journal of Political Economy*, 2003, *111* (4), 818–857.

Coibion, Olivier and Yuriy Gorodnichenko, "What can Survey Forecasts tell us about Information Rigidities?," *Journal of Political Economy*, 2012, *120* (1), 116–159.

_ , Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten van Rooij, "How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial," Working Paper 26106, National Bureau of Economic Research July 2019.

Cooper, Russell and Guozhong Zhu, "Household finance over the life-cycle: What does education contribute?," *Review of Economic Dynamics*, 2016, *20*, 63–89.

Crump, Richard K., Stefano Eusepi, Andrea Tambalotti, and Giorgio Topa, "Subjective intertemporal substitution," *Federal Reserve Bank of New York Staff Report 734*, 2015.

Deuflhard, Florian, Dimitris Georgarakos, and Roman Inderst, "Financial Literacy and Savings Account Returns," *Journal of the European Economic Association*, 2019, *17* (1), 131–164.

Doepke, Matthias and Martin Schneider, "Inflation and the redistribution of nominal wealth," *Journal of Political Economy*, 2006, *114* (6), 1069–1097.

Dräger, Lena and Giang Nghiem, "Are Consumers Spending Decisions in Line With an Euler Equation?," 2018.

Epstein, Larry G. and Stanley E. Zin, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 1989, *57* (4), 937.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar, "Expectations with Endogenous Information Acquisition: An Experimental Investigation," *The Review of Economics and Statistics*, 2020, pp. 1–54.

Jaravel, Xavier, "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector," *The Quarterly Journal of Economics*, 2019, *134* (2), 715–783.

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ **and Giovanni Violante**, "The Marginal Propensity to Consume in Heterogeneous Agents Models," Technical Report 2021.

_ **and Sam Schulhofer-Wohl**, "Inflation at the Household Level," *Journal of Monetary Economics*, 2017, *91*, 19–38.

Krusell, Per and Anthony A. Smith Jr., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Lee, David S, "Training, wages, and sample selection: Estimating sharp bounds on treatment effects," *The Review of Economic Studies*, 2009, *76* (3), 1071–1102.

Lei, Xiaowen, "Information and Inequality," *Journal of Economic Theory*, 2019, *184*.

Lieb, Lenard and Johannes Schuffels, "Inflation Expectations and Consumer Spending: The Role of Household Balance Sheets," 2019.

Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, "Survey: Rational inattention, A disciplined behavioral model," 2018.

Madeira, Carlos and Basit Zafar, "Heterogeneous Inflation Expectations and Learning," *Journal of Money, Credit and Banking*, 2015, *47* (5), 867–896.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson, "The power of forward guidance revisited," *American Economic Review*, 2016, *106* (10), 3133–58.

Peress, Joël, "Wealth, Information Acquisition, and Portfolio Choice," *The Review of Financial Studies*, 2004, *17* (3), 879–914.

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

Tzamourani, Panagiota, "The Interest Rate Exposure of Euro Area Households," 2019.

Vellekoop, Nathanael and Mirko Wiederholt, "Inflation Expectations and Choices of Households," 2019.

Wiederholt, Mirko, "Empirical Properties of Inflation Expectations and the Zero Lower Bound," 2015.

Appendix

A Empirical Observations

A.1 Michigan Survey of Consumers

The Michigan Survey of Consumers (MSC) is one of the most established sources of data on households' expectations. Compared to our main data source it has a disadvantage in that it does not provide comprehensive data on the wealth of participants. It only reports the current value of individuals' stock market portfolios. We use this value as a proxy for financial wealth and repeat part of the analysis on DHS data for the Michigan Survey.

An advantage of the MSC is the long time series for which consistent data are available. Data on inflation expectations and stock investment are continuously provided since September 1998. Furthermore, the data is available at monthly frequency. This does not only increase the number of observations along the time dimension, but also allows for a more precise computation of the forecast error as we can pin down the exact month of the observation. Applying the same approach as discussed above for the DHS data, we assign observations to investment quintile groups based on their position in the stock portfolio distribution in the month of their observation. We compute the expectation error as the reported forecast minus the realized inflation rate in the 12 months following the month of observation.

Figure A.1 reports the within quintile group standard deviation of expectation errors as well as the mean absolute forecast error by quintile group. Similar to the DHS data both are declining in investment value, a pattern that is statistically significant. Note that the first quintile now begins at zero investment as naturally there are no observations reporting a negative value of their stock market portfolio. Hence, we cannot observe any drop for negative wealth levels. Interestingly, we also cannot observe a flattening out of the decline for high levels of stock investment. Again, the pattern is robust to controlling for age or education. Figure A.2 shows that the standard deviation of errors split by education and age groups. As in the Dutch data college education reduces disagreement about future inflation rates. The findings are similar for mean absolute errors, as presented in Figure A.3.

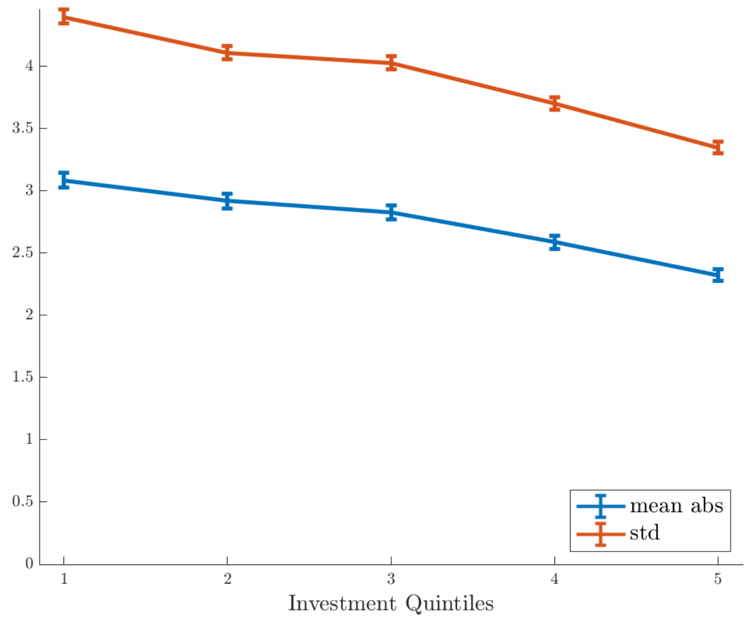

Figure A.1: Expectation Errors by Investment Quintiles (Michigan Data)

The figure plots the within quintile group standard deviation of errors and the mean absolute forecast error. Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from Michigan Survey of Consumers waves 09/1998-04/2018.

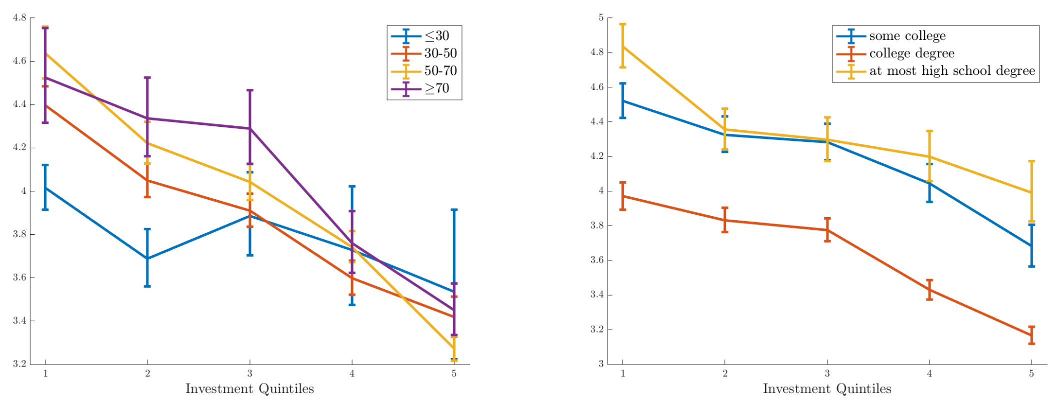

(a) by age

(b) by education

Figure A.2: Standard Deviation of Expectation Errors by Investment Quintiles (Michigan Data) – Controls

The figure plots the within quintile group standard deviation of errors by age (a) and education groups (b). Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from Michigan Survey of Consumers waves 09/1998-04/2018.

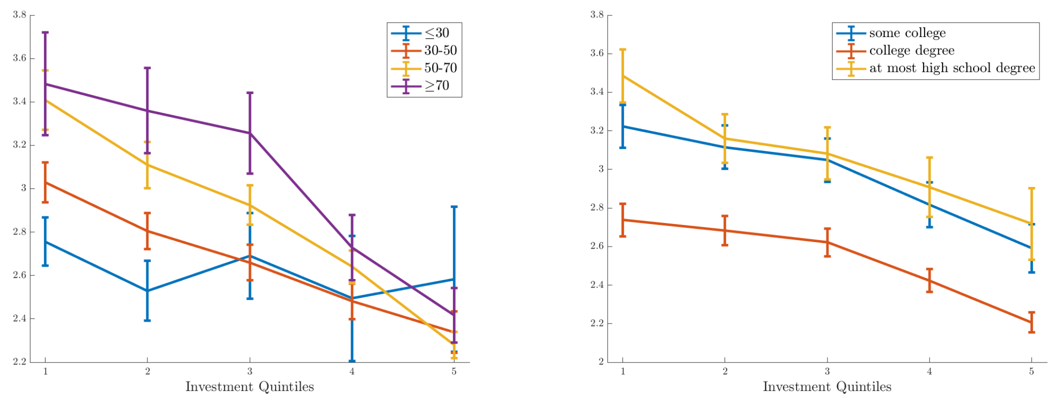

(a) by age (b) by education

Figure A.3: Mean Absolute Expectation Error by
Investment Quintiles (Michigan Data) – Controls

The figure plots the mean absolute forecast error for each investment quintile group by age (a) and education groups (b). Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from Michigan Survey of Consumers waves 09/1998-04/2018.

A.2 Additional Empirical Results

Table A.1: Net Financial Wealth Decile Groups – Summary

	Decile	N	Mean Assets	Forecast Errors Mean	Sd	Mean abs	N missing
deciles by wave	1	1,272	-27,513	1.18	1.84	1.58	149
	2	1,267	-54	1.47	2.11	1.84	170
	3	1,262	1,599	1.27	1.90	1.63	128
	4	1,264	4,574	1.24	1.80	1.59	123
	5	1,261	9,670	1.04	1.70	1.47	122
	6	1,270	17,134	0.99	1.57	1.37	110
	7	1,260	27,175	1.02	1.57	1.39	104
	8	1,265	45,373	0.97	1.56	1.33	84
	9	1,264	84,787	0.86	1.49	1.29	66
	10	1,260	289,130	0.91	1.48	1.30	53
deciles pooled sample	1	1,268	-27,640	1.17	1.83	1.57	142
	2	1,261	-33	1.47	2.10	1.83	172
	3	1,265	1,519	1.30	1.92	1.67	121
	4	1,264	4,497	1.25	1.80	1.60	130
	5	1,266	9,469	1.10	1.74	1.51	118
	6	1,267	16,938	0.97	1.53	1.34	113
	7	1,261	27,115	0.96	1.55	1.36	98
	8	1,268	45,120	0.97	1.56	1.33	96
	9	1,261	84,871	0.85	1.50	1.27	64
	10	1,264	289,088	0.90	1.48	1.29	55
	Total	12,645	45,061	1.09	1.72	1.47	1,109

Data from DNB Household Survey waves 2010-2018. Summary statistics by net financial wealth decile groups. Net financial wealth refers to net wealth ex housing, mortgages, businesses and vehicles. The first block sorts households into deciles by year of observations and then pools deciles across waves. The second block pools all observations and computes deciles based on the full sample.

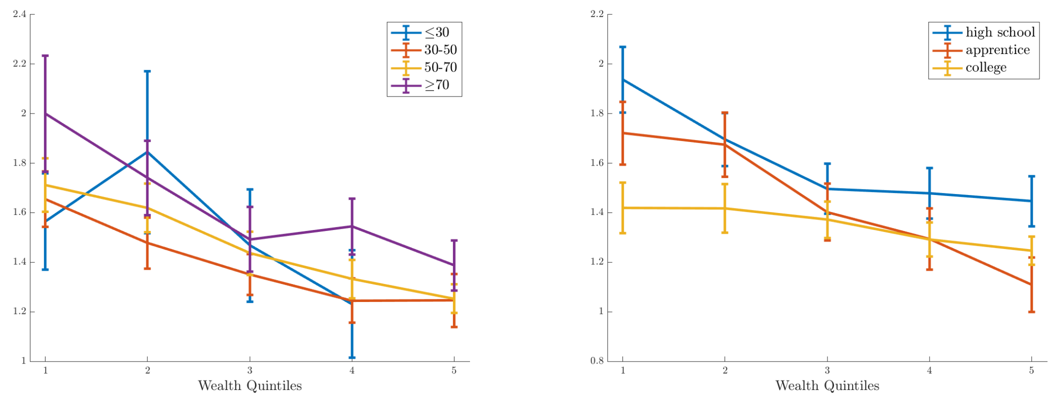

(a) by age

(b) by education

Figure A.4: Mean Absolute Expectation Error by Wealth Quintiles – Controls

The figure plots the mean absolute forecast error for each wealth quintile group by age (a) and education groups (b). Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018. Combination of youngest age and highest wealth quintile omitted due to lack of observations.

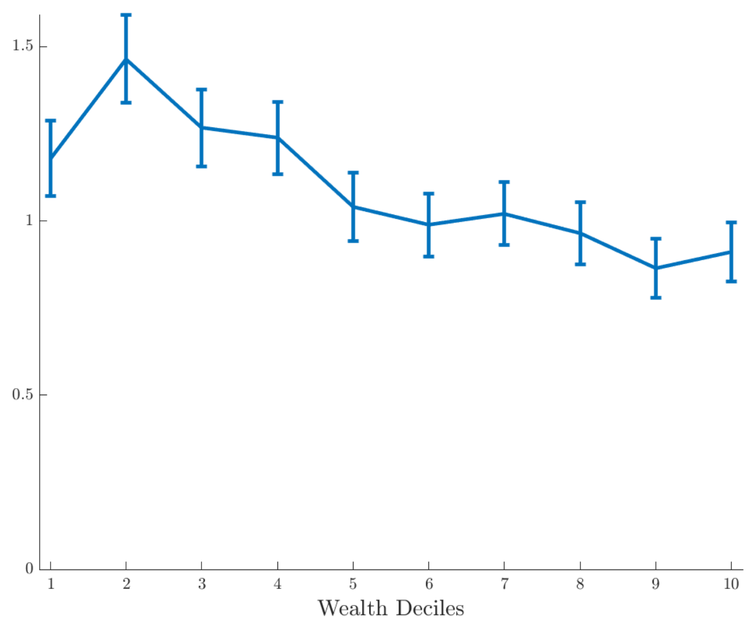

Figure A.5: Expectation Errors by Wealth Decile Groups – Mean

The figure plots the average expectation error by net financial wealth decile group. Bars provide confidence bands at the 95% level. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018.

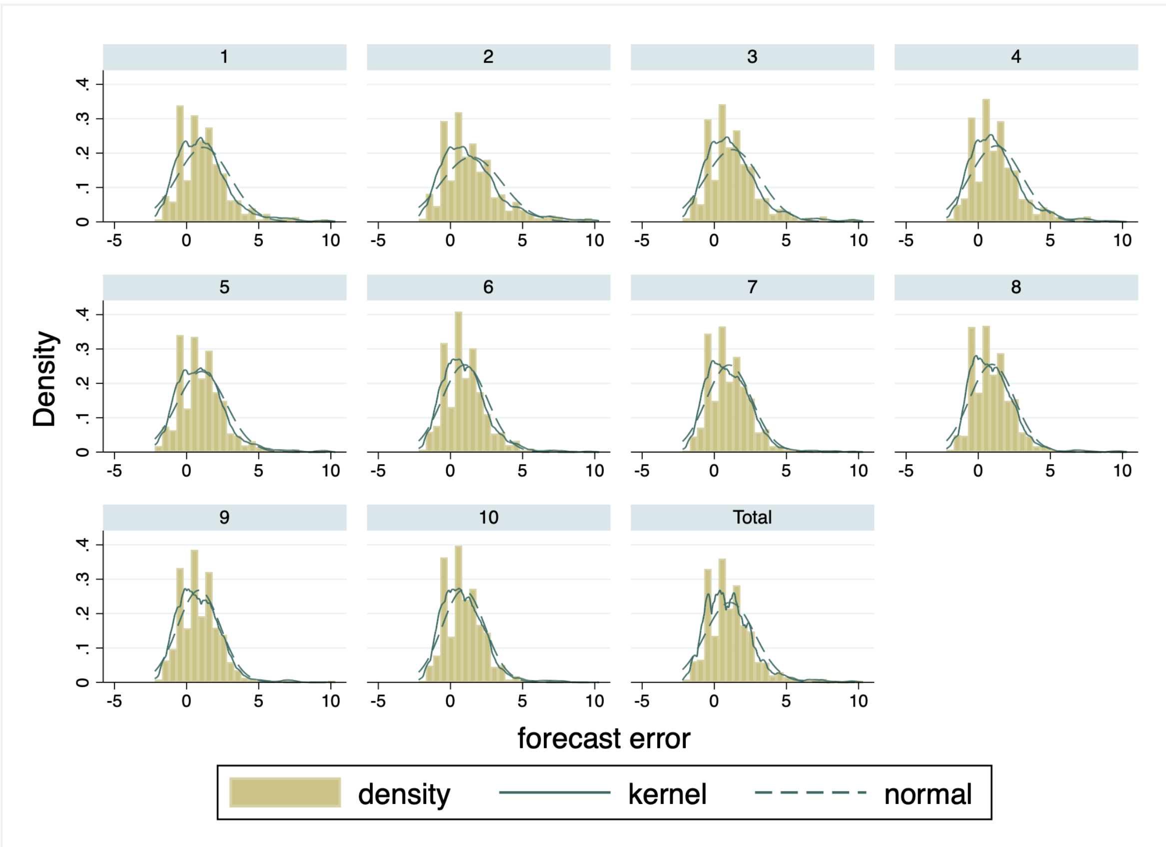

Figure A.6: Distribution of Errors by Wealth Decile Groups

The figure plots histograms of the distribution of expectation errors by wealth decile group along with fitted kernel densities and normal densities with identical mean and standard deviation as a reference point. Expectation errors are ex-ante point forecasts minus ex-post realizations. Data from DNB Household Survey waves 2010-2018.

Table A.2: Individual Absolute Forecast Errors

	$abs(err^i_{t+1})$
net financial wealth decile 2	0.208**
	(0.092)
net financial wealth decile 3	0.027
	(0.083)
net financial wealth decile 4	-0.021
	(0.079)
net financial wealth decile 5	-0.109
	(0.075)
net financial wealth decile 6	-0.205***
	(0.074)
net financial wealth decile 7	-0.195**
	(0.077)
net financial wealth decile 8	-0.229***
	(0.074)
net financial wealth decile 9	-0.272***
	(0.073)
net financial wealth decile 10	-0.253***
	(0.076)
high school	-0.012
	(0.130)
apprenticeship	-0.126
	(0.131)
college	-0.219*
	(0.127)
age 30-50	-0.089
	(0.083)
age 50-70	-0.062
	(0.084)
age >70	0.075
	(0.088)
home owner	-0.103**
	(0.045)
constant	1.817***
	(0.147)
Observations	11532
Adjusted R^2	0.0215

Data from DNB Household Survey waves 2010-2018. Household level regression of absolute forecast errors on households' wealth decile, education and age of the household head and an indicator for owning the primary residence. Standard errors (in parentheses) clustered at the household level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

B Theoretical Framework

B.1 Extension - Fundamental Disagreement

In our baseline model of expectation formation we abstract from fundamental disagreement about the underlying model of inflation and its parameters or any other exogenously imposed heterogeneity in beliefs to focus on endogenous expectation formation. Nevertheless, additional sources of disagreement among households can be important to capture moments of the data our baseline model fails to explain, such as e.g. a positive mean error or a positive error covariance.[27] We therefore extend our empirical analysis in order to evaluate the potential impact of other sources of heterogeneity in expectations on our findings.

To test for robustness towards including fundamental disagreement, we adjust our baseline model of expectation formation to incorporate heterogeneity in beliefs about the long run mean of inflation μ. Household i's belief about μ is denoted μ^i and assumed to be distributed normally among households. Furthermore, we assume $\mu^i \perp\!\!\!\perp s_t^i \ \forall i, t$. With all other notation as before, household i's inflation expectation and expectation error are now given as

$$\mathbb{E}_t^i[\pi_{t+1}|\hat{e}_{t+1}^i, n_t^i] = (1-\rho)\mu^i + \rho\pi_t + \omega_{t+1}^i(n_t^i)\hat{e}_{t+1}^i \tag{16}$$

$$\begin{aligned} err_{t+1}^i &= \mathbb{E}_t^i[\pi_{t+1}|\hat{e}_{t+1}^i, n_t^i] - \pi_{t+1} \\ &= (1-\rho)(\mu^i - \mu) + (\omega_{t+1}^i(n_t^i)s_{t+1}^i - (1-\omega_{t+1}^i(n_t^i))e_{t+1}). \end{aligned} \tag{17}$$

The error now includes an additional term accounting for households' misperception of the long run mean. Denote the average belief about the long term mean of a group g of households as $\bar{\mu}^g$ and its variance as σ_μ^{g2}. Assuming, as before, that households in group g exert the same effort $\bar{n}_t^g$, the variance of errors across households in group g and over time becomes

$$\begin{aligned} \mathrm{Var}^g(err_{t+1}^i) &= (1-\rho)^2\mathrm{Var}(\mu^i) + (\omega_{t+1}^g(\bar{n}_t^g))^2\sigma_s^2(\bar{n}_t^g) + (1-\omega_{t+1}^g(\bar{n}_t^g))^2\sigma_e^2 \\ &= (1-\rho)^2\sigma_\mu^{g2} + \frac{\sigma_e^2\sigma_s^2(\bar{n}_t^g)}{\sigma_e^2 + \sigma_s^2(\bar{n}_t^g)} = (1-\rho)^2\sigma_\mu^{g2} + \overline{SU}_{t+1}^{g2} \end{aligned} \tag{18}$$

[27] See Figure A.5 and B.1.

where now the endogenous subjective uncertainty term $\overline{SU}^{g2}_{t+1}$ is adjusted by the within-group fundamental disagreement about μ. Disagreement among households can hence be decomposed into disagreement about the long run mean and households' subjective uncertainty. We can also compute the covariance of the ex-post errors across time. This is given as

$$\text{Cov}^g(err^i_{t+1}, err^i_t) = (1-\rho)^2\,\mathbb{E}[(\mu^i - \mu)^2] - (1-\rho)^2(E[(\mu^i - \mu)])^2 = (1-\rho)^2\sigma^{g2}_{\mu}. \quad (19)$$

Together, (18) and (19) allow us to identify the endogenous component of error dispersion in the presence of fundamental disagreement from the difference between variance and covariance of forecast errors as

$$\overline{SU}^g_{t+1} = \sqrt{\text{Var}^g(err^i_{t+1}) - \text{Cov}^g(err^i_{t+1}, err^i_t)}. \quad (20)$$

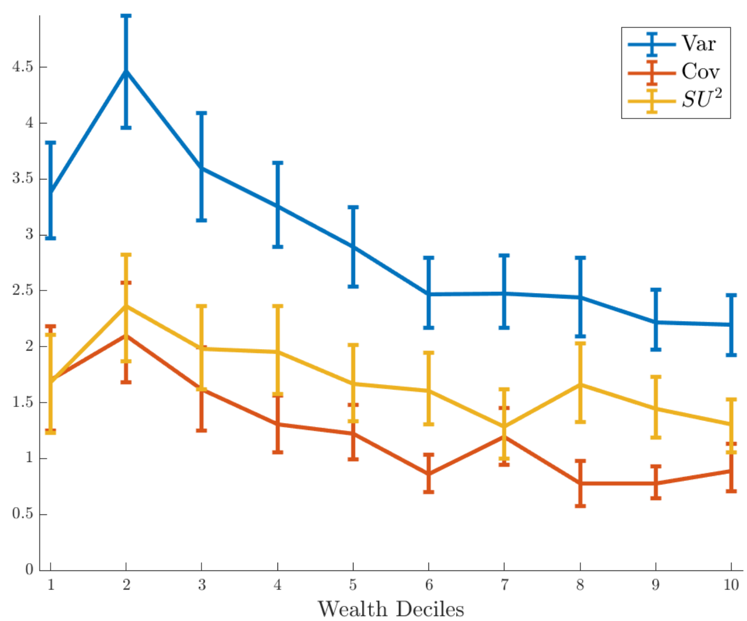

Figure B.1: Expectation Error Variance – Decomposition

The figure decomposes the variance of expectation errors across households (Var) by wealth decile groups into error covariance (Cov) and the square of subjective uncertainty (SU^2) as in (20). Data from DNB Household Survey waves 2010-2018. Bootstrapped 95% confidence intervals.

Intuitively, the covariance of errors is a sufficient statistic to measure heterogeneity in beliefs about the long run mean as we assume noise to be uncorrelated over time.

Persistent beliefs about misreporting in current inflation (household i assuming actual inflation $\tilde{\pi}_t^i = \pi_t + \bar{\pi}^i$), as well as dispersion in beliefs about the mean of the signal s or the shock e can be treated similarly as long as they are constant over time at the household level.

We apply equation (20) to our data and compute the implied subjective uncertainty of households by subtracting for each wealth decile group the error covariance over time from the within group variance. The result is presented in Figure B.1. The implied subjective uncertainty exhibits a similar pattern as our benchmark results. It is slightly increasing between the first and second decile group and broadly decreasing for further increases in wealth. The covariance, which according to the extended model is driven by the dispersion of beliefs about the long-run mean, is equally higher among households with lower wealth and decreasing alongside subjective uncertainty. Both the decreases in subjective uncertainty and covariance between their respective peaks and lowest points are significant at the 95% level . Of the overall drop in the variance of expectation errors across households, about half is attributable to the fall in exogenous disagreement about the long run mean (covariance) and half to a fall in endogenous subjective uncertainty.

Figure B.2 provides the decomposition of error variance across households by investment quintile group into covariance and subjective uncertainty as of equation (20) in the MSC. Even after allowing for fundamental disagreement almost all of the decline of within quintile group error variance is attributed to a decline in subjective uncertainty, while error covariance declines only modestly.

We take these findings as evidence that existence of the mechanism in our benchmark model is robust to incorporating fundamental disagreement.

B.2 Dynamic Budget Constraint - From Nominal to Real

Starting with nominal assets $\hat{a}$

$$
\begin{aligned}
Pc + \hat{a}' &= (1 + r^n)\hat{a} + y - P\mathcal{F}(n, i) \\
c + \frac{\hat{a}'}{P} &= (1 + r^n)\frac{\hat{a}}{P} + y - \mathcal{F}(n, i)
\end{aligned}
$$

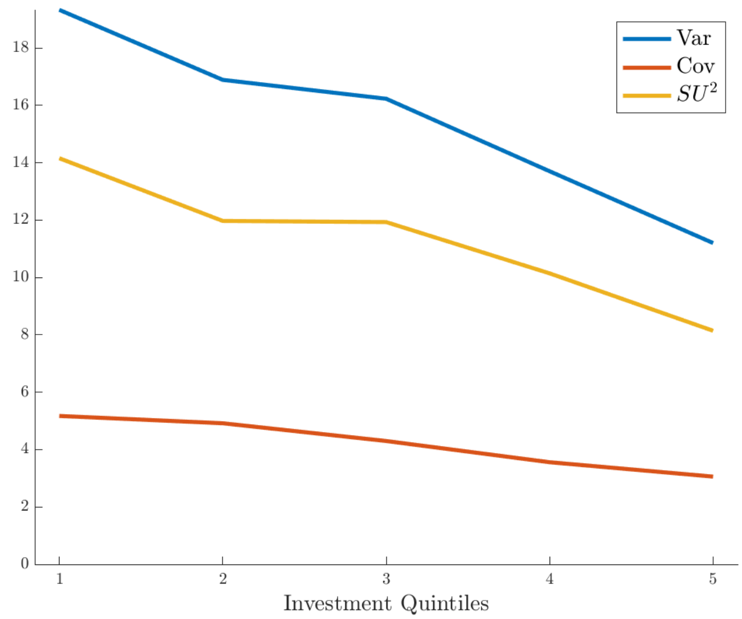

Figure B.2: Expectation Error Variance – Decomposition (Michigan Data)

The figure decomposes the cross-sectional variance of expectation errors (Var) by investment quintile groups into error covariance (Cov) and subjective uncertainty (SU) as in (20). Data from the Michigan Survey of Consumers waves 09/1998-04/2018.

Define $a' = \frac{\hat{a}'}{P}$, i.e. tomorrow's nominal assets in today's real consumption, and inflation rate $1 + \pi = \frac{P}{P_{-1}}$

$$c + a' = (1 + r^n)\frac{P_{-1}}{P}a + y - \mathcal{F}(n, i)$$
$$c + a' = \frac{1 + r^n}{1 + \pi}a + y - \mathcal{F}(n, i)$$

C Endogenous Expectations in a Two Period Model

To highlight the mechanism through which households' wealth levels impact their expectation formation, it is instructive to analyze the properties of a two period model. In the interest of a simpler exposition, we abstract from inflation entirely and focus directly on risk to the real interest rate. This is without loss of generality, since fluctuations in inflation translate into fluctuations in the real interest rate as long as nominal rates are not assumed to adjust one-for-one with inflation. Furthermore, their impact on real interest rates is the only channel through which fluctuations in inflation are relevant to the household's problem as long as additional (labour) income is assumed to be in real terms. These are the same assumptions we impose in the dynamic model where we consider inflation explicitly, making the two approaches comparable.

C.1 A Two Period Model

A household lives for two periods and maximizes utility by choosing consumption in both periods (c_1 and c_2) as well as savings a between periods. In both periods he receives a deterministic and constant income y. Additionally, at the beginning of the first period the household receives initial assets A. Preferences of the household are recursive, following Epstein and Zin (1989).

The real interest rate r between the two periods is stochastic. Before choosing savings in period 1, the household receives a noisy signal $\hat{r}$ about the interest rate. The distribution of the interest rate and the signal are given as

$$r \sim \mathcal{N}(\bar{r}, \sigma_r^2) \qquad \hat{r} = r + s \qquad s \sim \mathcal{N}(0, \sigma_s^2(n)), \tag{21}$$

where s is pure noise. Before receiving the signal, the household can influence the variance of the noise by exerting some effort n, for which he has to incur a monetary cost $\mathcal{F}(n)$. Based on the signal, the household forms a Bayesian posterior belief about the true interest rate r, attaching weight $\omega(n)$ to the signal received. Hence, conditional on n and $\hat{r}$, the posterior distribution is given as

$$\begin{aligned} r_{|n,\hat{r}} &\sim \mathcal{N}((1-\omega(n))\bar{r} + \omega(n)\hat{r}, \omega(n)\sigma_s^2(n)) \\ \omega(n) &= \frac{\sigma_r^2}{\sigma_r^2 + \sigma_s^2(n)}. \end{aligned} \tag{22}$$

We will refer to the standard deviation of a household's posterior belief about r (given by $\sqrt{\omega(n)\sigma_s^2(n)}$) as his subjective uncertainty about the future interest rate.

The household's effort choice problem is then given as

$$\tilde{V}(A) = \max_{n} \mathbb{E}_{\hat{r}}[V(A, n, \hat{r})|n]. \tag{23}$$

Conditional on having exerted effort n and receiving signal $\hat{r}$, the consumption-savings problem is given by

$$\begin{aligned} V(A,n,\hat{r}) &= \max_a \left(c_1^{1-\gamma} + \beta \left(\mathbb{E}_r[c_2^{1-\alpha}|\hat{r},n] \right)^{\frac{1-\gamma}{1-\alpha}} \right)^{\frac{1}{1-\gamma}} \\ c_1 &= A + y - a - \mathcal{F}(n) \\ c_2 &= (1+r)a + y \quad \forall\, r \end{aligned} \tag{24}$$

For the cost of effort and the relationship between effort and noise in the signal we assume functional forms

$$\sigma_s(n) = \frac{\chi}{1+n} \qquad \mathcal{F}(n) = (\theta n)^{\phi}. \tag{25}$$

These choices yield convex cost of and convex gains from exerting effort.[28] Note that with these functional forms χ is the variation in the noise if zero effort is exerted, i.e. the maximum variation possible, and that zero effort implies zero cost.

To highlight some properties of the proposed mechanism, we calibrate the model outlined above. The calibration is ad-hoc and for instructive purposes only. It is provided in Table C.1.

Table C.1: Two Period Model – Calibration

Parameter	Value
γ	2
α	2
β	0.98
y	4
$\bar{r}$	0.02
σ_r	0.01
ϕ	2
θ	0.005
χ	0.03

Calibration for the two period model. Values are ad-hoc and only for instructive purpose.

[28] $\sigma_s'(n) < 0,\ \sigma_s''(n) > 0$ and $\mathcal{F}'(n) > 0,\ \mathcal{F}''(n) \geq 0,$ iff $\phi \geq 1$.

C.2 Information Incentives

To study households' incentives to form precise expectations, we begin by taking the effort choice n as exogenously given. In order to do so, we drop the max-operator in (23) and set the cost in (24) to $\mathcal{F}(n) = 0 \; \forall \; n$. After solving the households' problem for given n we can compute a certainty equivalence consumption level cec_n, satisfying

$$\tilde{V}_n(A) = \left(cec_n^{1-\gamma} + \beta \left(cec_n^{1-\alpha} \right)^{\frac{1-\gamma}{1-\alpha}} \right)^{\frac{1}{1-\gamma}}, \tag{26}$$

where $\tilde{V}_n(A)$ is the value of (23) for exogenously given n and zero cost of effort. We use this certainty equivalent to construct a measure of the benefit of decreasing the noise in the signal as

$$\Delta cec_n = \frac{cec_n}{cec_0} - 1, \tag{27}$$

which is the percentage change in the certainty equivalence consumption level if effort is increased from 0 to n, and hence the standard deviation of the noise is decreased from χ to $\sigma_s(n)$.

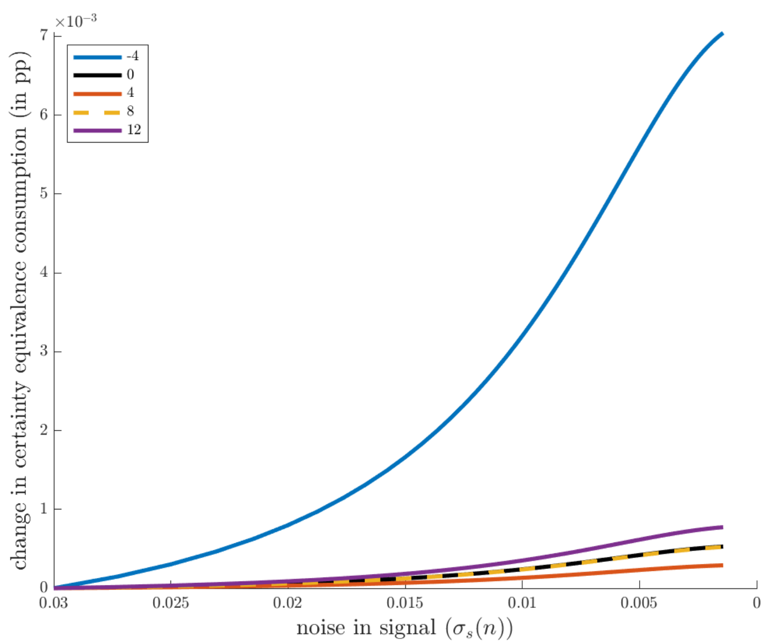

Figure C.1: Change in Certainty Equivalence Consumption

The figure plots the percentage gain in certainty equivalence consumption (cec_n, as defined in (26)) of decreasing the standard deviation of the noise in the signal from χ to $\sigma_s(n)$. Each line represents a different initial asset level A.

Figure C.1 plots results from the calibrated model for a range of initial asset values A. The gain from decreasing the variation in noise is highest for households starting with debt. It decreases as initial asset levels increase towards zero and modestly positive values for A and increases again once A becomes substantially positive. Note that the gains of decreasing the variation in the noise are small, for the given calibration below 0.01% of the certainty equivalent consumption level. This is evidence that already small cost of forming precise expectations might deter households from doing so.

The pattern of noise in wealth can be explained by two forces, governing households' incentives to form precise expectations: Exposure and absolute risk aversion. Exposure is given by the absolute value of a household's savings or borrowing between the two periods. It determines the relevance of the risk for a household. The higher absolute savings, the larger are expected fluctuations in period 2 consumption due to fluctuations in the interest rate. In the presence of risk aversion, fluctuations in future consumption reduce expected utility. Hence households with larger fluctuations in their future consumption due to the risk have stronger incentives to reduce the perceived risk and form more precise expectations. The exposure effect is therefore higher for households with either higher initial debt or higher (positive) initial assets, who engage in borrowing/saving between periods, but low for households with A close to zero, as these households save/borrow little between $t = 1$ and $t = 2$. Absolute risk aversion, as usual, implies that any absolute fluctuation in consumption has higher cost in terms of expected utility to households with a lower average consumption level. This effect is hence highest for households with higher debt (A substantially negative), as these households have the lowest consumption levels, and decreases as A increases.

To highlight the two effects on the change in certainty equivalence consumption, we conduct two quantitative experiments. For the first, we eliminate differences in the absolute risk aversion of households with different A to focus solely on exposure. This is achieved by compensating each household to obtain the same average consumption level as a benchmark household, which we chose to be a household with initial assets $A = -4$. More specifically, we fix the savings choice of a household at the optimal choice without any compensation. Conditional on the exogenously set effort n and the signal received

$\hat{r}$, each household receives a deterministic transfer for both periods, satisfying

$$\begin{aligned}\Delta c_1(A,n,\hat{r}) &= c_1(-4,n,\hat{r}) - c_1(A,n,\hat{r}) \\ \Delta c_2(A,n,\hat{r}) &= \underset{r}{\mathbb{E}}[c_2(-4,n,\hat{r})|n,\hat{r}] - \underset{r}{\mathbb{E}}[c_2(A,n,\hat{r})|n,\hat{r}].\end{aligned} \tag{28}$$

As this equalizes consumption levels across households, any difference in the remaining effect on the certainty equivalence consumption should be due to different exposure.

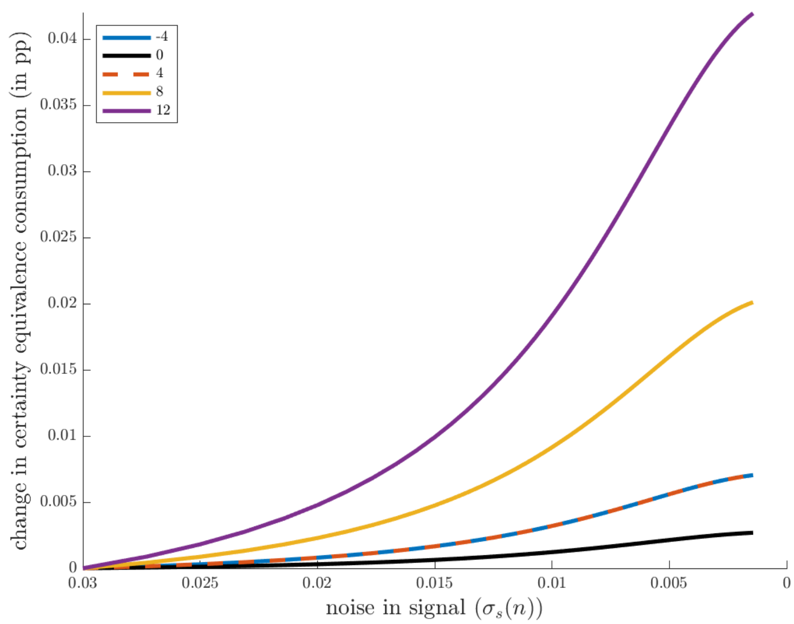

Figure C.2: Change in Certainty Equivalence Consumption – Exposure

The figure plots the adjusted percentage gain in certainty equivalence consumption (cec_n, as defined in (26)) of decreasing the standard deviation of the noise in the signal from χ to $\sigma_s(n)$. Adjustment equalizes average consumption levels across households to the level of a households with $A = -4$, as given in (28), while leaving the savings choice unchanged. Each line represents a different initial asset level A.

Figure C.2 plots the quantitative results. As expected, the change in the certainty equivalence consumption level is monotonically increasing in the absolute value of A, which is directly related to the absolute value of households' savings between periods. Note that the effect of decreasing the variation in noise is almost identical for households with $A = 4$ and $A = -4$. This reflects their, in absolute values and on average across signals, almost identical savings choices, implying a similar exposure to interest rate risk.

To control for the exposure effect and highlight the influence of absolute risk aversion, we can conduct a similar experiment by normalizing households savings choice. We assign every households the savings choice of a household with $A = 10$ (i.e. $s(10, n, \hat{r})$), controlling for n and $\hat{r}$. We additionally assign transfers, such that the household has the same average consumption level as before. These are given as

$$\begin{aligned} \tilde{\Delta} c_1(A, n, \hat{r}) &= s(10, n, \hat{r}) - s(A, n, \hat{r}) \\ \tilde{\Delta} c_2(A, n, \hat{r}) &= \mathop{\mathbb{E}}_r[c_2(A, n, \hat{r})|n, \hat{r}] - \mathop{\mathbb{E}}_r[c_2(10, n, \hat{r})|n, \hat{r}]. \end{aligned} \tag{29}$$

The results can be interpreted as the gain from decreasing the variation in noise for households with identical savings choice but varying consumption levels.

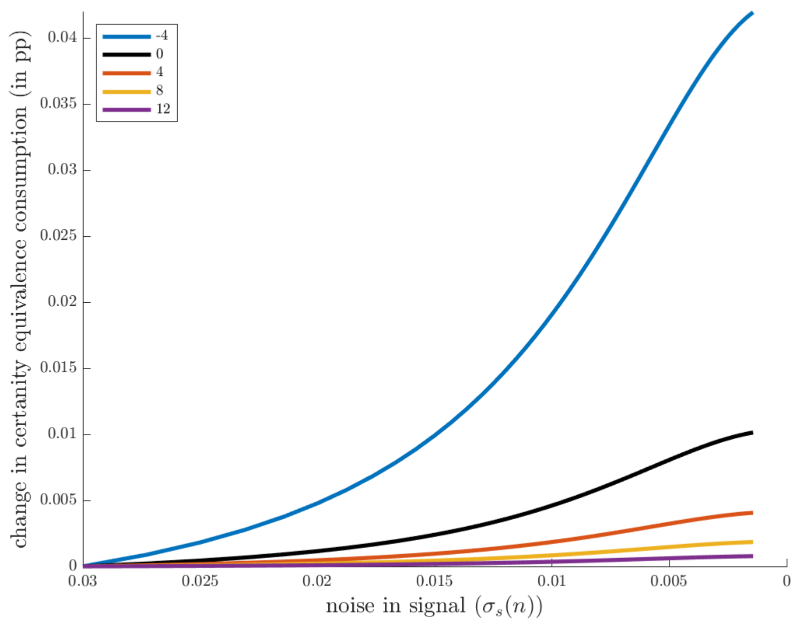

Figure C.3: Change in Certainty Equivalence Consumption – Absolute Risk Aversion

The figure plots the adjusted percentage change in certainty equivalence consumption (cec_n, as defined in (26)) of decreasing the standard deviation of the noise in the signal from χ to $\sigma_s(n)$. Adjustment equalizes savings across households to the level of a households with $A = 10$ as given in (29) while leaving the average consumption level of the household unchanged. Each line represents a different initial asset level A.

Figure C.3 plots the quantitative results. Unsurprisingly, when controlling for the savings choice, households with lower consumption level (and hence higher absolute risk aversion) profit more from a reduction of uncertainty. The gain from increasing n / reducing $\sigma_s(n)$ is decreasing in A.

C.3 Information Choice

We can summarize the findings above to make predictions about how households decide on effort n when the choice is endogenous. The exposure effect is increasing in households absolute initial wealth, as their future absolute savings will be equally increasing. This implies, that starting at a wealth level of zero, the further away we move in any direction along the wealth distribution the more effort households should want to exert due to the exposure effect. This effect is almost symmetric for positive and negative values of initial assets A. Absolute risk aversion is, however, monotonically decreasing in wealth. It reinforces the exposure effect, but more so for negative asset levels. The effect of absolute risk aversion is hence asymmetric in positive/negative wealth. We should hence expect the chosen noise in the signal to peak around zero wealth, decline as we move away from zero wealth in any direction, but decline steeper for negative wealth than for positive wealth. All discussion above assumes that effort is equally costly for all households. With the specification for effort to have monetary cost, this is not true in utility terms, as the same monetary costs transmit into higher utility cost for households with lower consumption levels. This adds an additional dimension of heterogeneity in incentives.

We confirm the predictions of our exercise by moving on to an endogenous choice of effort according to (23) and (24), subject to the cost function and return to effort as outlined in (25). The calibration remains the same as before. Figure C.4 plots the standard deviation of the noise implied by effort choice $n(A)$ across a range of initial asset level A. The findings confirm our conjecture from Section C.2. With increasing absolute wealth level (positive or negative), households decide to exert more effort to reduce the noise in the signal, driven by the exposure effect. Additionally, households with negative wealth choose to exert more effort (reduce the noise further) than households with similar positive wealth. This is due to the asymmetric impact of absolute risk aversion which has equally been discussed above.

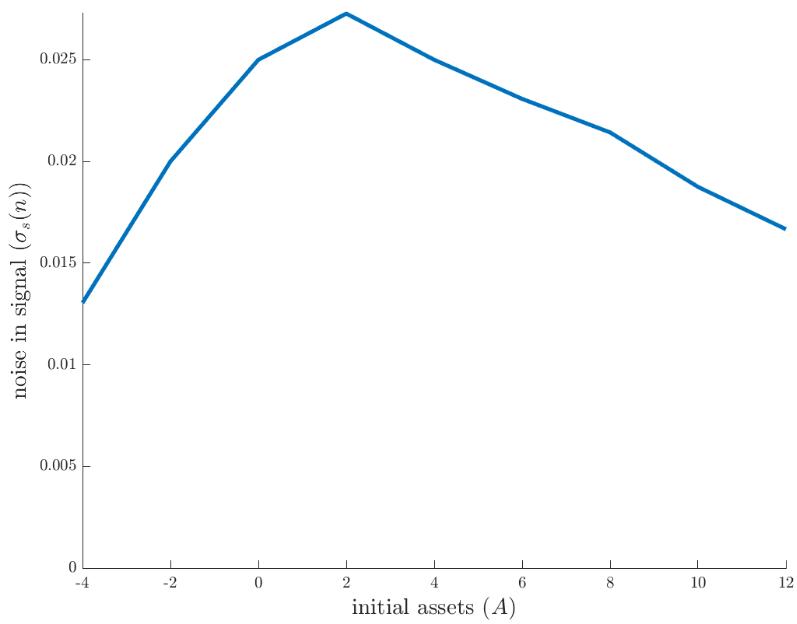

Figure C.4: Endogenous Effort – Chosen Standard Deviation of Noise

The figure plots the standard deviation of the noise implied by the endogenous choice for n, solving (23) for given initial asset level A.