Essays in Consumer Heterogeneity and Expectations

Loughlan O'Doherty

September 7, 2025

A thesis presented for the degree of Doctor of Philosophy

Department of Economics

London School of Economics and Political Science

United Kingdom

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science (LSE) is solely my own work other than where I have clearly indicated that it is the work of others.

This thesis consists of approximately 46,800 words.

Acknowledgements

I am grateful to my supervisor, Ricardo Reis, for his patience, support, and guidance. I would also like to thank my advisor Silvana Tenreyro for her kind counsel.

I greatly benefited from the exceptional community of researchers within the Centre for Macroeconomics at LSE. Conversations with Benjamin Moll, Joe Hazell, Alwyn Young, and Xavier Jaravel were always fruitful. Maarten de Ridder, in particular, always found ways to encourage me when I needed it the most. I also gratefully acknowledge the financial support of the Economic and Social Research Council (ESRC).

Special thanks should be extended to my colleagues and friends at LSE: Hugo Reichardt, Peter Lambert, Andres Fajardo-Ramirez, Yannick Schindler, Bernardo Mottironi, and Arnaud Dyèvre.

This thesis is dedicated to my partner Tatiana, and my mother Rosemary.

Abstract

This thesis comprises three essays that examine the origins, cross-section, and consequences of household forecast errors. Drawing on harmonised micro-data from the Michigan Survey of Consumers, the New York Fed Survey of Consumer Expectations, and the ECB Consumer Expectations Survey, the first essay uncovers a single latent "pessimism" factor that declines monotonically with income. Next, the chapter investigates forecast pessimism at a more granular level, and demonstrates that the effect of income is quantitatively large. Exploiting individual level income shifts, it shows that the relationship between income and forecast pessimism survives. Lower-income households systematically over-predict inflation, unemployment, and other macroeconomic outcomes.

The second essay investigates one explanation for forecast pessimism: ambiguity aversion. We ask: if ambiguity aversion drives forecast pessimism, what is the source of this ambiguity? The paper outlines two testable hypotheses and tests them. First, using survey-based Ellsberg experiments linked to expectation data, we find that more ambiguity averse agents have larger, positive, forecast errors, but that the elicited preferences covary with income themselves. Second, we find that the response of consumer expectations on receipt of a signal is asymmetric, and that this asymmetry declines with income. Taken together, these two results are consistent with low-income households acting as if they are more ambiguity averse, and the source of this ambiguity being the strength of signals households receive.

The final essay embeds these distorted beliefs in a portfolio-choice framework with nominal risk. Non-homothetic robustness is employed, consistent with the conclusions of Chapter 2. When we calibrate the model, we match the empirical profile of inflation forecast errors, as detailed in Chapter 1. We next estimate the welfare cost of distorted inflation beliefs in the model and find that a modest aggregate forecast error has the potential to generate large welfare losses.

Contents

List of Figures

List of Tables

Chapter 1

Income, Sentiment and the Cross-section of Consumer Forecast Errors

1.1 Introduction

As measured in consumer expectation surveys, households on average expect more inflation than actually occurs. Furthermore, there is significant cross-sectional variation in the degree to which households overestimate inflation. These two observations raise further questions: Which household characteristics best explain consumer forecast errors? Second, do such patterns also exist for expectations of other economic outcomes beyond inflation? If so, what best characterises consumer forecast errors? Finally, is the magnitude and cross-sectional variation in forecast errors quantitatively meaningful? This paper addresses these questions by presenting three new results from surveys on consumer expectations. First, we find that across household characteristics, income is the strongest predictor of forecast errors. Second, we show that the phenomenon of persistent forecast errors holds true across a range of economic expectations, not only inflation, and across different surveys. We characterise the direction of persistent forecast errors: households tend to over-predict outcomes they perceive to be negative and under-predict those they perceive to be positive. We describe this as households exhibiting forecasting "pessimism". Moreover, household forecast pessimism declines with income: the poorest households are most likely to display the most forecast pessimism. As a supplementary result we present suggestive evidence that the link between pessimism and income does not just hold in the aggregate, but at the level of the individual. Third, and finally, we find that the magnitude of these results - both the aggregate degree of pessimism and the cross-sectional variation - is quantitatively large.

We begin with a recent finding in the literature showing consumer expectations exhibit a low-factor structure, which can be interpreted as “sentiment”. We show that this measure of “sentiment" is most strongly determined by income group, out of all available demographic covariates. This is our first of three main results, and motivates our focus on household income throughout later sections. Consumers in lower-income groups display expectations of the future that are typically associated with worse economic outcomes in the future, relative to those in higher-income groups. This result is a direct refinement of the results of Kamdar and Ray (2024).

Given the importance of inflation expectations within macroeconomics, we proceed to investigate this first result further. Working on the basis that households dislike inflation[1], we find that there is a large degree of heterogeneity in the extent of inflation expectation “pessimism”. This provides further credence to our analysis of categorical expectations and supports the first fact. In particular, the degree to which households form expectations of inflation, greater than realised, is larger at the low end of the income distribution than the top. To clarify: poorer households display positive, and larger, inflation forecast errors on average than more affluent households. Above and beyond the current literature, we show that the result for inflation is particularly robust: we employ data from the US, using the Michigan Survey of Consumers and the New York Fed’s Survey of Consumer Expectations, as well as Europe, using the European Central Bank’s Consumer Expectations Survey. The results presented hold true across all three surveys.

We next turn to our second result and provide evidence that forecast pessimism - which declines with income is a feature of expectations of aggregate outcomes beyond just inflation, as has typically been studied to date. We do so on an expectation-by-expectation basis, rather than reduce the dimensionality of (categorical) household expectations to one dimension. Across a wide range of economic variables such as interest rate changes, unemployment changes and narrower price indices, we show that on aggregate consumers display expectations which overweight outcomes typically perceived to be welfare negative. Once again, we find that this phenomenon is most significant for low-income households. We then show that this pattern holds not just across expectations of aggregate outcomes, but across expectations of personal outcomes too. Using the recently published ECB Survey of Consumer Expectations, which retains respondents for a longer time than both the MSC and SCE, we are able to construct

[1] Stantcheva (2024) provides evidence for this. We take it as given for the purposes of our paper.

measures of household expectations over personal consumption outcomes. These too display the same patterns as expectations over aggregate outcomes.

These results taken together highlight that across three commonly used consumer expectation surveys, households persistently form expectations that can be classified as "pessimistic". In practice, this means that the forecast errors for households are persistently positive for variables that households perceive as increases to be detrimental. Our third result is that the size of this discrepancy is large - more affluent households have an inflation forecast error roughly half the size of poorer households, albeit still positive. This result holds even after controlling for demographic covariates.

As a final exercise, we also consider whether or not the relationship between income and forecast error exists at the level of the individual. The results described so far exploit variation in expectations between income groups across a moment in time. We therefore leverage the panel dimension of the MSC and use the variation in income at the individual level. We construct a proxy measure of an income shock, and test whether or not forecast errors grow (decline) with a negative (positive) income shock. That is to say, we investigate the possibility that the relationship between income groups and inflation expectation forecast errors is perhaps more than merely a cross-sectional correlation. Although our proxy for income shocks has limitations, we find suggestive evidence that after a positive income shock, consumers' forecast errors decline by roughly 10 percentage points.

These three results taken together are important for two reasons. First, surveys of consumer expectations matter: a body of evidence details their importance for household decision making[2]. Our results therefore help to explain cross-sectional variation in household consumption and savings behaviour in response to shocks. Second, they provide a set of moments which future theoretical work can benchmark models with: models of business-cycle dynamics - particularly those studying the cross-sectional effects of dispersed beliefs - should hope to match that consumers typically display pessimistic forecasts, and that this effect declines with income. Moreover, models of expectation formation should be able to explain *why* we see such patterns in the data. This paper is a presentation of the facts: we leave questions as to the possible causes of our empirical facts to later work[3].

Literature. A large literature has investigated the empirical properties of surveys of consumer expectations. A robust finding is the systematic upward bias of household inflation expectations has been detailed in US and European households in both the MSC and the SCE (see Carroll (2003), Gorod-

[2]See D'Acunto and Weber (2024) for a review of the evidence.

[3]Chapter 2 of this thesis explores one possible explanation - ambiguity aversion - and provides a refinement of the types of ambiguity households face that are consistent with the facts presented in this paper.

nichenko and Sergeyev (2021), Weber et al. (2021), Bhandari et al. (2024)). This paper illustrates that, when framed as "pessimism", persistent forecasting errors also exist not only in inflation but also in several other important economic variables, in both the US and Europe. Considering the cross-section, D'Acunto, Malmendier, et al. (2021) provide evidence that the inflation expectations of women is typically higher than men. D'Acunto, Hoang, et al. (2023) find evidence that forecasting ability is increasing with measured IQ, and with financial literacy. The work of Malmendier and Nagel (2016) highlights the correlation with age (cohort): those who have lived through periods of high inflation typically display larger forecast errors. Fofana et al. (2024) highlight a number of demographic determinants of consumer inflation expectations, of which income is consistently one, if not the, strongest. Relative to these findings, our first result - that income is among the strongest demographic predictor of an individual's expectations *generally* - spotlights the role of income in expectation determination.

In keeping with our result, there is a smaller body of work that has emphasised household economic conditions as a determinant of household inflation forecasts. For example, work such as Nord (2022) documents that the accuracy of forecasts increases with wealth in the Netherlands. Masolo and Monti (2024) find evidence suggesting that the inflation forecasts of savers are systematically larger than those of borrowers. Angelico and Giacomo (2024) employ a model of expectation formation under salient memory retrieval and provide evidence to suggest that differences in realised inflation experiences across the income distribution can help explain differences in expected inflation. Our findings are in line with these contributions. However, the framework in this paper also helps to explain our empirical evidence that forecast errors of other variables - beyond inflation only - also show "pessimism", and also decrease with income.

A closely related strand of the literature has used household survey data to investigate deviations from Full Information Rational Expectations, beginning with Mankiw, Reis, and Wolfers (2003). For example Coibion and Gorodnichenko (2015) introduce tests for information rigidity, and find that both professional forecasters and consumers exhibit predictable errors. Work such as Bordalo, Gennaioli, Ma, et al. (2020) highlight that forecast errors and forecast revisions are typically negatively correlated[4]. This implies that agents overreact to economic news. Other papers considering the dynamics of inflation forecasts include Angeletos et al. (2020), Kohlhas and Walther (2021) and Broer et al. (2022). Fofana et al. (2024) illustrates that the *dispersion* of inflation expectations is systematically affected by identified shocks, finding that the second moment of household expectations rise in the quarters following a shock. Pfäuti (2023) measures the inflation attention threshold of individuals in the Michigan Survey

[4] For a discussion as to why this may be misleading, see Section III of Reis (2021a).

of Consumers and finds that individuals typically become attentive to inflation when inflation exceeds 4%. Our second and third set of results highlights that theories of consumer expectation formation must adhere to the above established properties, but also be able to explain why consumer inflation forecasts also correlate with income. We also contribute to this literature by illustrating that the deviations from FIRE are not unique to inflation expectations alone.

Other work has emphasised patterns across *all* consumer expectations however. Kamdar and Ray (2024) finds empirical evidence of a low factor structure of consumer beliefs (which they term as "sentiment"), and illustrate how this can rationalise the correlation of inflation and unemployment survey responses and the supply-side reasoning of households (e.g. Candia et al. (2020) and Andre et al. (2022)). Haidari and Nolan (2022) provides similar evidence of the correlation between sentiment measures and expectations in survey data. In showing our first result - that a low dimensional representation of consumer expectations most strongly correlates with income- we follow the methodology of Kamdar and Ray (2024) closely to derive our measure. In this context our result can be seen as a refinement of theirs.

All of this work is important to the extent that surveys of household expectations are meaningful: that consumers actually behave in accordance with the expectations elicited. Weber et al. (2022) use an information treatment experiment to alter consumers' expectations, and show that this has quantifiable effects for their subsequent purchasing behaviour. Crump et al. (2022) is similarly able to find that consumption decisions are correlated with expectations of inflation, in a manner as described by the typical IS curve. With this in mind our results suggest that the heterogeneity in consumption and savings behaviour across income groups is driven at least in part by forecast pessimism.

Outline. This paper proceeds as follows. In Section 1.2 we describe our data. In Section 1.3 we present our first result: that the strongest demographic explanatory variable for expectations is income. In Section 1.4 we illustrate our second result. We show that i) persistent forecast errors exist across a number of other economic variables; ii) that the direction of these errors always points in a direction consistent with a notion of forecast pessimism; iii) that forecast pessimism declines with income. In Section 1.5, we show that the difference in forecasting errors across income is quantitatively large. After detailing our three main results, we investigate whether these aggregate relationships may hold at the level of individual: we explore the effect of individual income changes on forecast errors in Section 1.6. Section 1.7 concludes.

1.2 Data

We first outline the data sources we use in the remaining sections of the paper.

1.2.1 Survey Details

We draw on three commonly used surveys of Expectations in the US and Europe. These are:

New York Federal Reserve Survey of Consumer Expectations (SCE) The SCE is an online rotating panel survey of approximately 1300 consumers per month and has been conducted at a monthly frequency since 2013. The respondents remain in the sample for up to 12 months. The survey contains a core monthly module on expectations over various economic variables, as well as questions over detailed demographic information. The core survey is periodically supplemented with more detailed modules, which ask participants further questions on more specific topics. The survey is designed to be representative of the US population, matching the demographic targets reported in the American Consumer Survey[5]. Survey weights are provided to match population characteristics. The survey is of particular use for its breadth of expectations elicited over a number of differing economic variables. The final survey wave included in the dataset is September 2023.

Michigan Survey of Consumers (MSC) Given the short time span of the SCE, we also incorporate data from the MSC, which began in 1978. This affords us nearly 35 years of more data. The MSC is largely a repeated cross-sectional survey, and the same individuals are generally not tracked for extended periods of time[6]. Survey weights are provided to ensure that the analysis is representative of the US population. The MSC is a smaller scale survey than the SCE, which presents us with two disadvantages. First, the size of the survey is smaller, with only 500 respondents per month. Second, the breadth of the questions is also smaller. This means that we cannot present evidence on the same number of variables as we do in the SCE. Nonetheless its long timespan and ubiquity within the literature merit its inclusion. The final survey wave included in the dataset is March 2024.

European Central Bank Consumer Expectations Survey (ECB-CES). To test whether pessimistic forecasting biases are unique to US consumers, we also incorporate survey evidence from the European Union. In early 2020, the ECB launched its new Consumer Expectations Survey[7]. This is also a monthly rotating panel, representative of the populations in each country the survey is fielded. Survey weights

[5]See Armantier, Bruine de Bruin, Topa, et al. (2015) for more detail.

[6]Roughly one third of the total sample participated in a second round of the survey 6 months later

[7]See Bańkowska et al. (2021) for more.

are provided to match population characteristics. By design, the ECB-CES asks a range of questions that are generally very similar to those asked in the SCE. Despite the short time-series length, it offers two advantages over the US surveys. First, it is extremely large and aims to sample roughly 20,000 EU households per month, in 11 countries[8]. Secondly, respondents are not removed from the survey after 12 months. This allows for the construction of forecasting errors over household-level outcomes: More specifically, given that respondents are asked about both their expected expenditure growth over the coming 12 months and their expenditure growth over the previous 12 months, we can construct expectation errors over outcomes at the individual level, for those respondents who remain in the survey for 13 or more rounds. Naturally, a concern for European evidence will be that much of the survey period incorporates the Covid pandemic period. Although we attempt to provide some robustness tests to ameliorate this issue, our results must be viewed with this limitation in mind. The final survey wave included in the dataset is March 2024.

Comparing across datasets. Using different surveys comes with the advantage of providing external validity to any results found within one particular survey. However, this sizeable benefit comes with two caveats. The first is that the precise set of questions varies across surveys. All surveys contain the most typical demographic covariates (i.e. age, geography, gender, race, education, marital status). But beyond this the surveys differ in the demographic detail they provide. This is important, as from Section 1.4.1 onwards, we typically want to include a full set of controls. This leaves us with a dilemma: do we maintain the same set of (harmonised) controls across all surveys, and increase the possibility of omitted variable bias? Or do we include all the demographic information available, and perhaps reduce the inter-compatibility of our results? For the sake of this paper we value the cost of the first option to outweigh the benefits, and therefore fall in the second camp. However given this, we highlight in the main text the demographic controls used where required. Full detail on the breadth and structure of our controls can be found in Appendix A.1.

The second caveat is that the exact wording of questions, and the elicitation procedure, may differ across questionnaires. For example, for inflation, the MSC and ECB-CES ask individuals about the "prices in general", whereas the SCE instead asks about the "rate of inflation". These differences in survey design may generate differences in respondent engagement and comprehension, and so also generate differences across surveys. There is less we can do to mitigate these concerns. We therefore highlight where there are differences in question wording when relevant.

[8] These are Belgium, Germany, France, Spain, Italy, Netherlands, Austria, Finland, Ireland, Greece and Portugal.

1.2.2 Construction of Forecast Errors

In Section 1.3, we focus mainly on the raw expectations of individuals alone. However in Section 1.4 onwards, we turn to exploring forecast errors in particular. In these sections, we typically define a forecast error as:

$$\text{FE}(x)_t^i \equiv \mathbb{E}_t^i[x_{t+12}] - x_{t+12}$$

In words, the 12 month ahead forecast error over a variable x_t is defined as individual i's expectation of x_{t+12} in period t, minus the realised value of that variable. So, if respondents perceive an increase in a variable x_t to be negative for welfare, a persistent positive forecast error can be seen as evidence of forecast pessimism. In Appendix A.1 we provide a full list of the outcome variables we measure each forecast error against.

There are two possible deviations from this construction to highlight. First, for most variables for which the consumer expectation surveys elicit expectations, increases in x_t can typically be taken to be perceived as negative by the respondent. However, for some variables, this is not the case. For example, more likely than not, the average respondent perceives an increase in their personal income to be positive welfare. Indeed, this is the case in Section 1.3. Therefore for such variables y, we invert our forecast error measure (i.e. $\text{FE}(y)_t^i \equiv y_{t+12} - \mathbb{E}_t^i[y_{t+12}]$). This means that the forecast errors we present are always consistent with increases representing more pessimistic beliefs. For the full list of variables for which this is done, please see Appendix A.1. However, we will highlight in the main text where this has been done too.

The second deviation is where the elicitation procedure does not provide us with a straightforward way to calculate the forecast error. For example, the ECB-CES asks respondents to provide a point-forecast for one-year ahead unemployment, whereas the SCE asks for the percentage chance that the unemployment rate is higher in 12 months from now. We therefore have to construct alternative measures of a forecast error for cases such as these. Again where we do this is detailed extensive in Appendix A.1, but we will again discuss this in the main body of the text.

Outliers. Surveys of expectations, like any survey, are open to the possibility of measurement error. In order to reduce the effect of outliers (potentially due to input error), we trim all point forecasts at the 5-95% level by month. This also helps mitigate the effect of *true* idiosyncratic extreme expectations on our point estimates. Naturally we may worry about this possibility for our right-hand side variables too. This is generally less of a problem here, as many of the controls are categorical. For continuous income measures (which is the measure provided in the raw MSC), we may be concerned again about the effect of outliers. We therefore apply the conservative approach of first dropping any observations outside of the 5-95 percentiles of income by month, and then binning our measure into categories. This also allows for easier comparison across surveys, as income in the SCE and ECB-CES is grouped.

1.3 Sentiment and Income

This section illustrates our first result: that the cross-sectional variation in consumer expectations is most explained by household income. Recent work (e.g. Kamdar and Ray (2024)) has emphasised the low factor structure of consumer beliefs as reported in surveys. Across both expectations of personal and aggregate results - as well as forward and backward-looking beliefs - a remarkable result of their paper is that beliefs are well approximated by a single latent factor[9]. The low factor structure of consumer beliefs drives the notion that a major determinant of consumer expectations is "sentiment". The natural question to ask is what drives the differences in sentiment *between* households. We motivate our later findings by showing that a key determinant of what the literature has labelled "sentiment", is household income.

In this section, we follow their paper and construct a similar measure of sentiment across the MSC. Our sole departure from their baseline measure is that we only include forward-looking expectations that have been in the survey since 1980. We conduct a multiple correspondence analysis (MCA)[10] across the entire range of forward-looking beliefs that were elicited in each survey. For questions where a survey elicits a continuous response, we bin these into terciles. We report the loadings on each of the questions in Figure 1.1. In Appendix A.2 we carry out the same exercise on the SCE data and draw the same conclusions.

The first component explains a large proportion of the variation in these variables. In the MSC, 75.5% of the variation is explained by the first factor alone (and 6.12% is explained by the second factor)[11]. As in Kamdar and Ray (2024), the first factor loads on the variables in question in a manner consistent with a "pessimistic" vs. "optimistic" interpretation[12]. Consumer beliefs over the future are highly cor-

[9] Across the 15 variables which the MSC has regularly questioned respondents since 1980, this factor explains over 80% of the variation in reported subjective beliefs.

[10] The categorical analogue of the perhaps more familiar principal component analysis (PCA).

[11] This is marginally lower than the findings of Kamdar and Ray (2024) in the MSC. This is driven entirely by our inclusion of PX5 and the exclusion of backward-looking beliefs. In the SCE, the first and second factors explain 71.8% and 7.94% of the variation in the beliefs of the respondents about the future, respectively.

[12] Kamdar and Ray (2024) name this "sentiment". In their appendix, they show that the first dimension of their MCA analysis, when aggregated across individuals, is highly correlated with other typical measures of sentiment in the literature. In particular, Gallup's "economic confidence index", the OECD's "consumer confidence index" and even the "news index" of Shapiro et al. (2022).

Figure 1.1: MCA loadings on Expectations in the MSC

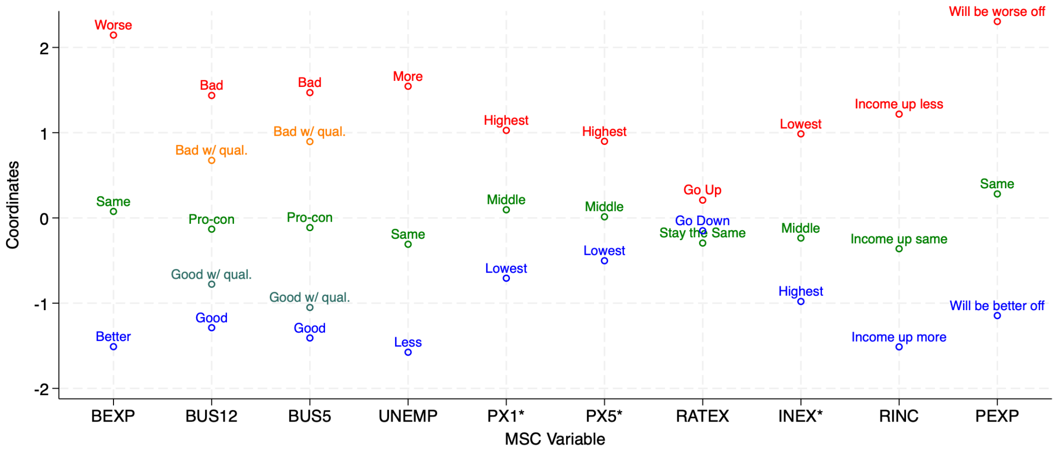

Each point represents the estimated loading of the first component for a categorical response in the MCA. Each variable on the x-axis represents a different question from the MSC. They are as follows: business conditions in one year relative to now (BEXP), business conditions over the next year (BUS12), business conditions over the next five years (BUS5), unemployment over the next year (UNEMP), inflation over the next year (PX1), inflation over the next 5 years (PX5), interest rates over the next year (RATEX), family income over the next year (INEX), family real income over the next one to two years (RINC), personal financial conditions in one year (PEXP). Inflation expectations (PX1 and PX5) and family income over the next year (INEX) are quantitative beliefs elicited in the MSC - these are therefore binned into terciles by survey month. As before, these are also trimmed at the 5-95th percentiles (within each survey month). Answers consistent with a notion of consumer "pessimism" have been coloured red, and those consistent with "optimism" blue (with the exception of RATEX, for which this notion is less clear). Here we have normalised the first-dimension of the MCA analysis such that higher values are consistent with more "pessimistic" expectations of future outcomes.

related, and responses associated with "pessimistic" outlooks have high, positive loadings. Consumers concerned about the future of aggregate variables or their own personal finances typically view high inflation and high unemployment (and higher interest rates) as more likely than their peers. Put more starkly, not only do those consumers who expect high inflation also typically expect high unemployment, but they also exhibit pessimistic beliefs over notions as diverse as business conditions and family income.

The natural question arises: How does this measure of sentiment vary with the characteristics of the respondent? To answer this, we regress the demographic characteristics of the household in the first dimension of the MCA analysis. We estimate the following regression:

$$\text{MCA}(1)_{it} = \beta_0 + \boldsymbol{X}_{it}\boldsymbol{\beta} + d_t + \varepsilon_{it}$$

where $\text{MCA}(1)_{it}$ is the first factor of the MCA analysis, and $\boldsymbol{X}_{it}$ is our vector of demographic covariates. The coefficients of interest are therefore the vector $\boldsymbol{\beta}$. Alongside binned income as discussed above, we include age, age squared, marital status, gender, education and census region in our baseline specification. The last four of these variables are categorical[13]. We include time-fixed effects to account for aggregate factors possibly affecting the underlying expectations.

We show the results of this regression graphically. For the full regression output, please see Appendix A.2. In figure 1.2 we plot the coefficient estimates of this regression. Age, marital status, and region have minimal explanatory power over our MCA measure of sentiment. However, gender, education, and income tercile do appear to co-vary strongly with our measure. Strikingly, women are considerably more pessimistic than men [14]. "Pessimistic" expectations are decreasing in both educational attainment and income, with the largest marginal effect of all covariates falling on the latter. Appendix A.2 illustrates the same exercise in the SCE which shows much the same patterns[15].

Figure 1.2: Marginal Effects of Demographics on MCA First Factor

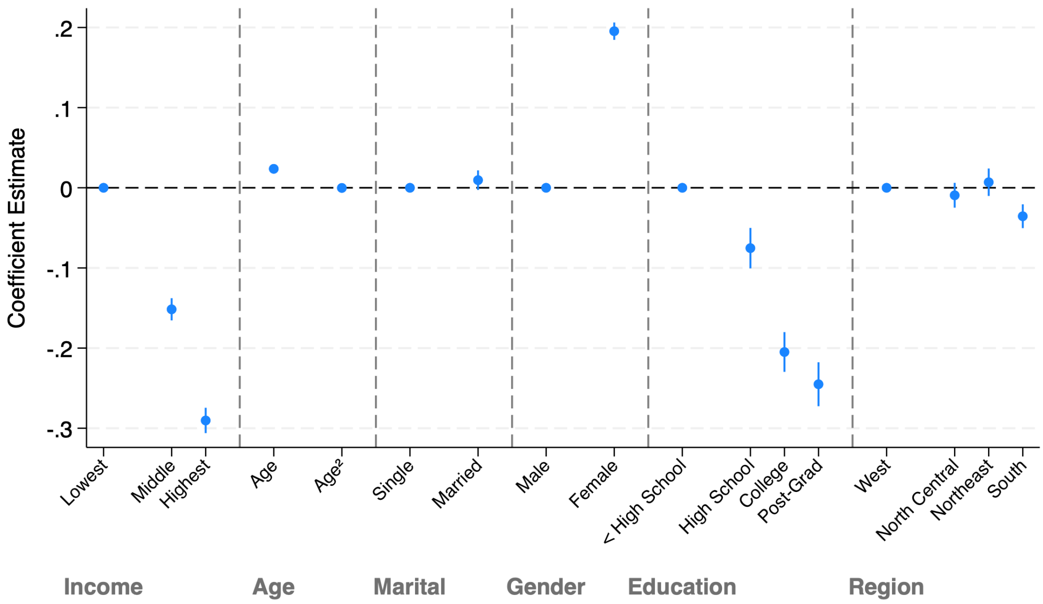

This graph plots the coefficients of the regression of the first MCA dimension on demographic characteristics and income tercile, in the MSC. Each of the covariates listed is an indicator variable, with the exception of age and age squared. The regression includes month fixed effects, and standard errors are clustered at the individual level. Survey weights are employed, and the number of observations is 136,443. The error bars denote the 95% confidence interval. Note that the first column for categories income, marital status, gender, education and census region are all the baseline categories, with which the others should be compared - a coefficient of zero is therefore by construction. For the full regression output, see Table A.2.

To explore further the relationship between the measure of sentiment and income, we employ the MSC's continuous measure of income. We deflate this measure using the BLS CPI and to remove the effect of potential outliers we at the 5-95% mark by month. We follow Cattaneo et al. (2024) and use a binned

[13]See Appendix A.1 for exact definitions.

[14]this can be viewed as a corollary result to the common finding in the literature that they typically have higher forecasts of inflation.

[15]Interestingly, both race and numeracy also have predictive power in the SCE, but income group remains the strongest.

scatter. This has the benefit of allowing us to control for time-fixed effects (and demographic covariates), while also being able to calculate consistent standard errors and confidence intervals[16].

Figure 1.3 illustrates the relationship between these two variables non-parametrically. Once again, the result is striking: the MCA measure of sentiment in the MSC is strongly declining in income. It should be stressed that these are marginal effects: It is not that richer people are typically more educated that drives these results. Furthermore, this set of results suggest that it is unlikely that the relationship between income and pessimistic expectation formation is driven by measurement error. By extracting our measure of sentiment from categorical responses (and coarsely binning continuous responses), we likely reduce any possible correlation of user-input error and income to zero.

Figure 1.3: Bin-scatter of Income Against MCA First Factor

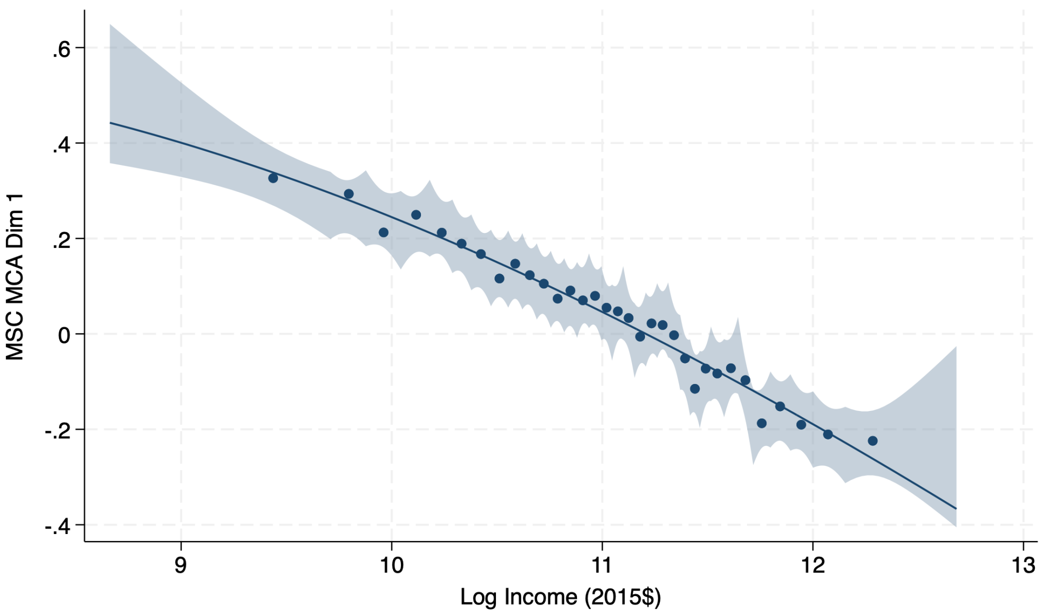

Here we plot a binned scatter graph of log real income against the first MCA factor in the MSC, following Cattaneo et al. (2024). We control for age, age squared, marital status, gender, educational attainment and region. Month fixed effects are employed, and standard errors are clustered at the individual level. A third-order polynomial is fitted and the shaded area denotes the confidence interval at the 95% level. The binning scheme is determined optimally in a data-driven way - please see the referenced paper for more detail.

As pointed out by Kamdar and Ray (2024), this has an important implication for theories which attempt to explain how consumers form their expectations. In their paper, they stress that the low-factor structure of beliefs implies expectations must be well explained by one variable in a given model. That this variable declines across income groups gives rise to a refinement of that conclusion: A theory of expectation formation which matches consumer survey data must not only be explained well by a single (or very few) variables in a model, but this variable must decline by income.

[16] Please see Cattaneo et al. (2024) for more.

1.4 Forecast Pessimism and Income

The fact that lower-income households are typically more "pessimistic" relative to higher-income households reveals little about whether households are pessimistic or optimistic *in general*. To shed light on this specifically, we turn to measuring consumer forecasting errors. In this section we illustrate our second set of results: we show that: i) persistent forecast errors exist across a number of other economic variables; ii) the direction of these errors always points in a direction consistent with a notion of forecast pessimism; iii) forecast pessimism declines with income.

At the individual level, one observation of a positive forecast error is not indicative of pessimistic forecasting alone. In any given period, the forecast error will be a function of both the individual's expectation formation process, unforeseen aggregate shocks, and individual-level forecast error. For our measure to accurately reflect forecast pessimism, we rely on group-level results to remove the effect of individual-level forecasting mistakes. However, this does not account for the effects of unexpected aggregate shocks. To interpret the level of our measure, we therefore require the survey sample to be long enough such that the law of large numbers bites, and so reasonably believe the effect of aggregate shocks on forecasting errors approach zero over the period in question. For the MSC, the survey with the longest sample in our analysis, this argument will be most convincing. For shorter surveys, this may be less convincing. However, since aggregate shocks in any given period are homogeneous across groups (by definition), comparing across groups removes the aggregate effect as in a difference-in-difference regression specification. For these cases, we therefore emphasise the results across income-groups, thereby focussing on the group-level heterogeneity.

We now detail the evidence of pessimistic expectations across various economic outcomes. As inflation expectations are central to business cycle dynamics in many macroeconomic models, and have therefore garnered the most attention within the empirical literature studying consumer expectations, we investigate forecasting errors across inflation separate to other variables.

1.4.1 Inflation Forecasting Errors

The results of Section 1.3 suggest that consumers who expect higher inflation typically also feel worse about future economic outcomes. Other recent work, such as Stantcheva (2024) and Macaulay (2022) provides more direct evidence that consumer perceptions of higher inflation are bad for personal welfare. Given this, we work from the assumption that consumers who have higher inflation expectations typically have a more *negative* view of inflation.

Aggregate Evidence Across different datasets, countries, and time periods, the systematic upward bias in inflation expectations is well documented (see Weber et al. (2022), Candia et al. (2020)). The results shown in Section 1.1 further support this.

Table 1.1: Aggregate Inflation Forecast Errors

	Mean	Median	St. Dev	N
MSC				
Full Sample	0.24	0.03	4.18	252702
1990-2020	0.93	0.48	3.25	159141
2013-2020	1.21	0.87	3.10	47168
SCE (Point)				
Full Sample	3.08	1.77	6.20	143000
2013-2020	3.20	1.82	5.63	108224
SCE (Density)				
Full Sample	1.56	0.90	6.10	154265
2013-2020	1.86	1.05	5.71	116519
ECB-CES				
Full Sample	-0.14	-1.31	6.24	543736
2022-	1.94	0.82	6.76	329419

Comparison of aggregated inflation forecast errors across the three surveys, over differing time periods. Point forecasts are trimmed at the 5-95% level, by date. Summary statistics constructed using survey weights.

Consider first the two US surveys. Our measure of the forecast error for each respondent as the forecast of expected inflation 12 months ahead minus realised inflation in 12 months. We use the national BLS CPI as our measure of inflation[17]. Survey weights are used to construct the summary statistics.

Over the full sample aggregate inflation forecast errors as measured by the MSC and the SCE differ substantially – a large proportion of this difference is due to inflation regime shift that occurred in the US during the 1980's, when inflation was higher, more unstable, and therefore harder to predict. Over the period 1990-2020, when inflation in the US was generally low and stable, the size of this bias is roughly 1.2% in aggregate in the MSC data. Using the pre-Covid period during which the two surveys overlap, the average aggregate inflation forecasting errors are much closer to one another[18].

[17]We prefer CPI for two reasons. First, the CPI basket weights are updated every two years, rather than quarterly as for the PCE. The CPI measure is therefore generally higher than the PCE (Bullard (2013), CEA (2023)), as it under-captures the substitution effect of consumers switching to cheaper items when prices increase. Using CPI inflation therefore serves to make the forecasting error smaller, making our measure of forecasting pessimism conservative. Secondly, it only considers expenditures made directly by consumers, rather than also incorporating purchases made by third parties (e.g., medical products or services purchased on behalf of a consumer by an insurer). It is therefore likely to be the most salient measure of inflation for a consumer and is therefore more appropriate when constructing forecast errors.

[18]Some of the disagreement across surveys may also be generated by differences in the quality of the elicited expectations. In the SCE, respondents are asked about "the rate of inflation", while in the MSC they are asked about "prices in general". As Bruine de Bruin, Klaauw, et al. (2012) and Van Der Klaauw et al. (2008) show, surveying consumers about "the rate of inflation" generates more homogeneous interpretations of the question and is more likely to lead the respondent to consider general inflation in the US and less likely to evoke considerations of specific price changes.

In the ECB-CES, respondents are asked about 'prices in general' (as in the MSC) in their country of residence. We measure inflation expectation errors using country-specific realised CPI inflation 12 months ahead[19]. Looking at the density-implied error only in the full sample, the ECB-CES inflation forecast error is mildly negative. This is unsurprising, given that both the Covid pandemic and the inflationary shock of the Russian invasion of Ukraine comprise most of the period and are both events associated with very large aggregate shocks (particularly for inflation, in the latter case). When considering the period after the European inflationary peak (that is, mid-2022), we can see that the systematic upward bias of inflationary expectations is restored.

For inflation, both the SCE and the ECB-CES ask respondents to report the probability that inflation in 12 months lies in particular ranges. Using these probabilities, a distribution can be fitted using the procedure of Engelberg et al. (2009). These measures have a range of benefits[20]. We also report the aggregated density-implied mean. Across the entire sample, the average forecast error using the density-implied measure is closer to the MSC. However, for the ECB-CES, the opposite is true. It is not clear which measure provides a more conservative estimate of the systematic bias in both surveys. For the sake of comparison (with both the MSC and other variables), we will generally use the point forecast but will highlight where the density implied forecast is used[21] .

At least in aggregate, consumers in the US and the EU typically overestimate inflation over the next 12 months. This phenomenon is robust to the choice of survey, the choice of the elicitation of inflation expectations, or the choice of a time-period sample. Next we turn to how these persistent inflation forecasting errors vary within the cross-section.

Cross-sectional Evidence We now turn to inflation forecast errors by income group. We group respondents by reported household income into 11 different groups, both for nominal income and income measured in 2015 dollars.

In Figure 1.4 we plot the average inflation forecast error by group, in both the MSC and the SCE from 2013 onwards[22]. We define the forecast error as in Section 1.2: $\text{FE}_{it} \equiv \mathbb{E}_t^i[\pi_{t+12}] - \pi_{t+12}$. We denote the

[19]The results are robust to using either CPI inflation, or HICP inflation. See Appendix Section A.3 for more.

[20]see Armantier, Bruine de Bruin, Topa, et al. (2015) for a discussion.

[21]All results we show hold for both measures, nonetheless. See Appendix A.3 for more information. The procedure for generating density-implied forecasts bounds the expectation between -40% and 40%, and requires the sum of probabilities to be 100. The potential for input error is therefore greatly reduced, but we continue to trim this measure for consistency.

[22]See Appendix A.3 for a similar plot for the MSC from 1990 onwards.

sample average forecast error of group g as $\widehat{FE_g}$ and calculate this as:

$$\widehat{FE_g} = \frac{1}{T}\sum_t \sum_j \frac{w_j}{\sum_j w_{jt}}(FE_{gt} - FE_{jt})\ , \quad \text{where} \quad FE_{gt} \equiv \sum_{j \in g} \frac{w_j}{\sum_j w_{jt}} FE_{jt}$$

where w_j denotes the survey weight of observation j [23]. That is to say, it is the sample average of the forecast error of group g, accounting for a time-fixed effect. In words: Figure 1.4 shows the average predicted inflation forecast error within the sample of income group g. The distinction between nominal and real incomes does not affect the results. There is substantial heterogeneity in the magnitude of the upward bias: inflation expectations for those who earn less than \$10,000 dollars per year are approximately 1.8-2.3% higher every month that realised inflation on average, while for those earning more than \$200,000 this bias is in the range of 0.25-0. 5%.

Figure 1.4: Inflation Forecast Errors by Income, MSC and SCE

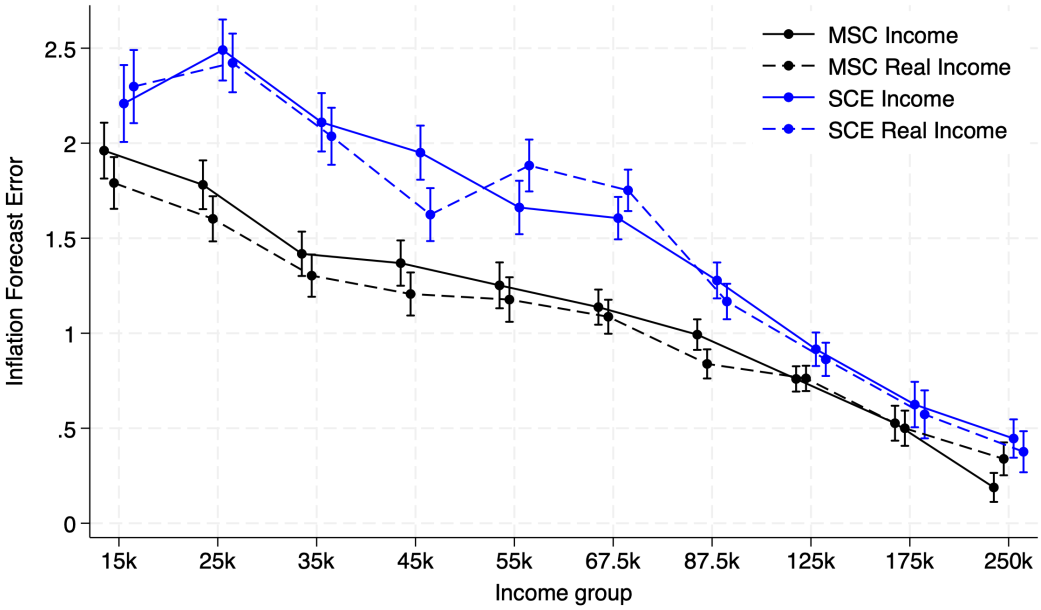

Average inflation forecasting error across income groups within each survey, for 2013 onwards (this is done for the sake of comparison). The solid lines are nominal income groups, and the dotted lines are in 2015 dollars. The x-axis label represents the mid-point of the income group. Here, the SCE Density-implied expectation is used. The graph employs survey weights. The standard errors are clustered at the individual level and the confidence intervals are at the 95% level.

Of course income is correlated with other characteristics the literature has highlighted as pertinent for expectation formation. Malmendier and Nagel (2016) highlight the role of past inflation experiences in inflation expectations and find that age, or more specifically cohort, is a strong predictor of inflation expectations.D'Acunto, Malmendier, et al. (2021) find that the systematic bias in inflation expectations is

[23]This is simply the *margins* command available in Stata, as well as other statistical software programmes.

higher for women than for men. D'Acunto, Hoang, et al. (2023) find that this bias is lower for individuals with higher IQ. We therefore seek to control for these confounding variables in the following regression specification.

We estimate the following regression:

$$\text{FE}_{it} = \sum_k \beta_k Y_{it}^k + \gamma \boldsymbol{X}_{it} + d_t + \varepsilon_{it} \tag{1.1}$$

This is the baseline specification. The dependent variable in our regression is the inflation expectation error. Our explanatory variables are: dummy variables Y_{it}^k which equal unity if a household is within (real) income group k, as well as a vector of controls $\boldsymbol{X}_{it}$[24] Finally, time fixed effects are also included, to control for aggregate factors which may influence expectations of inflation. The inclusion of these fixed effects allows us to focus on the heterogeneity in expectation biases across income groups, as in a difference-in-difference specification.

In Figure 1.5 we illustrate the results of this regression graphically[26]. We use a similar decomposition as in Figure 1.4, but now instead of only controlling for a time-fixed effect, also control for our vector of covariates $\boldsymbol{X}_{it}$ too. The graphs therefore show the sample mean of residualised forecast errors conditional on income.

Across all three surveys, inflation expectations show a similar cross-sectional profile. Our results those earning more than \$100,000 per year have forecasting errors which are roughly 0.5% lower than those earning less than \$50,000 per year on average in the SCE. For the MSC, the income gradient is more pronounced. Those at the bottom of the distribution have a forecasting error roughly 1 percentage point higher than those at the top.

The results for the ECB-CES data show a similar pattern. Here, only the household income quintile is available to the researcher—our results therefore group households into quintiles at the wave-country level by net income. We continue to control for demographic variables[27] and we report regressions using

[24]Our coefficient of interest is the coefficient beta associated with these income group indicators. We control for education, race, marital status, gender, labour force status, state of residence, age, and age squared. The SCE includes questions designed to estimate the financial literacy of the respondent. These are also included. Finally, there is evidence to suggest that prolonged participation in surveys is associated with lower forecasts of inflation[25]. If survey continuation rates are correlated with income, excluding tenure would likely bias our results. Survey tenure is therefore also included. For more detail on precisely how these controls are defined, please see Appendix 1.4.

[26]Regression output is reported in Appendix A.3.

[27]Our set of controls is the same as in the SCE excluding race and financial literacy, which are not included in the survey. Instead of state of residence, we can only control for country of residence.

Figure 1.5: Marginal Effect of Income on Inflation Forecast Errors

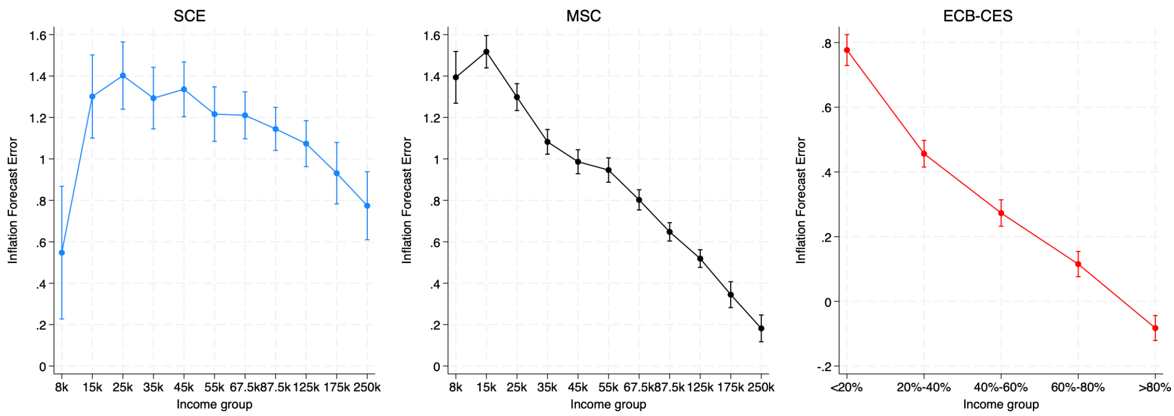

Here we plot the marginal effect of real income group on the inflation forecasting error across the three surveys. All use the regression specification outlined above and include survey weights. Standard errors are clustered at the individual level. One two or three stars denote significance at the 10%, 5% and 1% level respectively. For the SCE and MSC, the x-axis label denotes the mid-point of the income bracket.

time-fixed effects[28]. The results here show a similar pattern. Comparing households in the top 20% of income with those in the bottom 20%, we see that those at the top of the distribution have expectation biases that are approximately 0.8% lower than those in the bottom 20% of the distribution.

Discussion and Robustness. As inflation expectations are of particular interest in the literature, we pay particular attention to demonstrating the robustness of our results here.

Alternative Measures and Samples. The results of this section are robust to using either the density-implied forecast error or the point forecast error. Across both the MSC and the SCE, the results are also robust to using alternative forecasting horizons. Furthermore, we may also be concerned that specific periods of high inflation (such as the post-Covid inflation shock) may be driving our results. We show that our results are robust to the exclusion of 2020 and 2021 from our sample. See Appendix A.3, for more.

Measurement Concerns. We may be concerned that systematic mismeasurement of inflation expectations is generating the results. In our baseline regressions we go some way to guard against this possibility. We drop the 10% of households with extreme expectations to guard against this. We also bin income (in the MSC) for this reason too. Our results also hold for the density implied forecasts, of which one reported benefit is reducing such measurement error. Nonetheless, we explore alternative routes to reassure the reader that our results are robust to concerns of mismeasurement.

[28]Using country-time fixed effects has negligible effects upon the results.

Work by Gorodnichenko and Sergeyev (2021) has highlighted the presence of a zero-lower bound for inflation expectation survey data. They find that across multiple surveys, very few respondents ever report that they expect deflation, despite deflationary episodes coming to pass. Whether or not this is a concern for measurement requires taking a stand on the source of the lower bound. If consumer expectations are lower-bounded for some fundamental reason - if, say, they form expectations in such a way that generates lower-bounded expectations - then we should not be concerned. In such a case, surveys measure expectations correctly, other criticisms of surveys notwithstanding[29]. Another possibility is that households do not understand the concept of negative inflation, perhaps because they have limited economic or mathematical understanding, then controlling for financial literacy as we do above may help ameliorate these concerns. Even in this case it is not clear whether we mismeasure their expectations when we conduct a survey.

The case which is most concerning for the empiricist is one where there is something about the way in which surveys are conducted that makes expectations of deflation harder to elicit for some reason. Survey designers try to avoid such outcomes, but the problem may persist. Note that this will only be a problem for our cross-sectional estimates if both: 1) elicited expectations are lower-bounded, such that they do not reflect *true* expectations of the respondent; 2) the degree to which this happens covaries with income. Another possibility is that mismeasurement exists at the top-end of the distribution. It is certainly true that inflation expectations are skewed upwards in the cross-section. If this: 1) is not truly reflective of a respondents beliefs, and; 2) is correlated with income, this may provide another cause for concern. For example, one might think that lower income households are more likely to make typographic errors - where they add extra zeros, say - for some reason. This again would be a problem for our results.

Even still, in Appendix A.3 we carry out exercises to provide robustness to these possibilities. We find some evidence to suggest that variance, and skew, of inflation forecast errors declines as income increases, although the effect is not particularly pronounced. To check for the possibility of zero-boundedness biassing our estimates, we estimate similar regressions over population sub-samples for whom we are least concerned about this issue[30]. Our results are robust to this sub-setting. To check for the possibility of right-tail skew mismeasurement being correlated with income, we conduct the same regression over those respondents whose forecast errors fall within the interquartile range, per survey

[29] To some extent, this is part of the effect of ambiguity on signal information content, as in Chapter 2. Relative to positive news of the future, on receipt of negative news, expectations do not decrease as much. It is this asymmetry, along with the assumption that signals are mean-zero, which generates a positive forecast error across many individuals (or for one individual, over time).

[30] These are: those respondents who ever report a higher probability of deflation than inflation during their time in the survey; those respondents whose weight on deflationary outcomes varies the most - as measured by being in the 75th percentile or above by standard deviation of probability assigned to deflation; the intersection of these two samples

wave - here too, the results remain robust. Finally we estimate a quantile regression, to estimate the effect of income on the median inflation forecast error per group. Here the marginal effect of being in the top versus bottom income results in a lower *median* inflation forecast error of 0.49 percentage points. This is in line with the magnitude presented in our baseline specification.

Alternative Controls. There are a number of alternative controls that we may be interested in including, too. As Malmendier and Nagel (2016) demonstrate, inflation experience has predictive power over inflation expectations and we may be concerned that age is not a sufficient control for inflation experience in a panel setting. Fofana et al. (2024) find that the relationship between age and inflation expectations in the MSC has changed over time, possibly reflecting the inflation experiences of different cohorts. In Appendix A.3 we include the cohort as a control instead of age; the results remain unchanged.

We may also be concerned that the risk attitudes of households inform their expectations. This would be of particular concern if respondents reported their stochastic discount factor weighted expectation. If risk-averse individuals were to do this, then states of the world with high inflation (low consumption) would be associated with higher marginal utility and would be so overweighted. In Appendix A.3 we also proxy risk preferences using self-reported risk appetite as measured in the SCE. This also does not affect our results.

Heterogeneous Inflation. Recent work in the literature has highlighted the different inflation experiences of households across the income distribution, that is, different inflation experiences in the cross section[31]. Jaravel (2019) finds that between 2004 and 2015, the highest income decile nationally experienced an inflation rate that was approximately 0.35% on average lower than the lowest decile. Using scanner data, other work has found larger differences between retail goods[32]. More recent work by Hartley and Mejia (2024) has found even larger differences in inflation at the regional level.

Despite the fact that the SCE survey design is careful to attempt to elicit expectations about national inflation, rather than an individual's inflation experience, we may be concerned that respondents report their expectations about personal inflation in any way. With this in mind, we employ the cross-sectional inflation measures of Mejia (2022) and construct our inflation forecast error using these instead. This data has two benefits: it is constructed in such a way that it aggregates to a measure which follows BLS-CPI closely, and it is available at both the national decile level and the region-quartile level. We find

[31] See Jaravel (2021) for a review.

[32] Jaravel (2019) estimates 0.5-1% difference on average, and Angelico and Giacomo (2024) estimate the difference to be 1.4%. Although this work allows the researcher to observe prices paid Kaplan and Schulhofer-Wohl (2017), it comes at the cost of only representing roughly 40% of household expenditure on goods (Jaravel (2019)). It also over-represents some goods - nearly 60% of the Kitts-Nielsen Consumer Panel are food or beverages.

that compared to our baseline results, using either series leaves our results unchanged (see Appendix A.3 for more detail)[33]. However, we find that controlling for cross-sectional differences in inflation rates leaves the results largely unchanged. Differences in inflation forecast errors across income groups appear to be not driven by differences in inflation experience, but something else.

Inflation Exposure. Finally, we may be concerned that the degree to which households display systematic bias is driven by differential inflation exposure. As Doepke and Schneider (2006) and Auclert (2019) highlight, the balance sheet exposure of households to inflation varies substantially in the cross section[34]. The net exposure of households to inflation is measured as their net nominal position (NNP). Households with negative NNPs gain when there are unexpected increases in the price level, and those with positive NNPs lose. If households are pessimistic in their expectation formation, we would expect those who benefit in inflationary states of the world – i.e. those with negative NNPs – to underweight the probability of inflation persistently and have lower systematic biases than those with positive NNPs. In the limiting case where households face risk through the Fisher channel only, we would expect negative NNP households to show expectations that are systematically *negatively* biased and those with positive NNP to be *positively* biased.

We use SCE's additional survey modules to construct (proxy[35]) measures of household exposure to inflation. We run three additional regressions to investigate how the expectation error of the household varies with the composition of their balance sheet. We control for: 1) mortgage ownership, 2) mortgage ownership and stock of household debt, and 3) a (limited) measure of NNP exposure[36]. We find that mortgage owners typically have forecast errors that are 0.1-0.3% lower on average. We also find that households with higher NNP positions typically have higher forecast errors. Both of these results are indicative of households with positive exposure to inflation risk having lower systematic biases. However, given the small sample size and the unreliability of the data, these regressions should only be considered indicative. Datasets of expectations with more reliable balance sheet information would be

[33]Of course it should be made clear that this only goes some way to accounting for differences in inflation experience – we are not aware of a dataset that contains information on both individual level inflation and inflation expectations. Angelico and Giacomo (2024) come closest, and use the Kilts-Nielsen Scanner data matched to the MSC data using four income groups. Their model of expectation formation based on salient memory retrieval explains 20% the disagreement in expectations across income quartiles.

[34]Younger (and richer) household's balance sheets are primarily comprised of nominal mortgage debt and so benefit during bouts of inflation. Older (and richer) households hold much of their wealth in nominal assets – usually cash - and see their real wealth decline during inflationary episodes.

[35]It should be noted that SCE's balance sheet data are limited and potentially unreliable, and the additional surveys are only completed by a small sub-sample of the survey as a whole. We therefore restrict these regressions to the subset of households whose stock of debt or nominal assets/liabilities is less than three times their income. We choose the ratio 3x income to remove outliers: Auclert (2019) estimates that the maximum NNP asset/liability to income ratio found in either Panel Survey of Income Dynamics (PSID) or Consumer Expenditure Surveys (BLS-CEX) was approximately 2.

[36]Regression outputs and further detail on how measures of balance sheet exposure are constructed can be found in Appendix A.3.

a useful resource for the literature.

1.4.2 Forecast Pessimism Across Other Expectations

We now turn to consumer expectations over other forecasts elicited across the three surveys.

Unemployment We consider first unemployment. We assume that households perceive higher unemployment to be a worse state of the world. Due to differences in the questions posed to the respondents, our measures of forecasting errors for unemployment in the three surveys are not directly comparable. For more detail on how we construct these different measures, please see Appendix A.4. However, consumers on aggregate persistently overestimate the probability of high unemployment (or unemployment increasing) across all three surveys. In the SCE sample, the probability that households place on unemployment is higher in 12 months is 16% higher than realised. In the MSC, a similar measure generates a mean value unemployment forecast error of 14%. In the ECB-CES, consumers expect unemployment to be 5.0% higher in 12 months than materialise.

Figure 1.6: Marginal Effect of Income on Unemployment Forecast Errors

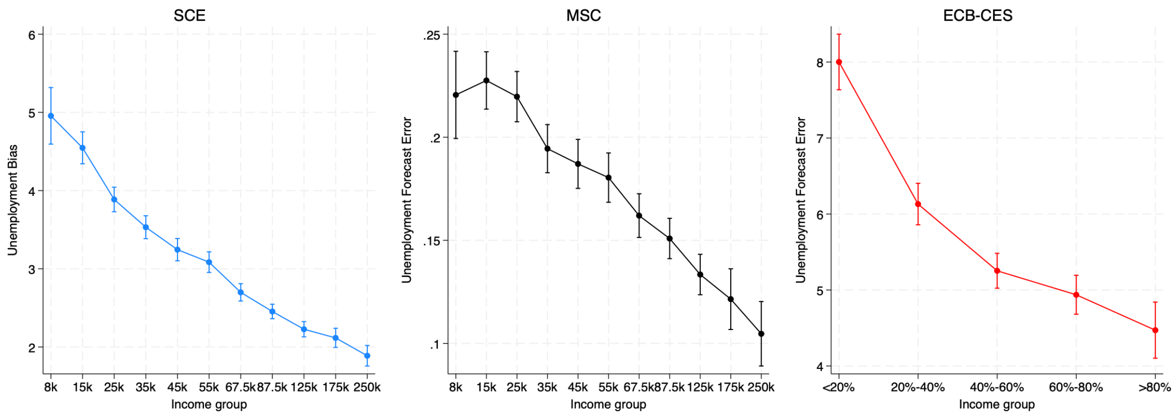

Here we plot the marginal effect of real income group on the unemployment forecasting error across the three surveys. All underlying regressions include demographic controls and time fixed-effects, and employ survey weights. Standard errors are clustered at the individual level. For the SCE and MSC, the x-axis label denotes the mid-point of the income bracket.

We utilise the same regression as (1.1), continuing to control for demographic factors, survey tenure, and geographic and time fixed effects. Here, the effects of income group are less stark but pronounced. In the SCE, those in the \$50k-\$100k and \$100k+ income groups expect approximately a 1.5% lower unemployment rate than occurs relative to the bottom income group. In the ECB-CES, those in the top quintile of the income distribution have a 2.2% lower forecast error than those in the bottom quintile (see Appendix A.4 for regression output). Figure 1.6 shows the marginal effect of real income on unemployment forecast errors in the three surveys, having been controlled for demographic factors and fixed time effects[37] All show a declining profile.

Interest Rates In all three surveys, expectations over changes in nominal interest rates are elicited[38]. However, similar aggregate patterns hold here. In the SCE, respondents typically underestimate the probability of *savings* rates increasing, by 10% over the sample. A similar measure from the MSC shows that consumers overestimate the probability of *borrowing* rates increasing by 0.44% in the sample. In the ECB-CES, consumers expect their home country mortgage rates to be 0.75% higher than materialises in 12 months.

Figure 1.7: Marginal Effect of Income on Interest Rate Forecast Errors

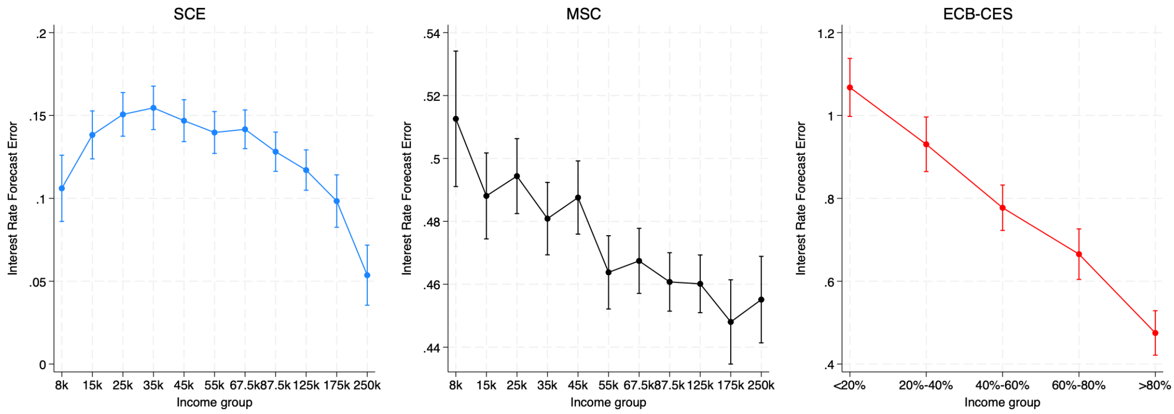

Here we plot the marginal effect of real income group on the nominal interest rate forecasting error across the three surveys. All underlying regressions include demographic controls and time fixed-effects, and employ survey weights. Standard errors are clustered at the individual level. For the SCE and MSC, the x-axis label denotes the mid-point of the income bracket.

Here, both the SCE and the ECB-CES data show strong effects of income group on interest rate forecast. The SCE average forecast error is roughly 40% lower for those earning more than $100,000 compared to those earning less than $50,000. The effect in the ECB-CES is even more pronounced. The top income quintile has an interest rate forecasting error 80% lower than the bottom quintile. See Appendix A.5 for the regression output. Figure 1.7 displays the marginal effects of the income group on the error in the interest rate forecast from the baseline specification. Both the MSC and ECB-CES show declining profiles. The SCE data are markedly more hump-shaped, but decidedly decline from the $> \$35k$ income

[37] The process for visualising these results is the same as Figures 1.4 and 1.5.

[38] As with expectations over unemployment, differences in survey questioning mean that the forecasting errors across surveys are not directly comparable. In particular: the SCE asks respondents for expectations over the *"average interest rate on savings accounts"*; the MSC for expectations over *"interest rates for borrowing money"* and the ECB-CES for expectations over *"the interest rate on mortgages in the country you are currently living in"*. We assume (consistent with the evidence of Section 1.3) that increases in savings rates are perceived as good to the respondent, where as borrowing cost increases are bad. We therefore invert the SCE forecast error measure to be consistent with respondent pessimism, but leave the MSC and ECB-CES measures unchanged. For further detail on the exact construction of these measures, please see Appendix A.5.

groups onwards. As the definition of each forecast error varies between surveys, quantitative estimates should not be compared between surveys.

Real Consumption In the ECB-CES, it is possible to create a measure of the expectations for real consumption. Across other variables, the connection between expectations over aggregate variables and personal welfare relies on extra assumptions – we also require individuals to perceive higher inflation or unemployment to be detrimental for their own welfare. Expectations of real consumption are not open to this criticism.

To construct a measure of expectations of real consumption, we exploit the long time dimension of the ECB-CES. Respondents report both their expected change in expenditures over coming 12 months[39], as well as the percentage change in their actual expenditure over the previous 12 months. As almost half of all observations are in the survey for greater than 12 months, we can construct forecast errors for real consumption after deflating by expected inflation and realised inflation respectively. We invert our measure of expected real consumption growth such that higher forecast errors are consistent with more pessimistic expectations.

Households are asked to report their realised and expected expenditure changes with a numeric answer. If they respond with “Don’t Know”, they are then requested to provide a categorical response based on intervals. We run our regression specification over those who reported a point estimate, as well as a combined sample in which we include the bracketed responses also[40].

The forecast pessimism displayed in Table 1.2 across aggregate outcomes in the previous sections appears to also be present for expectations over *individual* outcomes too. Poorer households display expectations which are pessimistic relative to their richer counterparts. Column 1) shows the specification including country and time fixed effects, using the measure excluding "Don’t Knows". Column 2) includes demographic controls. Column 3) performs the same regression, but also includes respondents who gave a bracketed response. Focussing at columns 2) and 3), those in the top quintile of the distribution have a forecast error that is roughly 1.2% smaller than those at the top of the distribution. This is roughly 75% of the mean forecast error. The similarity between columns 2) and 3) illustrate that the results are not sensitive to including "Don’t Knows", nor omitting these respondents completely. Whether or not we control for demographic factors or use the density implied forecast to deflate expected expenditure

[39] Households estimate how much higher or lower they expect their *“monthly household spending on all goods and services to be 12 months from now”* in percentage terms.

[40] The intervals are less than 20%, -20% to -16%, -15% to -11%, -10% to -7%, -6% to -4%, -3% to -2% -2% to 0%, 0% to 2%, 2% to 3%, 4% to 6%, 7% to 10%, 11% to 15%, 16% to 20% and greater than 20%. We use the mid-point of each bracket as the households expected and reported expenditure change. For the open intervals we use -20% and 20% respectively.

Table 1.2: Real Consumption Forecast Error Regressions

	Real				Nominal
	(1)	(2)	(3)	(4)	(5)
20-40pct	-0.382	-0.428*	-0.407**	-0.453	-0.198
	(-1.61)	(-1.83)	(-1.97)	(-1.51)	(-1.56)
40-60pct	-0.376	-0.401	-0.656***	-0.269	-0.268**
	(-1.53)	(-1.59)	(-2.98)	(-0.86)	(-2.11)
60-80pct	-0.573**	-0.696***	-0.785***	-0.647**	-0.276**
	(-2.53)	(-2.90)	(-3.64)	(-2.06)	(-2.21)
80-100pct	-0.965***	-1.177***	-1.227***	-1.641***	-0.442***
	(-3.86)	(-4.22)	(-4.89)	(-4.81)	(-3.09)
Obs	215,614	215,495	262,846	216,511	186,027
Adj. R^2	0.060	0.070	0.070	0.050	0.060
Mean dep var	1.72	1.72	1.62	3.64	2.1
First obs	2020m4	2020m4	2020m4	2020m4	2020m4
Controls		✓	✓	✓	✓

Here, our measure of the real consumption forecasting error is regressed on income quintiles using the specification outlined earlier. Column 1) shows the baseline specification only including country and time fixed effects, using the measure excluding "Don't Knows". Column 2) includes the set of demographic controls. Column 3) uses the same regression but includes the respondents who gave a bracketed response. Column 4) uses the density implied forecast rather than the point forecast to deflate expected expenditure. The final column uses expected nominal expenditure growth as the dependent variable (i.e. does not deflate by the inflation forecasting error). All specifications include country and time fixed effects. Standard errors are clustered at the individual level and survey weights are employed. One two or three stars denote significance at the 10%, 5% and 1% level respectively.

(column 4), the forecast error profile is decreasing in income[41].

We may be concerned that the patterns we see across real consumption forecast errors are driven by inflation forecast errors alone. To account for this possibility, we construct a measure of the forecast error of household expenditure[42]; a nominal, rather than a real, consumption forecasting error. Here too, the forecast error for nominal expenditure also declines with income, suggesting that the results over real consumption expectations are not driven by the income profile of inflation expectations alone.

We view the results over persistent real consumption forecast errors as a particularly convincing test of pessimistic expectations: of all the variables considered so far, personal real consumption growth is mostly clearly associated with an individual's welfare[43].

[41]When using the density implied forecast as the deflator, the difference between the top and bottom income quintile relative the mean forecast error is smaller, at 45%.

[42]We have also inverted this measure. Higher forecasting errors are therefore associated with underestimating future nominal expenditure relative to its realisation. Of course, this need not necessarily be associated with lower welfare for an individual, as it could be entirely driven by increasing consumption, which presumably would be associated with higher welfare, and is therefore indicative of optimism. We invert the measure for the sake of comparison to the real measures of columns 1)-4) alone. The results of this regression are displayed in column 5).

[43]These results have interesting implications for forward-looking welfare. If low income households persistently under forecast their own real consumption in the future, then in a framework such as Baqaee et al. (2024) - which attempts to measure forward looking welfare in the cross section - we significantly over estimate welfare at the bottom of the income distribution if we do not take this into account. An agent with 40 years left of life, who underestimates their real consumption by 1.2% *each year* requires a *much* greater equivalent-variation transfer today to be made equally well off.

Other Variables Finally, we construct forecast error measures in seven more expectations measures available within the SCE and the MSC. These are stock prices, government debt growth, gas price growth (available in both surveys), food price growth, medical care price growth, the cost of rent and house prices. For more detail on the exact construction of these variables, please see Appendix A.6. We invert stock-prices, assuming that consumers typically perceive increases in stock prices to be good[44]. For rent prices and house prices whether or not increases benefit one individual or another depends on their home-ownership status. We therefore restrict our estimates across these two variables to renters and homeowners only[45]. We assume that homeowners perceive aggregate increases in house prices to be net positive for their individual welfare. Given this, we invert the measure for house prices too.

Table 1.3: Summary Statistics for Other Forecast Errors

	Mean	Median	St. Dev	N
House Prices (owners)				
SCE (Point)	1.80	1.13	8.14	110065
SCE (Density)	2.80	2.24	8.12	101580
MSC (Point)	2.47	3.24	7.57	75989
Rental Prices (non-owner)				
SCE	4.60	2.33	6.84	32481
Government Debt Growth				
SCE	4.44	2.26	11.82	127354
Food Prices				
SCE	3.75	2.76	5.94	128957
Medical Prices				
SCE	7.88	6.09	8.27	127354
Gas Prices				
SCE	4.67	7.58	23.74	129191
MSC	12.37	8.40	33.55	114190
Stock Prices				
SCE	0.42	0.50	0.45	158730

This table displays summary statistics for all of the forecast error variables calculated throughout section 1.4. The summary statistics are constructed using survey weights. The SCE is the full sample, and the MSC from 1990 onwards.

Table 1.3 displays the summary statistics for this final set of forecast errors. All variables show positive aggregate forecast errors, consistent with the evidence presented in previous sections. Moreover many of these effects are quantitatively large. For example, even for something as salient as gas prices, the average household within the SCE sample expects gas prices to rise by more than 4% than is realised per year. This number is even larger in the MSC.

[44] This is an assumption, and of course, may not be true for all households at all times. The possibility that consumers perceive the stock market to be a bubble dangerous for financial stability, or that some households hold net-short positions on stock market indices, is of course possible. These edge-cases notwithstanding, this does not appear controversial for the average respondent at any given time.

[45] In the MSC, only homeowners are questioned about their house price expectations either way.

As before, we now consider the cross-sectional variation. We use the same specification as before using the additional expectations. The regression output is shown in Table 1.4. All forecast errors introduced in this section - with the exception of house prices - show a declining profile in income. The effect of income is particularly pronounced for price categories we might broad categorise as necessities: Rent, Medical Costs and Food. That disparities across expectations of these categories across income are large is perhaps unsurprising. We may be able to rationalise such a result if the reason why pessimism declines across income is due to perceived welfare in given state[46]. If some goods are required for subsistence, and forecast errors are a function of expected utility across states, then we would expect pessimism over these variables to strongly be a function of income.

Table 1.4: Marginal Effect of Income on Other Forecast Errors

	SCE								MSC	
	House (P)	House (D)	Rent	G. Debt	Gas	Food	Med.	Stocks	House	Gas
50-100k	0.930***	0.166	-1.355***	-2.003***	-0.598***	-1.004***	-1.542***	-0.0196***	-0.275***	0.225
	(7.18)	(1.10)	(-6.52)	(-9.99)	(-4.34)	(-9.55)	(-9.81)	(-4.84)	(-6.13)	(1.05)
100k+	1.550***	0.711***	-2.015***	-2.899***	-0.559***	-1.401***	-2.090***	-0.0481***	-0.271***	-0.999***
	(10.62)	(4.38)	(-7.20)	(-12.77)	(-3.26)	(-10.80)	(-12.10)	(-10.39)	(-5.84)	(-4.50)
Obs	109,826	101,349	32,336	126,971	128,793	128,570	126,957	158,227	75,092	111,305
Adj. R^2	0.456	0.409	0.042	0.252	0.882	0.149	0.078	0.737	0.759	0.507
Mean dep var	2.06	2.86	4.14	3.61	4.53	3.41	7.51	.402	2.44	12.3
First obs	2013m6	2013m6	2013m6	2013m6	2013m6	2013m6	2013m6	2013m6	2007m1	1991m1

This table shows regressions of our final set of forecast errors on real income groups across surveys. The regressions employ the baseline specification of (1.1). All regressions therefore include the full set of demographic controls available in the MSC and the SCE respectively. Time fixed effects are also included. All t-statistics are calculated using standard errors clustered at the household level. Survey weights are employed in all regressions. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

The only exception to the pattern of declining pessimism by income is house prices. For the point forecast, there is a strongly increasing profile of forecast pessimism using the point forecast, which stands in opposition to the pattern in MSC data[47]. One possible explanation is the differing samples. The SCE survey began after the experiences of the U.S. housing market crash during the Great Financial Crisis, whereas the MSC sample used begins in 1990.

1.5 Quantifying the Effect of Income on Forecast Pessimism

In this section, we illustrate our third result: that the effect of income on the degree of forecast pessimism exhibited by households is large. To compare the relative size of the various forecast errors discussed

[46] As in the investigation of Chapter 2 or the model of Chapter 3.

[47] House price forecast error is conditional on ownership (whereas the stock price forecast error is not. Therefore it may be that increased exposure of the richest households to house price inflation induces their more pessimistic expectations. However this does not explain why the gradient has the opposite sign in the MSC data.

so far, we normalise each forecast error by the aggregate sample mean. That is, we compute:

$$\text{FE}(x)_{it}^{\text{norm}} \equiv \frac{\text{FE}(x)_{it}}{\overline{\text{FE}(x)}}, \quad \text{where} \quad \overline{\text{FE}(x)} = \frac{1}{T}\sum_t \sum_i \frac{w_{it}}{\sum_i w_{it}} \cdot \text{FE(x)}_{it}$$

where w_{it} denotes the sample weight of respondent i at time t, and x denotes the expectation variable (inflation or unemployment, for example). Our normalised forecast error $\text{FE}(x)_{it}^{\text{norm}}$ therefore has a mean of one, for all x. For example, if $\text{FE}(x)_{it}^{\text{norm}} = -0.505$ (as in the SCE), this implies those in the bottom income group ($< \$50k$) have an inflation forecast error that is 50.5% of the sample mean forecast error greater in magnitude than those in the top income group ($> \$100k$).

We use the same regression specification as (1.1), but instead replace our dependent variable with the normalised measure. That is, for all expectations thus discussed, we estimate:

$$\text{FE}(x)_{it}^{\text{norm}} = \sum_k \beta_k^{\text{norm}} \cdot Y_{it}^k + \gamma \boldsymbol{X}_{it} + d_t + \varepsilon_{it} \tag{1.2}$$

Here Y_{it}^k takes the value one if respondent i belongs to income group k. As we keep the lowest income group as our baseline group. Our coefficient of interest is therefore $\beta_{\bar{k}}^{\text{norm}}(x)$, where $\bar{k}$ is the index of the maximum group within each survey's measure of income, and x denotes a given measured expectation within a particular survey. This has a simple interpretation: it is the relative difference between the top and bottom income group's forecasting error, *as a percentage of the average forecasting error*. Due to data limitations, the precise definition of the income groups differs across surveys. For the SCE and the MSE this is the difference between those earning $> \$100k$ and those earning $< \$50k$ (in 2015\$). For the ECB-CES, this is the inter-quintile difference (measured within each survey wave, per geography). For all of these regressions, we include the full set of demographic controls available within each survey in $\boldsymbol{X}_{it}$, alongside time fixed effects.

In Table 1.5 we report the point estimates and standard errors for all of these coefficients. To maintain the interpretation that larger values indicate greater pessimism, we report: $-\beta_{\bar{k}}^{\text{norm}}(x)$. A positive coefficient in Table 1.5 therefore indicates that the lowest income group has a higher forecast error than the highest income group. We show these same results graphically as well, in Figure 1.8.

Across the thirteen expectation variables for which at least one survey elicits a forecast, ten display a positive and highly significant coefficient, confirming a remarkably robust income gradient in pes-

Table 1.5: High versus Low Income Forecast Error Magnitudes

	SCE		MSC		ECB	
Inflation	0.505	(0.038)	0.733	(0.028)	2.781	(0.264)
Inflation (D)	0.480	(0.080)			0.547	(0.056)
Unemployment	0.086	(0.031)	0.342	(0.028)	0.610	(0.045)
Int. Rates	0.332	(0.040)	0.071	(0.010)	0.757	(0.060)
Real Cons.					1.846	(0.267)
Expenditure					0.281	(0.091)
Stocks	0.075	(0.011)				
Gov. Debt	0.448	(0.054)				
House Pr.	-0.717	(0.082)	0.111	(0.019)		
House Pr. (D)	-0.219	(0.060)				
Rental Pr.	0.352	(0.063)				
Food Pr.	0.291	(0.037)				
Medical Pr.	0.225	(0.023)				
Gas Pr.	0.135	(0.039)	0.084	(0.018)		

This table reports $-\beta_{\bar{k}}^{\text{norm}}(x)$, as taken from equation 1.2. The rows show different expectation variables, and the columns differing surveys. The standard errors are reported in parentheses, and are clustered at the individual level. Entries are left blank in surveys for which the relevant expectation is not elicited.

simism[48]. We do not apply too much weight to the large estimate in the top–bottom quintile gap for the ECB-CES, given the atypical inflation dynamics over its relatively short time span, except to note that the coefficient is positive. Beyond these values in the ECB-CES, most variables display large differences. In the SCE, the average person earning $< \$50,000$ in real terms has an inflation forecast error that is 50% higher relative to mean forecast error, than a respondent earning $> \$100,000$. Even for comparatively "familiar" outcomes such as food or gasoline prices, the bottom-income group still overshoots the top by roughly 30 percent of the mean error (SCE Food = 0.29; Gas = 0.14). Taken together, these results show that the income gradient is not confined to headline inflation but permeates a broad set of economic expectations.

Median Responses. We may be concerned that measurement error is potentially correlated with income group, or that the results of Table 1.5 are driven by differing degrees of skewness in the expectations data within income groups. Therefore to try to mitigate the effect of possible skewness or outliers on the calculation of the means we conduct the following robustness exercise. We estimate a similar regression, but where we estimate the median response of each group relative to the median forecast error. To do so, we slightly modify our normalised forecast error and instead use the sample median forecast error as our normalisation factor[49]. We then employ the same regression specification, but using a quantile regression method.

[48]Only for house price expectations as elicited in the SCE is the relationship between pessimism and income inverted: $-\beta_{\bar{k}}^{\text{norm}}(x)$ is negative. This suggests that richer homeowning households have more pessimistic expectations than than lower income households. What makes house price expectations different to other expectations in this regard is left to future research.

[49]Where the median forecast error in the sample is zero - which is only the case for unemployment in the MSC, due to its construction, we do not normalise the measure. We still estimate the same quantile regression specification.

Figure 1.8: Relative Forecast Errors Across All Expectations

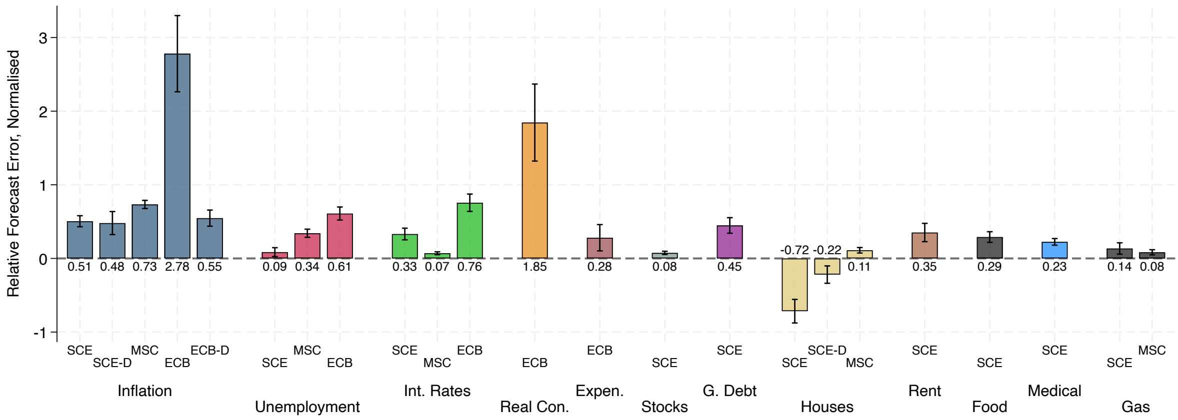

This graph shows the relative forecasting error across all variables, where we have normalised the measure by the sample mean forecast error for each variable, for the sake of comparison. All regressions include time fixed effects and control for the full set of demographic controls available within each survey, as well as geographic fixed effects and survey tenure. A positive estimate suggests that low income groups have larger forecast errors than high income groups. For the SCE and MSC, these groups are households earning less than \$50,000 and those earning more than \$100,000. For the ECB-CES variables, this is the first income quintile and the fifth income quintile. The confidence intervals shown are at the 5% significance level, and test the hypothesis that the normalised relative forecast error is different from zero. The standard errors are clustered at the individual level. The labels below (or above) the x-axis denote the value of the estimate. The inverted measures are: Interest Rates (SCE), Stocks(SCE) House prices (SCE, MSC), Real Consumption (ECB), Nominal Expenditure (ECB).

We use the method of moments quantile regression method of Machado and Santos Silva (2019). This enables us to continue to include our time-fixed effects, and is simple to implement with clustered standard errors. Therefore the regression specification of (1.2) remains the same. This quantile regression method imposes a linear parametric assumption on the conditional distribution. Quantile τ is parameterised as: $Q(\tau|\boldsymbol{X}) = \boldsymbol{X}(\boldsymbol{\beta} + \boldsymbol{\gamma} q_\tau)$, where q_τ is the quantile of the latent error term, and $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ are parameters to be estimated. This is a "location-scale" model, and common within panel quantile regression methods. We therefore calculate marginal effects at the median as: $\boldsymbol{\beta}(x) + \boldsymbol{\gamma} \cdot q_{50}(x)$. The interpretation of this marginal effect is simple enough: it compares the median forecast error between the top and bottom income groups, relative to the sample median. In keeping with the previous set of results, we display the marginal effect multiplied by negative one. A higher value is therefore consistent with poorer households being more pessimistic than more affluent households.

In Table 1.6 we display these results. Again, for all variables, the effect of income on the relative median forecast is pronounced. The magnitudes are roughly in accordance with our measure at the mean. If anything, this has increased the quantitative effect of income on inflation pessimism. It is also reassuring for our hypothesis that for the vast majority of estimates, the sign has not changed either - forecast pessimism almost always decreases with income, whether measured at the mean or the median. Relative to Table 1.5 the sign of 3 of the 23 estimates have changed. These are SCE interest

Table 1.6: High vs. Low Income Forecast Error Magnitudes, Quantile Regression

	SCE		MSC		ECB	
Inflation	0.705	(0.058)	1.419	(0.057)	-0.828	(0.093)
Inflation (D)	0.542	(0.121)			1.807	(0.222)
Unemployment	0.054	(0.020)	0.057	(0.005)	0.839	(0.067)
Int. Rates	-2.063	(0.252)	0.043	(0.004)	1.401	(0.127)
Real Cons.					-1.585	(0.266)
Expenditure					0.374	(0.133)
Stocks	0.063	(0.010)				
Gov. Debt	0.710	(0.093)				
House Pr.	-1.060	(0.129)	0.085	(0.014)		
House Pr. (D)	-0.216	(0.073)				
Rental Pr.	0.596	(0.114)				
Food Pr.	0.314	(0.042)				
Medical Pr.	0.237	(0.026)				
Gas Pr.	0.072	(0.022)	0.068	(0.019)		

This table reports the marginal effects of being in the highest income group relative to the lowest across the method of moments quantile regression. The specification remains the exact same as in as taken from 1.2. The rows show different expectation variables, and the columns differing surveys. The standard errors are reported in parentheses, and are clustered at the individual level. Entries are left blank in surveys for which the relevant expectation is not elicited.

rates, ECB-CES inflation and ECB real consumption. This last estimate switches sign because of the inflation estimate switches, as nominal expenditure remains large and positive. This suggests that there is a strong correlation between the skew of interest rate forecast errors in the SCE and point implied forecast errors in the ECB - as defined in this paper - and income.

1.6 Individual Forecasting Responses to Income Changes

The previous three sections each illustrated one of our three results. The evidence in Section 1.4 on individual outcomes - real and nominal expenditure in the ECB survey - highlight that the *aggregate* patterns of forecast pessimism across income apply to forecasts of *individual* circumstances. In this section we examine the MSC data to explore whether there is evidence for the relationship between forecast pessimism and income holding at the level of the individual itself. That is to say, do changes in household income precipitate changes in the degree of forecast pessimism?

To do so, we calculate an (albeit imperfect) proxy for income shocks and explore whether individual's forecasting errors show similar patterns after an unexpected income shock. We use a simple difference-in-difference specification and find meaningful effects of income changes on forecasting errors: those consumers who experience positive unexpected income shocks over a period of 6 months display lower forecasting errors than those whose income moves as expected. Similarly, those who face negative income shocks display larger forecasting errors after an unexpected income shock.

Unfortunately as a continuous measure of income is available in neither the SCE or the ECB-CES, we are limited to using the MSC. Nonetheless, the MSC does have a limited panel dimension[50]. A subset of respondents respond to the survey twice, 6 months apart. We therefore calculate a measure of the real income shocks they face. Consumers report their income over the previous 12 months and their expected income growth over the coming 12 months in percentage terms. In order to compare individuals who respond six months apart, we annualise their income changes. The real change in annual income is calculated as $\Delta y_{i,t}^{\text{ann}} = 2 \cdot \log\left(\frac{y_{i,t}}{y_{i,t-6}}\right)$, where real income ($y_{i,t}$) has been deflated using the CPI. We therefore calculate our income shock variable as:

$$\text{shock}_{i,t} = \Delta y_{i,t}^{\text{ann}} - \mathbb{E}_{i,t-6}\left(\log\left(\frac{Y_{i,t+6}}{Y_{i,t-6}}\right) - \pi_{i,t+6}\right) \tag{1.3}$$

In words: the shock in period t for the individual i is calculated as their (realised) annualised real percentage income change, above their expected real income growth 6 months ago[51]. We choose percentage changes (i.e. relative) rather than absolute (i.e. dollar) changes as a definition of a shock variable simply because it is consistent with the MSC's elicitation of income expectations[52]. We present a nominal version of the income shock variable - where we do not deflate using CPI or inflation expectations - later in this section.

Naturally in such an environment, we may be particularly concerned about the risk of input error biasing our shock measure. To this end, we continue to trim expectations and reported income at the 5th and 95th percentiles. We also exclude all shocks which imply unexpected income shocks greater than doubling or less than halving. Finally, we categorise our shock variable and pick a threshold of 10% changes in absolute magnitude[53]. It is for this reason that we view our measure of a shock not as a 'pure' income shock per se, but rather as a proxy[54]. Therefore, we define two categorical variables. We denote $\text{shock}_{i,t}^{+} = \mathbb{I}\left(\text{shock}_{i,t} > 0.1\right)$ as a binary variable for a positive shock, and $\text{shock}_{i,t}^{-} = \mathbb{I}\left(\text{shock}_{i,t} < -0.1\right)$ for a negative shock. A respondent who expects real income to fall by 5% in the coming year, but who in fact

[50]This is something that is sometimes over-looked in the literature. A large plurality of observations in the MSC are repeat respondents - albeit only for the second time. This is possibly obscured by the fact that the off the shelf MSC does not come in a panel format, and some data-cleaning and manipulation is required.

[51]Because the observations are only 6 months apart, naturally this requires an assumption on the timing of income changes: on average within the MSC, individuals do not consistently expect their income to change *only* in the 6-12 month interval after their response. Given the long panel and the monthly frequency, this is very unlikely to be the case.

[52]A relative rather than absolute shock definition is more consistent with the logic of the most typical utility specifications too (i.e. log, CRRA).

[53]This cut-off is arbitrary, but the results are robust to alternative choices.

[54]Future work may be able to identify "true" unanticipated income shocks by using natural experiments, such as unexpected tax rebate, in the tradition of Johnson et al. (2006) or as Sahm et al. (2010) does in the MSC to measure marginal propensities to consume. Whilst information on the household's exposure to the 2001 or the 2008 US tax rebate fiscal stimulus programmes is available in the MSC, its relatively low frequency and short panel structure make it unsuitable to investigate the effect on inflation expectations. Work investigating the true causal effect of income shocks on expectation formation would require identified exogenous, unanticipated shocks over a much longer time horizon than currently available.

reports an *increase* of 5% after annualisation would therefore be categorised as the marginal recipient of a positive income shock.

Across the full sample of MSC double-respondents, who twice submit valid responses for income, income expectations, and inflation expectations, we are left with approximately 152,000 different households. Of these, 37,320 are classified as having not received an income shock, 22,384 as receiving a negative income shock, and 44,1378 as receiving a positive income shock[55]. For more detail on the definition of the shocks and summary statistics, please see Appendix A.7.

We use this shock variable in a standard two-way fixed effect Difference-in-Difference setting[56]. The only departure our specification takes from the textbook treatment is that we measure the effects of two (mutually exclusive) shocks. Our regression specification is therefore of the following form:

$$\mathbb{E}_{i,t-12}(\pi_t) - \pi_t = \beta_0 + \beta^+ \cdot \text{treat}_{i,t} \times \text{shock}^+_{i,t} + \beta^- \cdot \text{treat}_{i,t} \times \text{shock}^-_{i,t} + \lambda_i + \delta_t + \varepsilon_{i,t}$$

We use a two-way fixed effect specification, with clustered standard errors at the individual level. The variable $\text{treat}_{i,t}$ therefore takes the value 1 during the second survey response per individual, and 0 in the first (i.e. is an indicator variable denoting the treatment period). As before, the time-fixed effects will capture aggregate conditions and the individual effects any underlying idiosyncratic differences in forecasting. The effect of an unexpected real income shock on inflation forecasting errors is identified using within-date and within-individual variation. Our coefficients of interest are therefore β^- and β^+, which here represent the average effect of treatment on the treated (ATT) for the positive and negative shocks, respectively. In our baseline specification, we find that households suffering negative shocks have inflation forecast errors that are 0.23% higher than those who experienced no shock within the same period, on average. Households experiencing positive shocks, on the other hand, reduce their inflation forecast error by -0.29%. Both estimates are significant at the 1% level.

For the sake of comparison, we also include the specification in which we consider the shock to be continuous - i.e. do not classify a household as receiving a shock or not based on a particular threshold, but rather take the shock as verbatim. This is shown in the second column of Table 1.7. The estimated coefficient has the expected sign, although it is quantitatively small: for a 10% annualised unexpected

[55]That more households receive positive shocks than negative shocks should come as a no surprise. As discussed later in the paper, households routinely underestimate their future real income.

[56]Despite working in an environment with heterogeneous timing of treatments, we need not employ the techniques of the recent dynamic difference-in-difference literature - as we only have two periods (pre- and post-treatment for each individual in the panel), there is no risk of negative weighting of pre- or post-treatment periods at the individual level. See Roth et al. (2023) for a review.

increase in income, the respondents' inflation forecast error falls by 0.04pp.

Table 1.7: Difference-in-Difference Regressions

	Baseline		Nominal		Anticipated
	Discrete	Continuous	Discrete	Continuous	Discrete
-ve shock	0.227***		-0.0624		-0.0402
	(5.11)		(-1.43)		(-0.92)
+ve shock	-0.288***		-0.0922**		-0.0907**
	(-7.47)		(-2.38)		(-2.46)
shock (cont.)		-0.402***		-0.0827*	
		(-8.99)		(-1.85)	
Obs	105,184	105,184	105,538	105,538	117,170
Adj. R^2	0.465	0.465	0.462	0.462	0.458
Mean dep var	.516	.516	.52	.52	.541
First obs	1980m7	1980m7	1980m7	1980m7	1980m7
TWFE	✓	✓	✓	✓	✓

Coefficient estimates and t-statistics of β^+ and β^- across five difference-in-difference specifications conducted in the MSC. They are as follows: column one shows the baseline regression; column three shows the same specification but employs the nominal shock definition instead; column 5 uses the anticipated definition of a shock. Columns 2 and 4 use the "continuous" version of baseline and nominal shocks, where we allow the intensity of the treatment to vary across households. Standard errors are clustered at the individual level, and survey weights are employed. For more on the precise definitions of each shock variable, please see Appendix A.7.

At this point, none of these results is unequivocally "causal" in any sense. Indeed, a concern with this specification may be that we have introduced a mechanical correlation of the shock variable with the inflation forecasting error, as the extent to which a real income shock is unanticipated *itself* depends on an individual's ability to forecast inflation changes correctly. Those respondents who expect inflation to be persistently higher than is realised would also expect real income to be lower in the future, ceteris paribus, and are therefore more likely to experience negative 'shocks' as constructed here. In these cases, negative shocks would be positively correlated with higher forecast errors and thus produce a more positive estimate β^-. Such a concern firmly pushes this specification into the realm of correlation, rather than causation.

To at least take a step towards causality - without making any claims that our finds are indeed causal - we employ a specification which relies solely on nominal income shocks and lean on the presence of time-fixed effects to capture movements in the aggregate price level. In this alternative specification of our shock, we construct our income shock as unanticipated changes in *nominal* income.

$$\text{nominal shock}_{i,t} = \Delta Y_{i,t}^{\text{ann}} - \mathbb{E}_{i,t-6}\left(\log\left(\frac{Y_{i,t+6}}{Y_{i,t-6}}\right)\right) \tag{1.4}$$

We continue to use a 10% cut-off for positive and negative shocks each[57], and employ the same specifi-

[57]Such that $\text{nominal shock}^+_{i,t} = \mathbb{I}(\text{nominal shock}_{i,t} > 0.1)$ and $\text{nominal shock}^-_{i,t} = \mathbb{I}(\text{nominal shock}_{i,t} < -0.1)$.

cation as outlined above. The average treatment effects of the nominal shock on the treated are shown in columns 3 and 4 of Table 1.7.

We find that the ATT of the unanticipated negative nominal income shock is statistically insignificant at the 95% level. Positive income shocks are significant, although with a coefficient roughly 40% of the magnitude of the baseline specification. Following an unanticipated nominal income "shock" - proxied by the our self-reported measure - the inflation forecast error of households drops by on average 0.1%.

In the final column of Table 1.7 we use the final alternative definition of our shock variable. We consider the case in which we do not adjust for the expectations of changing household income at all. We define the (possibly) "anticipated" shock using the households income alone, and define a shock as having occurred simply if the respondent's (nominal) income changes sufficiently. This does little to change the analysis. More detail on the construction of this variable can be found in Appendix A.7. The results using this shock are remarkable consistent with the "nominal" shock results using the discrete measure.

Naturally, in the context of our previous results, we may be interested in whether the effects of an unanticipated change in income are correlated with income itself. We investigate this in Appendix A.7. While the point estimates for the response of expectations to a positive shock are mildly increasing in income (i.e. the forecast error falls post-shock as income increases), the estimates are too imprecise to distinguish the effect across income groups statistically. Given our, at best, noisy measure of a shock and limited panel dimension, this is perhaps unsurprising.

While our definition of an individual level shock is imperfect - in that it relies on self-reported income, which requires annualisation - we find that at the individual level unanticipated income differences are correlated with movements in forecast errors. As households experience negative shocks, their forecast errors increase, whereas those receiving positive shocks decrease their inflation forecast error. There is some, albeit limited, evidence that this effect reduces as income increases. In specifications which reduce the concern of mechanical correlation, a similar pattern is observed, although the strength of the responses are more muted. Although we do not claim correlation, this evidence is at the very least consistent with the notion that pessimism and pessimistic forecasting decreases with income. This evidence suggests that pessimists do not sort themselves across the income distribution, rather it is something about having less income which induces pessimism itself. Put another way, individuals may reduce their forecast errors (i.e. become less pessimistic) as their income increases.

1.7 Conclusion

This paper makes three contributions. First, it establishes that a strong predictor of the common component in consumer expectations in general is income. Second, it establishes - using evidence from the MSC, the SCE, and the CES - that: i) persistent forecast errors exist across a number of other economic variables, not just inflation; ii) that the direction of these errors always points in a direction consistent with a notion of forecast pessimism; iii) that forecast pessimism declines with income. Third, it shows that these beliefs matter quantitatively: pessimistic households over-predict inflation, unemployment, and the probability of adverse personal-finance events by economically meaningful margins. Forecast errors of 1–1.5 percentage points for inflation, and up to 8 pp for unemployment probabilities, are common in the bottom decile. Moreover, as Section 1.5 shows, the differences *between* income groups are meaningful. Finally, we present suggestive evidence that income and expectations are linked at not only the aggregate scale, but at the level of the individual. Panel regressions reveal that when a household experiences an unanticipated positive income shock, its one-year-ahead inflation forecast error declines by roughly 0.1 pp.

The results carry two policy implications. First, they caution against reading "average" household expectations as representative. To the extent that income is correlated with the marginal propensity to consume, then an equal-weighted average of expectations obscures the fact that those with the largest MPCs also hold the least accurate beliefs. Second, they suggest that communication strategies aimed at lower-income households could mitigate belief distortions and enhance the effectiveness of stabilisation policy. The importance of addressing consumer pessimism rises as the magnitude of aggregate shocks increase. Aggregate shocks whose incidence fall particularly on the lowest income groups are likely to have outsized effects on inflation expectations, and aggregate demand.

Several paths for future research remain open. The obvious question remains: if lower-income households *are* more pessimistic when forming their expectations of the future, why might this be? In Chapter 2 we examine a possible explanation common in the literature: that households face Knightian uncertainty and are ambiguity-averse. On the empirical side, data which can link at the individual level any two of: household expectations, granular household consumption and high-quality wealth data would greatly benefit not just this paper, but a plethora of adjacent topics too. On the modelling side, partial equilibrium analyses of the implications of expectation heterogeneity and/or pessimism would be of interest. In Chapter 3 we explore the effects of such beliefs on portfolio allocation and wealth accumulation, but this is only one possible direction to take. Finally, a heterogeneous agent general-equilibrium treatment that could match the profile of household expectations in the cross section would be of use to the literature. However the challenge of combining the "correct" behavioural biases, heterogeneous agents, and general equilibrium is no small task.

A Supplementary Material for Chapter 1

A.1 Data Description

In order to investigate the cross-sectional dispersion of pessimistic expectations, a number of surveys of household expectations are used. This section describes the variables used in greater detail.

Controls

In our baseline specification, we include a number of controls. As the definitions of these controls vary slightly between surveys, in this section we detail the definition of these controls. The baseline controls are:

- **SCE:** Gender: records if the respondent is male or female. Education: records if the respondent's highest educational attainment was high school, some (but not 4 years) of college, or a college degree or higher. Race: records the race of the respondent across 6 possible groups. Numeracy: takes the value 0 or 1 depending on the number of financial literacy questions were answered correctly: these are Lusardi and Mitchell (2014) "Big-3", and question respondents on interest rates, inflation, and diversification. Marital Status: records if the respondent is married or not. Job Status: records the number of jobs between the respondent (and possibly their partner). Indicator variable on whether there are 0, 1, 1.5 (i.e. FT + PT) or 2 jobs in the household. Age and squared age are included. State of Residence. Survey tenure

- **MSC:** Gender: records if the respondent is male or female. Education: records if the respondent's highest educational attainment was some (but did not complete) high school, high school, college, or a post-graduate degree or higher. Region: records the census region (i.e. West, North Central, Northeast or South) of the respondent. Age and squared age.

- **ECB-CES** Gender: records if the respondent is male or female. Education: records if the respondent's highest educational attainment was primary or lower secondary, upper secondary, or tertiary (i.e. university) education. Marital Status: records if the respondent lives with a partner or not. Job Status: records if the respondent is working or not. Age: grouped as 18-34, 35-54, 55-70, 71+. Country of residence. Survey tenure.

Variable Construction

We first detail the precise series used to construct the forecast errors. These are presented in Table A.1.

Table A.1: Forecast Error Construction Details

Survey	Name	Expectation Measure	Outcome Series	Outcome Series Description	Inverted?	5% trimmed?	Notes
SCE	Inflation (Point)	Q8v2part2	CPIAUCSL	BLS CPI		by month	
	Inflation (Density)	Q9_mean	CPIAUCSL	BLS CPI			
	House Prices (Point)	Q31v2part2	CSUSHPISA	Case Shiller Home Price Index	yes	by month	Estimates use subsample of home owners only
	House Prices (Density)	C1_mean	CSUSHPISA	Case Shiller Home Price Index	yes		Estimates use subsample of home owners only
	Rental Prices	C4_5	CUSR0000SEHA	BLS CPI Rent of Primary Residence Price Index		by month	Estimates use subsample of non-home owners only
	Unemployment	Q4new	UNRATE	BEA Unemployment Rate			Respondents asked to report the probability with which they expect the unemployment rate to be higher next month. This is measured against a variable which takes the value 1 if the outcome is higher, and 0 otherwise.
	Interest Rates (saving)	Q5new	SNDR	FDIC National Deposit Rates: Savings	yes		Respondents asked to report the probability with which they expect interest rates on savings accounts to be higher next month. This is measured against a variable which takes the value 1 if the outcome variable is higher, and 0 otherwise.
			SAVNRNJ before 2021	FDIC Average rate paid by all insured depositary institutions			
	Gov. Debt	C3_part2	GFDEBTN	US Treasury Fiscal Service: Total Federal Debt		by month	GFDEBTN is a quarterly series, so we interpolate to create a monthly series
	Gas Prices	C4_1	GASREGW	US Energy Information Administration US Regular All Formulations Gas Price		by month	
	Food Prices	C4_2	CPIUFDSL	BLS CPI Food Price Index		by month	
	Medical Prices	C4_3	CPIMEDSL	BLS CPI Medical Care Series		by month	
	Stock Prices	Q6new	NASDAQCOM	Nasdaq Index	yes		Respondents asked to report the probability with which they expect the average stock price in the US market to be higher next year. This is measured against a variable which takes the value 1 if the outcome variable is higher, and 0 otherwise.
ECB	Inflation (Point)	c1120	3 letter country code + CPIALLMINMEI	OECD CPI (by country)		by country-month	
	Inflation (Density)	c1150_1-8 c1152_1-10	3 letter country code + CPIALLMINMEI	OECD CPI (by country)			
	Unemployment	c4031	LRHUTTTT + 2 letter country code + M156S	OECD Unemployment rate (by country)		by country-month	
	Household Spending (prev. 12m)	point: c6020 bracketed: c6030	*		**	by country-month	*Expected real consumption measured as expected nominal expenditure growth minus expected inflation (trimmed). Actual real consumption growth measured as reported nominal expenditure growth (previous year) minus realised inflation. Forecast error measured as expected real consumption growth minus actual real consumption growth.
	Household Spending (expected, 12m ahead)	point: c6120 bracketed: c6130	*		**	by country-month	**We invert our real consumption forecast error, and invert expenditure growth for the sake of comparison.
	Interest Rate (Mortgages)	c5111 (pre 06/22) c5113 (post 06/22)	MIR.M. + 2 letter country code + .B.A2C.A.R.A.2250 .EUR.N	ESCB average mortgage rate within a given country. Accessible at the ECB's data-portal.		by country-month	
MSC	Inflation (Point)	PX1	CPIAUCSL	BLS CPI		by month	
	Interest Rates (borrowing)	RATEX	MORTGAGE30US	Freddie Mac 30Y Fixed Rate Mortgage US Average			RATEX recoded to take the value 1 if borrowing rates expected to increase, 0 if borrowing rates expected to stay the same, and -1 if borrowing rates expected to decrease. Measured against a variable that takes the value 1 if actual mortgage rates do change greater than x, 0 if mortgage rates change is in [-x,x] and -1 if mortgage rates change less than -x. Here, x is 0.91 (0.5 std deviations of the monthly mortgage rate index between 1990 and 2019 inclusive).
	Unemployment	UNEMP	UNRATE	BEA Unemployment Rate			UNEMP recoded to take the value 1 if unemployment expected to increase, 0 if unemployment expected to stay the same, and -1 if unemployment expected to decrease. Measured against a variable that takes the value 1 if actual unemployment rate change is greater than x, 0 if unemployment change is in [-x,x] and -1 if unemployment change less than -x. Here, x is 0.8 (0.5 std deviations of the monthly unemployment rate between 1990 and 2019 inclusive).
	House Prices	HOMPX1	CSUSHPISA	Case Shiller Home Price Index	yes	by month	Estimates use subsample of home owners only
	Gas Prices	GAS1	GASREGW	US Energy Information Administration: US Regular All Formulations Gas Price		by month	

This table details the exact series that were used to construct the measures of forecast errors used throughout Section 1.4. For each survey we detail the expectation measured used, as well as the outcome variable these are assessed against. The outcome variables are all taken from FRED. We also included miscellaneous detail, such as if we have inverted the measured (such that a higher value is consistent with consumer "pessimism"), if we have trimmed the variable at the 5-95% level and additional notes. Please see the appendix section on interest rates for the details of constructing the interest rate forecast error in the ECB-CES.

A.2 Sentiment Appendix

This section supports the results of Section 1.3. We first provide the regression output supporting Figure 1.2. This is shown in Table A.2.

Table A.2: MCA Demographics Regression

	MSC		SCE	
Middle Income	-0.152***	(-21.45)	-0.121***	(-6.09)
Highest Income	-0.290***	(-35.97)	-0.195***	(-8.72)
Age	0.0236***	(23.83)	-0.000279	(-0.08)
Age²	-0.000158***	(-16.25)	0.000132***	(3.30)
Married	0.00953	(1.53)	0.0410*	(1.93)
Female	0.195***	(35.39)	0.142***	(10.60)
Education = 2	-0.0753***	(-5.85)	0.00996	(0.45)
Education = 3	-0.205***	(-16.20)	-0.105***	(-4.86)
Education = 4	-0.245***	(-17.53)		
LF: 1 Job			-0.0458*	(-1.77)
LF: 1.5 Jobs			-0.0520	(-1.49)
LF: 2 Jobs			-0.0324	(-1.05)
Race: Black			0.0854***	(2.86)
Race: Asian			0.138***	(4.48)
Race: Other (Non			0.134***	(4.40)
Low Numeracy			0.121***	(6.85)
Obs	136,443		43,030	
Adj. R^2	0.115		0.083	
Mean dep var	.00337		.000502	
First obs	1980m2		2013m6	
Time FEs				
Area FEs	✓		✓	

Regressions of the first component of the MCA analysis regressed on demographic covariates. The first column shows the MSC results, the second the SCE results. These are regression supporting 2.3 and A.2. For the education regressors, the categories are: 1 = < High School, 2 = High School, 3 = College, 4 = Post-Grad, in the MSC; 1 = High School, 2 = Some College, 3 = College, in the SCE. We omit the geographic fixed effects from the table for the sake of space. We also additional control for survey tenure in the SCE. All t-statistics are calculated using standard errors clustered at the household level. Survey weights are employed in all regressions. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

SCE. Here we perform similar exercises to those performed in the main body, but instead on the SCE dataset. We find much the same results.

In the SCE, as in the MSC, we include only forward-looking expectations. The only significant difference to the MSC is that many of the questions in the SCE elicit continuous responses. These are therefore binned into terciles (within each survey month). As before, these responses are trimmed at 5-95% level. In Figure A.1 we report the loadings on each question.

In terms of MSC, the MCA analysis shows a remarkable strong correlation between consumer expectations in various outcomes. The first factor explains 71.8% of the variation in the beliefs of the respondents,

Figure A.1: MCA Loadings on Expectations in the SCE

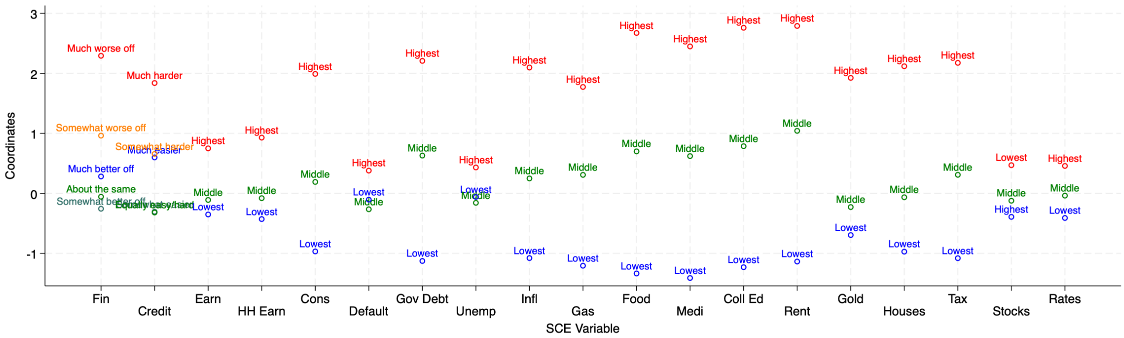

Each point represents the estimated loading of the first component for a categorical response in the MCA. Each variable on the x-axis represents a different question from the SCE. They are as follows: financially better off 12 months from now (Fin, Q2), Harder or easier for people to obtain credit or loans 12 months from now (Credit, Q29), expected earnings 12 months from now (Earn, Q23v2), Percentage growth in total household income 12 months from now (HH Earn, Q25v2), Percentage growth in total household spending 12 months from now (Cons, Q26v2), Probability miss debt payment over next 3 months (Default, Q30new), Percentage increase in level of government debt 12 months from now (Gov Debt, C3part2), Percentage chance unemployment higher 12 months from now (Unemp, Q4new), inflation over next 12 months (Infl, Q8v2part2), gas price inflation over next 12 months (Gas, C4_1), food price inflation over next 12 months (Food, C4_2), Medical price inflation over next 12 months (Medi, C4_3), College Education price inflation over next 12 months (Coll Ed, C4_4), Rent inflation over next 12 months (Rent, C4_5), price of gold inflation over next 12 months (Gold, C4_6), House price inflation over next 12 months (Houses, Q31v2part2), Percentage increases in total taxes in 12 months (Tax, Q27v2part2), Percentage chance average stock prices higher in US 12 months from now (Stocks, Q6new), Percentage chance average interest rate on savings accounts higher in US 12 months from now (Rates, Q5new). The corresponding question in the SCE has been included for convenience. Answers consistent with a simple notion of consumer "pessimism" have been colored red (and those with "optimism", blue).

and the second factor 7.94%. These numbers are very similar to those of Kamdar and Ray (2024).

To determine the demographic correlates of this factor, we perform a similar regression to that in the main body. We regress the first factor of the MCA on numerous demographic characteristics, as well as the income tercile of each respondent. Here we have additional controls available, and so we also control for the employment status of the household (where we group observations based on the number of jobs held by either the respondent or their spouse), their race, and their numeracy. This latter measure is determined by a series of questions testing financial literacy - see Armantier, Bruine de Bruin, Potter, et al. (2013) and Lusardi and Mitchell (2014) for more details. Time fixed-effects are included in the regression as well. For the SCE we also control for survey tenure and state of residence, although we suppress the reporting of these coefficients in the plot.

Table A.2 shows the regression output in Column 2. Figure A.2 plots the coefficients of this regression. Once again, the most pronounced demographic correlate is income group. Factors such as gender and education show patterns similar to those in the MSC data. Interestingly, both non-white respondents and those with low-financial literacy also display more "pessimistic" sentiment measures.

Figure A.2: Marginal Effects of Demographics on MCA First Factor (SCE)

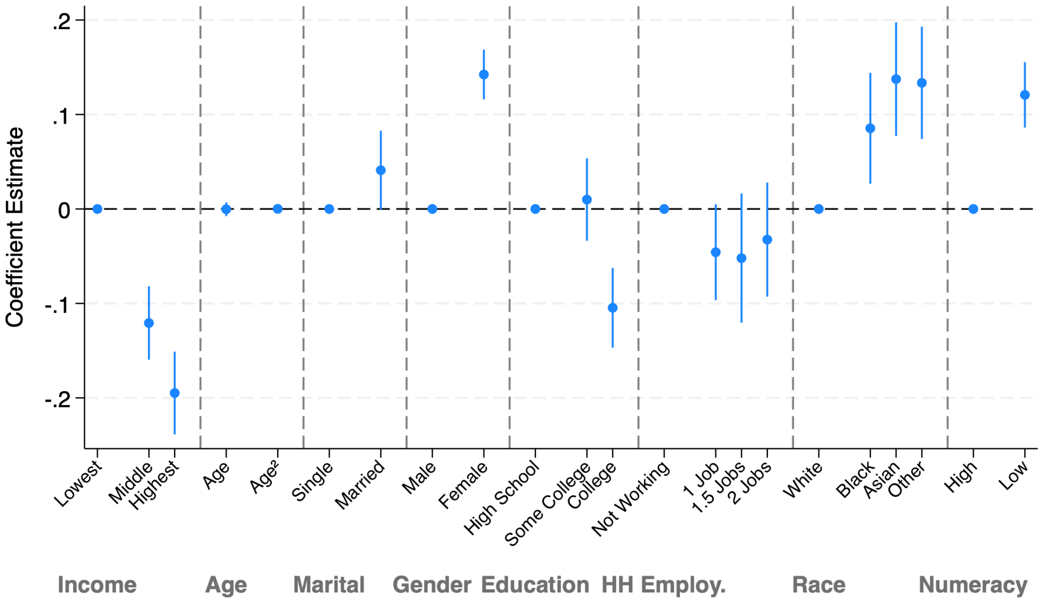

This graph plots the coefficients of the regression of the first MCA dimension on demographic characteristics, income tercile, household job status, financial literacy, state and survey tenure in the SCE. Each of the covariates listed is an indicator variable, with the expectation of age and age squared. The regression includes month fixed effects, as we as standard errors clustered at the individual level. Survey weights are employed and the number of observations is 43,030. The error bars denote the 95% confidence interval. The full regression output is in Table A.2.

A.3 Inflation

In this section, we include supporting details for Section 1.4.1.

Figure 1.4 illustrates average inflation forecasting errors in the US across both the MSC and the SCE for 2013 onwards, for the sake of comparison. Here we perform the same exercise on the MSC data, from 1990 onwards. This is shown in Figure A.3. It shows much the same cross-sectional pattern as the 2013- data, except the level of the inflation forecast errors has shifted downwards somewhat at the lower end of the distribution (from roughly a 2% average error to a 1.3% average error). This is consistent with Table 1.1.

Figure A.3: Inflation Forecast Errors by Income, MSC (1990-)

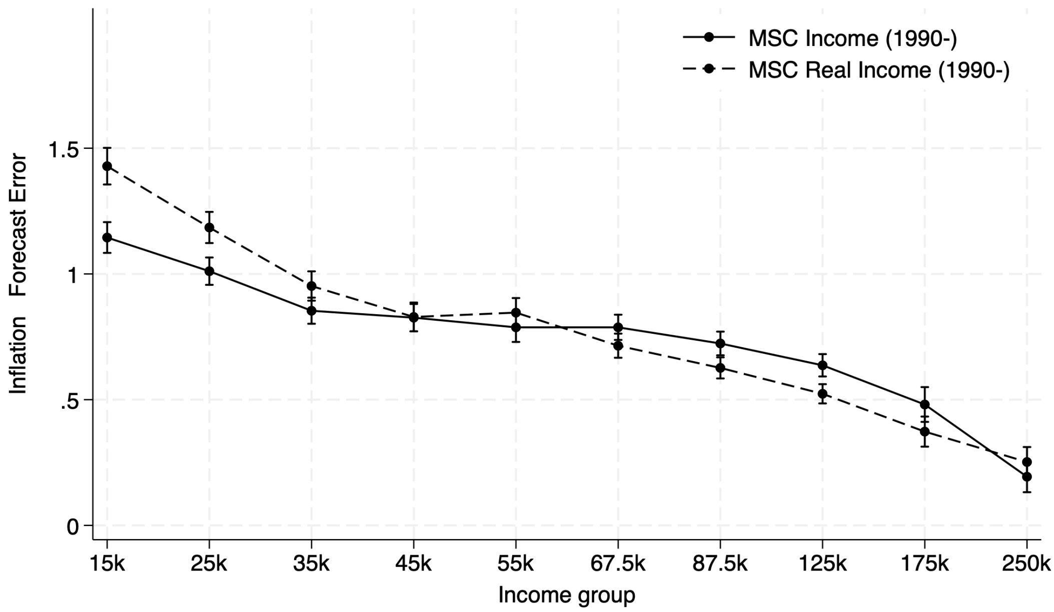

Average inflation forecasting error across income groups within the MSC from 1990 onwards. The solid lines are nominal income groups, and the dotted lines are in 2015 dollars. The x-axis label represents the mid-point of the income group. The graph employs survey weights. The standard errors are clustered at the individual level and the confidence intervals are at the 95% level.

In the main body of the paper, we plot the marginal effect of income group on the inflation forecasting error across the three surveys in Figure 1.5. For the sake of clarity we omitted the regression output table. Here we group households by three income bins: those earning less than \$50,000, those earning \$50,000-\$100,000 and those earning greater than \$100,000 per year before tax. These groups roughly correspond to income terciles in the US. The results of this regression are shown in Table A.3.

Table A.3: Regression Output of Inflation Forecast Error on Income

	SCE		MSC		ECB	
50-100k	-1.869***	-0.992***	-0.445***	-0.384***		
	(-18.34)	(-9.80)	(-20.91)	(-16.84)		
100k+	-3.060***	-1.502***	-0.736***	-0.638***		
	(-30.43)	(-13.00)	(-34.45)	(-25.95)		
20-40pct					-0.292***	-0.320***
					(-3.52)	(-4.21)
40-60pct					-0.395***	-0.504***
					(-4.64)	(-6.20)
60-80pct					-0.625***	-0.662***
					(-7.54)	(-8.12)
80-100pct					-0.751***	-0.859***
					(-9.45)	(-10.52)
Obs	145,048	144,089	179,487	177,263	693,625	662,014
Adj. R^2	0.130	0.200	0.215	0.217	0.250	0.332
Mean dep var	2.42	2.42	.836	.833	1.06	.965
First obs	2013m6	2013m6	1990m1	1990m1	2020m4	2020m4
Controls		✓		✓		✓

Regressions of inflation forecast on real income groups across surveys. Odd columns only include time-fixed effects, even columns include our vector of controls also. For the sake of comparison all expectations are measured using the (trimmed) point forecast. For the MSC, we use the period 1990 onwards. All t-statistics are calculated using standard errors clustered at the household level. Survey weights are employed in all regressions. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

Robustness

Cohort effects. As discussed in the paper, a robust finding in the literature is that inflation experience is a strong determinant of inflation expectations. To account for this possibility, we include an alternative specification where instead of including age and age squared as our controls, we instead include a cohort (i.e. birth year) fixed effect in our regressions. Given that we have such a small time dimension in the ECB-CES, we only perform this exercise for the US surveys. The results of this regression are shown in Table A.4.

Table A.4: Inflation Forecast Error Regressions, Controlling for Cohort

	SCE		MSC	
50-100k	-0.993***	-0.990***	-0.384***	-0.389***
	(-18.15)	(-18.17)	(-17.99)	(-18.29)
100k+	-1.503***	-1.483***	-0.638***	-0.639***
	(-24.59)	(-24.39)	(-27.87)	(-28.23)
Obs	144,087	144,087	177,263	177,263
Adj. R^2	0.200	0.203	0.217	0.218
Mean dep var	2.42	2.42	.833	.833
First obs	2013m6	2013m6	1990m1	1990m1
Cohort control		✓		✓

Here we provide an alternative specification of the baseline regression, where instead of controlling for age and age squared we control instead for cohort (i.e. year of birth). Columns (1) and (3) show the baseline specification, and Columns (2) and (4) replace the age controls with cohort controls. The standard errors are clustered at the individual level. Each coefficient is significant at the 1% level.

The results are almost identical; controlling for age or cohort does little to change the coefficients of interest.

Crisis periods. We may be concerned that the results are driven by periods of high inflation, and these periods alone, and that if we were to restrict our analysis to periods of more "normal" inflation, the differences across income groups would fade away. This would possibly be the case if households of different income groups reacted different during periods of high-inflation, but were formed their expectations in a similar manner to lower income households during "normal times". To investigate this possibility, we perform the same regressions as before, but this time excluding 2020 and 2021 from the sample. It is in these two years that forecast-errors are largest (although typically negative), due to the unforeseen inflationary shock of 2021 and 2022 that occurred in the wake of the Russian invasion of Ukraine. In Table A.5 we omit the years 2020 and 2021 from the sample.

While the average forecast error across the sample changes when the inflationary shock of 2020-2021 is excluded, the cross-sectional effect is left entirely unchanged. This suggests that the "pessimistic"

Table A.5: Inflation Forecast Error Regression, Excluding 2021-2022

	SCE		MSC		ECB	
50-100k	-0.993***	-1.036***	-0.384***	-0.387***		
	(-18.15)	(-18.21)	(-17.99)	(-17.83)		
100k+	-1.503***	-1.484***	-0.638***	-0.634***		
	(-24.59)	(-23.47)	(-27.87)	(-27.14)		
20-40pct					-0.363***	-0.389***
					(-10.98)	(-7.70)
40-60pct					-0.524***	-0.602***
					(-15.22)	(-11.47)
60-80pct					-0.662***	-0.770***
					(-19.13)	(-14.28)
80-100pct					-0.830***	-0.947***
					(-23.76)	(-17.30)
Obs	144,087	114,924	177,263	165,171	524,062	319,287
Adj. R^2	0.200	0.169	0.217	0.173	0.335	0.232
Mean dep var	2.42	3.09	.833	1.06	.453	2.47
First obs	2013m6	2013m6	1990m1	1990m1	2020m4	2022m1
Excluding 21-22?		✓		✓		✓

Columns (1), (3) and (5) replicate the baseline specifications as shown in Table A.3 across all three surveys. Columns (2), (4) and (6) perform the same regression (i.e. including demographic and time-fixed effects) but exclude the years 2020 and 2021 from the regression. Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

expectations are not an artefact of high inflationary periods, and extant during "normal" times too. Interestingly, in all cases, the R squared of the regression falls – this would suggest that during high-inflationary episodes, demographic factors (and time-fixed effects) explain more of the variation in forecast errors than during "normal" times.

Density forecasts. In the MSC, expectations are measured by asking respondents to report expected inflation as a single numerical response, i.e. to report a point forecast. In the SCE elicits both point forecasts and forecast probability densities. To achieve the latter, respondents are asked to assign probabilities to various possible inflationary outcomes. They are asked to assess the probability that the "rate of inflation" over the coming 12 months falls within 10 different intervals[58].The SCE follows Engelberg et al. (2009) and fits a generalised beta distribution to each respondent's stated histogram. We can then use the density mean as our measure of expected inflation 12 months ahead. One benefit of this measure over using the point forecast is that it makes interpersonal comparisons more straightforward. With a density implied measured there is no ambiguity over whether the respondent has reported their

[58] These intervals are: -12% or less, -12% to -8%, -8% to -4%, -4% to -2%, -2% to 0%, 0% to 2%, 2% to 4%, 4% to 8%, 8% to 12%, and 12% or more.

perceived mean, median, mode or some other characteristic of their perceived probability distribution as their point forecast. Furthermore, the density implied forecast is less susceptible to input error. Across both the SCE, the density implied mean is typically smaller in magnitude than the reported point forecast . It is therefore a more conservative measure of the bias. Indeed, some respondents in the MSC report point forecasts extremely large in magnitude, generating significant outliers. Similarly, the ECB-CES survey also elicits a probability density of inflation outcomes. As the ECB does not fit the distribution and calculate the mean themselves, we follow the same method as Engelberg et al. (2009) to fit the distribution and calculate the mean.

Table A.6: Inflation Forecast Error Regressions, Point vs. Density Implied Measure

	SCE		ECB			
50-100k	-0.993***	-0.0925***				
	(-18.15)	(-2.99)				
100k+	-1.503***	-0.300***				
	(-24.59)	(-8.42)				
20-40pct			-0.320***	-0.222***	-0.319***	-0.225***
			(-10.63)	(-4.79)	(-10.60)	(-4.82)
40-60pct			-0.504***	-0.510***	-0.499***	-0.508***
			(-16.13)	(-10.78)	(-16.00)	(-10.71)
60-80pct			-0.662***	-0.718***	-0.660***	-0.719***
			(-21.15)	(-15.00)	(-21.11)	(-14.99)
80-100pct			-0.859***	-1.336***	-0.856***	-1.335***
			(-27.20)	(-27.94)	(-27.14)	(-27.88)
Obs	144,087	140,485	662,014	645,792	662,014	645,792
Adj. R^2	0.200	0.283	0.332	0.174	0.369	0.199
Mean dep var	2.42	1.04	.965	2.82	.715	2.57
First obs	2013m6	2013m6	2020m4	2020m4	2020m4	2020m4
Point/Density?	Point	Density	Point	Density	Point	Density
Measure	CPI	CPI	CPI	CPI	HICP	HICP

Columns (1) and (3) replicate the baseline specifications as shown in Table A.3. Columns (2) and (4) use the density-implied one year inflation expectation to calculate the forecast error, rather than the point forecast. Columns (5) and (6) use the point and density-implied expectations respectively, but instead use the HICP rather than the CPI measure of country-level inflation to measure outcomes. Neither choice of point or density forecasts, nor CPI or HICP definitions of inflation change the results significantly. All specifications use the full set of demographic and time-fixed effects (as in the respective baseline specifications). Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

However, whether we use density measure or the point measure, the forecast error still declines in income. The results of the regression are shown in Table A.6. For the baseline specification, the difference between the top and bottom income groups is roughly 60% of the average forecast error in the SCE. When using the density implied expectation, this number declines to roughly 30%. For the

ECB-CES there is also a drop in the relative difference between top and bottom groups when using the density-implied forecast, but the magnitudes remain large - those in the top 20% of the income distribution have forecast errors that are roughly half those in the bottom 20%.

For the sake of robustness, in columns (5) and (6) we instead use the country-specific HICP measure of inflation (rather than the CPI) to measure inflation outcomes and calculate forecasting errors in the ECB-CES data. This too leaves the results unaffected (as should be expected with time and country fixed effects).

Long term forecasts. We may also be concerned that these findings are unique to inflation forecasts for one year and that differences between income groups disappear over longer-term forecasts. Both the MSC and the SCE also generate longer-term forecasts. In the MSC, respondents are asked for their expectations of average annual inflation over the next 5 years. In the SCE, they are asked for their expectation for inflation over the 2-3 horizon. Although longer-term forecasts are generated in the ECB-CES, we omit this from the exercises given the short-time dimension of the data currently. In the future performing the same exercise will be feasible. We calculate forecast errors as we do for the one year expectations, and show the results of the same regression in Table A.7.

Table A.7: Inflation Forecast Error Regression, Varying Forecast Horizon

	SCE		MSC	
50-100k	-0.993***	-0.926***	-0.384***	-0.330***
	(-18.15)	(-16.06)	(-17.99)	(-19.25)
100k+	-1.503***	-1.327***	-0.638***	-0.409***
	(-24.59)	(-20.93)	(-27.87)	(-22.75)
Obs	144,087	125,943	177,263	148,346
Adj. R^2	0.200	0.256	0.217	0.419
Mean dep var	2.42	1.54	.833	.731
First obs	2013m6	2013m6	1990m1	1990m4
Horizon	1y	2-3y	1y	5y

Columns (1) and (3) replicate the baseline specifications as shown in Table A.3 for the SCE and MSC respectively. Columns (2) and (4) instead use the 2-3 year and 5 year inflation expectations to calculate the appropriate forecast error, which are then used as the dependent variable in the regression. All specifications use the full set of demographic and time-fixed effects (as in the respective baseline specifications). Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

While the average forecast error across the entire sample reduces in both cases - consumers appear to be less pessimistic about further in the future, and are more accurate in their forecasts - the cross-sectional differences are largely unchanged. If anything, as the average forecast error decreases, relative to the population average, the effect of income becomes more pronounced.

Inflation Heterogeneity. As discussed in the main body of the paper, another concern we may have is that respondents do not report their expectation for US aggregate inflation (as specifically prompted to do so in the SCE), but rather report their expectation of inflation over the coming year across their *personal* consumption basket. Of course, we cannot observe the personal inflation rate experienced by members of the surveys, nor are we aware of a data-set that can do this[59].

Nonetheless, we attempt to proxy for the differential experiences of households by using the data of Hartley and Mejia (2024). The authors calculate income-decile specific CPI and attempts to very closely match the method of construction of the index as closely as possible to the BEA's construction. One advantage of this data is that it is also available at the quartile-regional level, and so also captures geographic variation in consumption differences. In their paper, they demonstrate that the geographic divergences in inflation experience are of a similar magnitude to income-related divergences as had been previously found in the literature. For more on the detail of the construction of these indexes, please see the original paper.

Table A.8: Inflation Forecast Regressions Using Income/Region Specific Inflation

	CPI	Decile-US	Quartile-Regional
50-100k	-0.992***	-0.939***	-0.961***
	(-9.80)	(-9.17)	(-9.37)
100k+	-1.502***	-1.487***	-1.490***
	(-13.00)	(-12.64)	(-12.54)
Obs	144,089	138,709	138,703
Adj. R^2	0.165	0.181	0.181
Mean dep var	2.56	2.38	2.35
First obs	2013m6	2013m6	2013m6
Controls	✓	✓	✓

Column (1) shows the results of the baseline specification as above. Column (2) uses an alternative dependent variable, where the inflation outcome is an income-decile specific measure calculated by Hartley and Mejia (2024). Column (3) uses an alternative specification once again, and uses an income-quartile by region measure. Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

Using this data (which the authors were kind enough to provide), we construct two alternative forecast-error measures. The first uses an income-decile specific measure of inflation to measure inflation expectations against[60]. The second uses quartile-region level inflation as the outcome variable[61]. We then perform the same regression as in the baseline specification on the SCE sample. The results of

[59] A possibility of course would be to link scanner data to a expectations survey.

[60] As we only have binned income in the SCE, we match households to deciles by taking the mid-point of the range reported in the SCE, and assigning the relevant disaggregated inflation series based upon the income cut-offs each year used by Hartley and Mejia (2024) to construct the series.

[61] Due to the small sample size of the CEX, the authors are unable to provide a greater level of disaggregation than this.

using these two alternative definitions of the dependent variable are show in Table A.8.

Although the aggregate population-level forecast errors reduce slightly, the cross-sectional effects are left almost unchanged. If anything, the heterogeneity in forecast errors across income groups increases somewhat (as a proportion of the population average). Naturally this is not a perfect measure of personal inflation experience, but the results appear to be robust to using one of the literatures best measures of disaggregated inflation. Of course, data that could link expectations directly to inflation experience would be of great interest to the literature.

Inflation Exposure. We may be concerned that the degree to which respondents display a systematic bias is driven by their inflation exposure. As Doepke and Schneider (2006) and Auclert (2019) highlight, the balance-sheet exposure of households to inflation varies substantially in the cross section[62]: this is the Fisher channel (Fisher (1933)). The net exposure of households to inflation is measured as their net nominal position (NNP). Households with negative NNPs gain when there are unexpected increases in the price level, and those with positive NNPs lose. If households are pessimistic in their expectation formation, we would expect those who benefit in inflationary states of the world – i.e. those with negative NNPs – to persistently underweight the probability of inflation and have lower systematic biases than those with positive NNPs. In the limiting case where households face risk through the Fisher channel only, we would expect negative NNP households to show expectations that are systematically *negatively* biased and those with positive NNPs to be *positively* biased.

To investigate this, we employ the additional SCE survey modules, which contain information on household's balance sheet positions. We run three additional regressions to investigate how the expectation error of the household varies with the composition of their balance sheet. We control for: 1) mortgage ownership, 2) mortgage ownership and stock of household debt, 3) a (limited) measure of NNP exposure[63] It should be noted that the SCE's balance sheet data is limited. We therefore restrict these regressions to the subset of households whose stock of debt or nominal assets/liabilities is less than three times their income, as measured by the mid-point of their bin[64]. Our sample is reduced for all

[62]Younger (and richer) households' balance sheets are primarily comprised of nominal mortgage debt and so benefit during bouts of inflation. Older (and richer) households hold much of their wealth in nominal assets – usually cash - and see their real wealth decline during inflationary episodes.

[63]We use the additional SCE module "Credit Access Survey" to construct these measures. Households are asked about various aspects of their balance-sheet. Mortgage ownership is directly questioned (N1_2). To construct a proxy for household debt, we sum credit card debt (N1_1), mortgage debt (N2_2), student loans (N2_3), home-based loans (N2_4), auto loans (N2_5) and other personal loans (N2_6). To construct a proxy for NNP, we take household debt excluding mortgages to proxy their nominal debt. We proxy nominal assets using the "Household Finance" additional module, and calculate it as the percentage of their assets in checking/savings accounts (d19_1), T-bills/CDs/MMFs (d19_2) and Bonds/Bond Mutual Funds (d19_4) and then times this by the total value of their savings and investments (d16new_1). We calculate NNP as nominal assets minus nominal debt divided by their income.

[64]We pick the ratio 3x income to remove outliers. Auclert (2019) estimates the maximum NNP asset/liability to income ratio found across either the PSID or BLS-CEX surveys was approximately 2.

three regressions, particularly case 3.

Table A.9: Inflation Forecast Regressions, Controlling for Inflation Exposure

	Baseline	(1)	(2)	(3)
50-100k	-0.351***	-0.245*	-0.333**	0.188
	(-3.34)	(-1.77)	(-2.34)	(0.51)
100k+	-0.754***	-0.606***	-0.748***	-0.249
	(-6.13)	(-3.54)	(-4.42)	(-0.50)
Owns Mortgage		-0.122	-0.283*	
		(-1.10)	(-1.72)	
log(Debt)			0.0688	
			(1.61)	
NNP				0.218*
				(1.76)
Obs	155,387	91,725	86,837	7,753
Adj. R^2	0.112	0.136	0.136	0.151
Mean dep var	1.39	1.21	1.24	1.74
First obs	2013m6	2013m6	2013m6	2013m9
Controls	✓	✓	✓	✓

Column "Baseline" denotes the baseline specification, with demographic controls and time-fixed effects. Column (1) adds mortgage ownership as a covariate. Column (2) additionally adds the log debt of the household. Column (3) instead uses the NNP of the household as a covariate. Here we have used the density-implied forecast to improve the precision of the estimates - as these are less prone to input-error - given the small sample of Column (3) in particular. Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

The regressions are shown in Table A.9. The NNP position of a household is strongly determined by mortgage ownership: in regressions 1) and 2), we see that those who own a mortgage have lower forecast errors of approximately 0.1 to 0.3%. In regression 3), we can see that those households with higher NNP positions typically have higher forecast errors. Both of these results are indicative of households with positive exposure to inflation risk having lower systematic biases. However, given the small sample size and unreliability of the data, these regressions should only be considered indicative: the estimates are only significant at the 10% level. Datasets of expectations with more reliable balance sheet information would be a useful resource for the literature.

Risk appetite controls. A further concern may be that instead of measuring a feature of a households expectation formation process, we may be picking up variations in their risk-aversion or their risk appetite. Indeed, if households are reporting their stochastic discount factor (SDF) weighted expectation rather than the expectation under the objective measure[65] then their forecasts of the future would be correlated with their individual risk-aversion.

To try to mitigate this possibility, in this section we include other covariates that may proxy for their

[65]In finance, these are sometimes referred to as Q, the "risk-neutral measure", and P, the "physical measure", respectively.

Table A.10: Inflation Forecast Error Regressions, Controlling For Risk-Aversion

	Baseline	(1)	(2)	(3)	(4)
50-100k	-0.351***	-0.350***	-0.176	-0.303*	-0.210*
	(-3.34)	(-3.35)	(-1.47)	(-1.88)	(-1.75)
100k+	-0.754***	-0.753***	-0.536***	-0.569***	-0.588***
	(-6.13)	(-6.15)	(-3.78)	(-2.81)	(-4.14)
Expectation Uncertainty		0.000112			
		(0.09)			
Financial Risk = 2			-0.407*	0.304	
			(-1.72)	(0.96)	
Financial Risk = 3			-0.871***	-0.302	
			(-3.66)	(-0.99)	
Financial Risk = 4			-0.819***	-0.125	
			(-3.50)	(-0.41)	
Financial Risk = 5			-0.897***	-0.233	
			(-3.83)	(-0.76)	
Financial Risk = 6			-1.064***	-0.466	
			(-3.75)	(-1.21)	
Financial Risk = Very Willing			-0.968**	-0.456	
			(-2.54)	(-1.05)	
Owns Mortgage				0.0919	
				(0.77)	
Exp. Income Growth				0.130***	
				(8.55)	
Risk Aver. Proxy (MCA)					-0.428***
					(-3.73)
Obs	155,387	155,387	120,164	64,956	120,144
Adj. R^2	0.112	0.112	0.097	0.138	0.096
Mean dep var	1.39	1.39	.87	.537	.869
First obs	2013m6	2013m6	2015m4	2015m4	2015m4
Controls	✓	✓	✓	✓	✓

Column "Baseline" denotes the baseline specification, with demographic controls and time-fixed effects. Column (1) includes the variance of the density implied expectation. Column (2) includes the households self-reported willingness to take financial risks (QRA1). Column (3) additionally includes mortgage ownership as well as the (density-implied) expected one-year income growth. Column (4) adds the first factor of an MCA analysis on self-reported willingness to take financial risks (QRA1) and "willingness to take risks in daily activities" (QRA2), to the baseline specification. As we use the density-implied variance of inflation expectations, our dependent variable here is the density-implied forecast error. Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

risk-aversion or more indirectly their SDF itself. To do this, we specify four extra regressions. In the first, we include the variance of the density implied expectation, to proxy for their uncertainty. In the second, we include another question from the SCE, in which respondents are asked to grade their willingness to take financial risks (QRA1) on a scale of 1-7. We believe this provides (a crude) proxy for their risk-aversion. In the third, we supplement this with whether or not they own a mortgage, as well as their expected nominal income growth over the coming year, to attempt to proxy for their expected future welfare across inflation outcomes. Finally in specification 4, we use an MCA analysis to extract the first factor determining their willingness to take financial risks as well as their "willingness to take risks in daily activities" (QRA2). The results of these regressions are shown in Table A.10.

Inclusion of the variance of inflation expectations does little to change the results (and the coefficient on the control is insignificant). Column (2) is more interesting: households who report being more risk-averse (i.e. report low financial risk tolerance) have higher forecasting errors. This is consistent with the notion that more risk-averse households place more weight on bad states of the world (high-inflation) if they are in fact reporting their SDF weighted expectation. However as this variable isn't strongly correlated with income, it doesn't change the cross-sectional effect of income materially. Attempting to proxy for their exposure to inflation (i.e. the utility term in the SDF, not just the curvature of utility) doesn't change the cross-sectional results materially either. The results of Column (4) are similar to that of (2) and (3), suggesting that general (rather than financial) risk appetite adds little to the analysis.

Again, while this is, of course, far from a perfect measure (indeed, a perfect measure would be hard to imagine for such an exercise), this may at least provide some reassurance that the results are not driven entirely by risk-aversion alone.

Measurement error and asymmetric expectations. In this section, we perform robustness exercises around the possibility of lower-bounded mismeasurement, to try to ameliorate concerns surrounding the possibility of the final case.

Why exactly might this be a problem for the results presented above? We might imagine this becoming a measurement issue if the variance of responses co-varies with our variable of interest – that is, income. If the cross-sectional variance of mean inflation expectations is higher for those at the bottom of the income distribution, but these responses are for whatever reason bounded below by zero, it would serve to mechanically generate higher average inflation forecast errors for those at the bottom than the top of the income distribution. Crucially, heteroscedasticity alone would not generate a biased result - only in combination with censored observations (whether actual, or in effect) does this create a source of bias.

In Figure A.4, we plot the densities of the inflation forecast errors conditional on income groups. It is certainly true that there is more dispersion in inflation forecasts at lower levels of income, and the effect is particularly pronounced for the point forecast measure. The densities are also almost always skewed too, with more weight placed on higher inflation outcomes than low inflation outcomes. However, the median forecast error within income groups is declining as income increases - it is not purely the skewness of inflation forecast errors driving the baseline regression results. Furthermore, both types of measures and all income groups there is non-trivial mass at zero inflation and under. Deflationary

Figure A.4: Inflation Forecast Errors Kernel Density Plot, by Income-Group

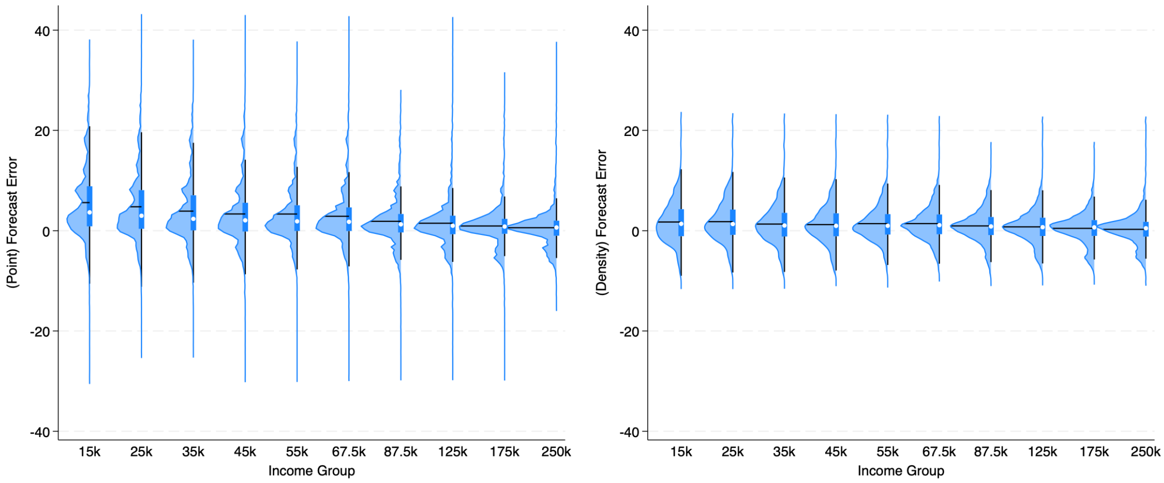

This graph plots the distribution of the forecast errors by income group (2015$) in the SCE (trimmed at the 5-95% level, by date). The left panel shows the point forecast, and the right panel shows the density forecast. The blue area is a kernel density of forecast errors The x-axis label represents the mid-point of the forecast errors, conditional on income group. The vertical black lines (whispers) represent 5-95th percentile range, and the boxes the inter-quartile range. The white dots represent the median. The densities employ survey weights.

expectations are less common, but nonetheless present across all income groups.

We focus now on the density-implied forecast error, which is less susceptible than the point forecast measure to bunching due to respondents preferring integer answers, or indeed input-error more generally. Across the entire SCE 43% of the respondents place at least some mass on the probability of deflation at any given time. Furthermore, 78% of respondents place at least some mass on the probability of deflation within their time answering the survey. There is therefore reason to believe that the density-implied measure is less susceptible to a zero lower bound measurement issue – most individuals report that they believe deflation is possible. It is reassuring that the cross-section patterns hold regardless of which forecast error definition we use. Nonetheless, we may be concerned that respondents report some probability of deflation because they are primed by the survey to do so, not because truly believe it to be a possible state of the world. Put another way, it may be that the variation in the density-implied measure is driven by the respondents changing the mass only on the positive part of the distribution, and leaving the probability that inflation is negative relatively unchanged.

To try to account for this, we perform two exercises. We run our baseline regression again, but over two sub-samples of the data. The first sub-sample is the group of respondents who ever report their median inflation expectation (as elicited by their subjective probability density) to be negative - this is roughly 27% of observations. This is a stringent sub-sample of the population, including only those who have - at least once in the survey to place more probability on deflationary state of the world. The second sub-sample is the group of respondents who are in the top 25% of respondents by variability in their probability assigned to deflation. This sub-sample represents 29% of observations, who display the largest changes in their expected probability of deflation. They are therefore the least likely to be mechanically responding with some weight on deflation purely because they have been primed to do. Granted, this is only a proxy for engagement with the question. The results of these regressions are shown in A.11.

Table A.11: Inflation Forecast Error Regressions, Robustness to Deflationary Response Sub-Samples

	Baseline	(1)	(2)	(3)	(4)	(5)
50-100k	-0.351***	-0.285*	-0.507***	-0.515**	-0.356***	-0.183**
	(-3.34)	(-1.65)	(-2.86)	(-2.49)	(-15.10)	(-2.02)
100k+	-0.754***	-0.365*	-0.668***	-0.569**	-0.617***	-0.491***
	(-6.13)	(-1.69)	(-3.13)	(-2.30)	(-23.53)	(-4.63)
Obs	155,387	44,611	40,945	33,705	77,557	155,387
Adj. R^2	0.112	0.080	0.073	0.073	0.723	
Mean dep var	1.39	-.401	.262	-.0665	.661	1.39
First obs	2013m6	2013m6	2013m6	2013m6	2013m6	2013m6
Controls	✓	✓	✓	✓	✓	✓

Column "Baseline" denotes the baseline specification, with demographic controls and time-fixed effects, using the density-implied forecast error. Column (1) performs the same regression over the subsample of respondents who ever report a higher probability of deflation than inflation during their time in the survey. Column (2) performs the same regression over the subsample of respondents whose weight on deflationary outcomes varies the most - as measured by being in the 75th percentile or above by standard deviation of probability assigned to deflation. Column (3) performs the same regression, but over the intersection of the two samples used in (1) and (2). Column (4) performs the same regression, but for only those households whose forecast errors are within the inter-quartile range. Finally, Column (5) performs a method-of-moments quantile regression, where *qtile* is evaluated at the 50th percentile (therefore the marginal effect on the median is given by: $\text{qtile}_i = \text{main}_i + \text{scale} \cdot q_{50}$, where q_{50} is the quantile of the standard errors at the 50th percentile. We report only qtile_i. Survey weights are employed and standard errors are clustered at the individual level. Two asterisks denote estimates significant at the 5% level, and three denote those significant at the 1% level.

Looking at column (1), even across the respondents who are particularly likely to report deflation (i.e., have at one point in their time in survey reported deflation as more likely than inflation), we still see the same qualitative effect of income on forecast error. Quantitatively, the effect between the top and bottom income groups is roughly halved. Among those respondents whose deflationary forecasts vary the most (so as to have perhaps engaged most with the question), the results are almost unchanged. This is also true for column (3), the subsample of respondents who adhere to both the criteria of columns (1) and (2). The results appear to be robust to the criticism that a "zero lower-bound" on expectations are driving the relationship between forecasting errors and income.

Columns (4) and (5) attempt to deal with the problem that the skewness of the inflation forecast is correlated with income more directly. In the subsample used in column (4) we completely remove all response which falls outside of the interquartile range for forecast errors (within each date). Again, this has little qualitative impact on the cross-sectional results. However, it does reduce the magnitude of the population average forecast error by roughly half, owing to the long upper tail of inflation forecasts. Finally, in Column (5) Finally, we also perform a method-of-moments quantile regression[66] to evaluate the marginal impact of income on the median forecast error, rather than the mean as in OLS. Otherwise, the specification remains the exact same. As this is a location-scale model, the marginal effect of the income group i on the median forecast error is: $\text{qtile}_i = \text{main}_i + \text{scale}_i \cdot q_{50}$, where q_{50} is median of the standardised error. Therefore, the middle income group has a median forecast error -0.183 lower than the lowest income group, and the upper income group -0.491. This is reported in the last set of results for the final column. This would indicate that, while the skew of the distribution of inflation forecasts does have an effect on our baseline estimates, the difference between income groups even at the median inflation forecast error remains. We continue to use the marginal effects on the mean in the main paper (and as our baseline) for its familiarity, and note that there is no intrinsic reason to prefer a median versus a mean (even when the latter is affected more by outliers).

[66]The basic location-scale quantile-regression model is: $Y_i = X_i'\beta + X_i'\theta\varepsilon_i$, where β and θ are referred to as the location and scale parameters respectively. See Machado and Santos Silva (2019) for more.

A.4 Unemployment

Each of the three surveys elicits expectations about unemployment in different ways. We therefore construct our measures of forecast errors as follows:

SCE: The SCE asks respondents for the "percentage chance that 12 months from now the unemployment rate in the US will be higher than it is now?". We use this to construct a measure of unemployment forecast error, where the expectation is measured against a variable which takes the value one if the BLS national unemployment rate is indeed higher 12 months in the future. This variable is therefore bounded in the interval $[-1\ ,\ 1]$.

MSC: In the MSC, respondents are asked whether they think there will be more, less or the same unemployment in the coming 12 months. We assign each of these outcomes a value of 1, 0 or -1 respectively. We measure this expectation against a variable that takes the value 0 if the future employment realised is within 1 standard deviation[67] of current unemployment, 1 if it is higher, and -1 if lower. This variable is therefore bounded in the interval $[-2\ ,\ 2]$. Of the three measures, this is the hardest to interpret, but nonetheless captures a sense in which expectations may display a systematic bias.

ECB-CES: In the ECB-CES, respondents report a point forecast of expected unemployment in 12 months in their country of residence. We therefore calculate the forecast error as standard[68].

For all of these measures, higher values are associated with more pessimistic expectations.

In Table A.12 we report the regression output of the simple specification (i.e. only including time-fixed effects) and the baseline specification estimated for each survey, across three income groups.

As discussed in the main body of the paper, all three surveys display a negative effect of income on unemployment forecast error. The easiest to interpret is the ECB-CES: here, the top income quintile expects roughly 2% lower unemployment than the bottom income quintile.

Perception errors. What if the forecast error of consumers is driven entirely by their *complete* inattention (i.e. to past, current, and future) aggregate economic conditions? Indeed, we may think that a respondents forecast error over unemployment is at least in part explained by the lower salience of the

[67] The standard deviation of the unemployment rate as measured over 1990-2020. Approximately 60% of the observations are within this interval throughout the sample of the MSC. The choice of interval size by which we assign unemployment to be "unchanged" does not affect the analysis meaningfully. Alternatively, one could use a more complex approach to estimate a suitable (perhaps time-varying) estimate of this interval, as in Bhandari et al. (2024). This does not affect the results either.

[68] We continue to trim observations at the 5% level by country-month.

Table A.12: Unemployment Forecast Error Regressions

	SCE		MSC		ECB		
50-100k	-0.0152***	-0.0236***	-0.0399***	-0.0381***			
	(-4.04)	(-6.00)	(-9.51)	(-8.41)			
100k+	0.00352	-0.0153***	-0.0634***	-0.0597***			
	(0.89)	(-3.17)	(-14.66)	(-12.14)			
20-40pct					-0.868***	-0.784***	-0.0585
					(-5.37)	(-5.33)	(-1.32)
40-60pct					-1.492***	-1.400***	-0.245***
					(-10.57)	(-10.08)	(-4.83)
60-80pct					-1.931***	-1.591***	-0.265***
					(-13.06)	(-10.95)	(-5.64)
80-100pct					-2.341***	-2.105***	-0.301***
					(-17.47)	(-13.95)	(-6.47)
$\mathbb{E}_t[u_t] - u_t$							0.946***
							(285.29)
Obs	159,337	158,221	207,282	204,422	691,659	660,075	645,063
Adj. R^2	0.744	0.751	0.377	0.379	0.027	0.179	0.808
Mean dep var	.166	.166	.172	.172	4.99	4.88	4.78
First obs	2013m6	2013m6	1990m1	1990m1	2020m4	2020m4	2020m4
Controls		✓		✓		✓	✓

Regressions of unemployment forecast errors on real income groups across surveys. The construction of each forecast error (and their differences) is detailed above. As they each differ in their interpretation, the quantitative estimates are not comparable across surveys. Odd columns only include time-fixed effects, even columns include our vector of controls also. The final column uses the same specification as column (6), but also includes the current unemployment rate error (i.e. the degree to which consumers are incorrect about the current unemployment rate). For the MSC, we use the period 1990 onwards. All t-statistics are calculated using standard errors clustered at the household level. Survey weights are employed in all regressions. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

unemployment rate relative to inflation. We have three comments to make on this. First, this need not be anathema to the possibility of pessimistic expectation formation: if the current unemployment rate is not within the individual's information set, it remains possible that individuals also form pessimistic expectations over current unemployment. Second, the salience of unemployment is likely to be strongest at the bottom of the income distribution, among those households who face the least job security and greatest income risk - our results demonstrate that if this effect is present, it is not sufficient to overturn the result. Third, across many European economies the salience of national unemployment is likely to be high given recent experience. Nonetheless, the final column of A.12 attempts to shed some light on this. A useful feature the ECB-CES dataset is that it contains information on the perceived current unemployment rate of households. I can therefore calculate a perception error using the perceived unemployment rate and the prevailing rate. This variable accounts for a (very) large proportion of the variance in forecasting error, unsurprisingly. It is striking that even accounting for this, top income quintiles have lower forecasting errors than bottom income quintiles by 0.3% on average.

A.5 Interest Rates

Each of the three surveys elicits expectations about interest rate changes in different ways. We therefore construct our measures of forecast errors as follows:

SCE: Respondents are asked whether *'the percentage chance that 12 months from now the average interest rate in savings accounts will be higher than it is now'*. We construct a forecast error measure using an indicator variable which takes the value one if higher nominal savings account rates materialises. It is assumed that higher savings account interest rates are perceived to be a benefit to the consumer, and so invert the measure such that higher readings are consistent with forecast pessimism.

MSC: Respondents are asked about their expectations for *"interest rates for borrowing money over the next 12 months"*. As for unemployment, they respond with whether or not they expect interest rates to increase, decrease, or stay the same. We therefore construct our forecast error measure in a similar manner as we do with unemployment. As this is a borrowing rate (rather than a savings rate), we do not invert this measure.

ECB-CES: Respondents are asked for a direct forecast. They are asked *'In 12 months from now, what do you think the interest rate on mortgages in the country where you live currently'*. We therefore construct a forecast error using average country specific interest rates[69]. Again, as this is a borrowing rate we do not invert this measure.

In Table A.13 we report the regression output of the simple specification (i.e. only including time fixed-effects) and the baseline specification estimated for each survey, across three income groups.

As discussed in the main body of the paper, all three surveys show a negative effect of income on the unemployment forecast error. The easiest to interpret is the ECB-CES: here, the top income quintile

[69] We construct these measures like so:

For the SCE we use the FDIC measure of the average savings account interest per annum as the outcome variable, downloaded from FRED. In 2021 the FDIC changed its publication of savings rates. We therefore use SNDR, and append SAVNRNJ before 2021.

For the MSC, respondents are asked about their expectations of borrowing rates. As 30 year fixed rate mortgages are commonly held (roughly 50% of households in the US own a mortgage, of which over 90% are 30-y fixed), and constitute the vast majority of household borrowing, we take the 30 year fixed rate mortgage to be the most suitable outcome variable. We use Freddie Mac's measure of average 30-y fixed interest mortgage rates (FRED code: MORTGAGE30US).

For the ECB-CES, we use ESCB data for average mortgage rates across European economies. The data can be accessed at the ECB's data-portal. For each country the series code is MIR.M.??.B.A2C.A.R.A.2250.EUR.N, where ?? is to be replaced by the 2-character country code.

Table A.13: Interest Rate Forecast Error Regressions

	SCE		MSC		ECB		
50-100k	-0.0273***	-0.00613	-0.0259***	-0.0262***			
	(-6.88)	(-1.47)	(-6.33)	(-5.94)			
100k+	-0.0848***	-0.0402***	-0.0356***	-0.0333***			
	(-19.52)	(-7.73)	(-8.75)	(-7.13)			
20-40pct					-0.177***	-0.137***	-0.102**
					(-3.83)	(-2.98)	(-2.16)
40-60pct					-0.331***	-0.290***	-0.263***
					(-7.56)	(-6.25)	(-5.58)
60-80pct					-0.491***	-0.402***	-0.368***
					(-10.77)	(-8.29)	(-7.44)
80-100pct					-0.695***	-0.593***	-0.555***
					(-16.59)	(-12.60)	(-11.57)
$\mathbb{E}_t[p_{t+12}^{\text{Homes}}]$							0.0316***
							(11.22)
Obs	159,344	158,228	205,786	202,978	466,667	454,251	419,519
Adj. R^2	0.790	0.800	0.402	0.403	0.074	0.117	0.123
Mean dep var	.113	.112	.469	.47	.769	.752	.689
First obs	2013m6	2013m6	1990m1	1990m1	2020m9	2020m9	2020m9
Controls		✓		✓		✓	✓

Regressions of interest rate forecast errors on real income groups across surveys. The construction of each forecast error (and their differences) is detailed above. As they each differ in their interpretation, the quantitative estimates are not comparable across surveys. Odd columns only include time-fixed effects, even columns include our vector of controls also. The final column uses the same specification as column (6), but also includes the expected growth in local house prices 12 months in the future. For the MSC, we use the period 1990 onwards. All t-statistics are calculated using standard errors clustered at the household level. Survey weights are employed in all regressions. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

expects mortgage rates to 0.75% lower than materialises in 12 months.

As the ECB-CES survey asks specifically about local mortgage rates, rather than savings or borrowing rates as the SCE and MSC respectively do, it may be that differences in interest rate expectations are an artefact of differential views on the housing market. In the final column of Table A.13, we also control for the respondents expectations of house price growth in 12 months. We find that those who expect higher house prices typically have larger forecast errors, but the results do not change the qualitative cross-sectional relationship between income and interest rate forecasts.

A.6 Other Variables

Here we detail the construction of the other expectation forecast errors. Most are taken from the SCE. They are:

Stocks: Respondents in the SCE are asked *"the percent change that 12 months from now, on average,* *stock prices in the US stock market will be higher than they are now"*. We construct a forecast error measure similar to that of SCE Unemployment or interest rates and is bounded in [-1, 1]. For stock prices, we assume an increase is perceived to be beneficial to the respondent and so invert the measure.

Government Debt Growth: Respondents answer *"By about what percent do you expect the level of U.S. government debt to [decrease/increase]?"*. For government debt, we assume that increases are perceived to be negative, although this comes with added caveats[70].

Other Prices: Consumers are asked their expectation of percentage price changes in 12 months for a gallon of gas, food, medical care and *"the cost of renting a typical house/apartment"*. They are also asked what they expect to happen to *"the average home price nationwide"*[71]. For gas, food and medical care prices, we assume increases are perceived to be detrimental to the respondent. For rent prices and house prices whether or not increases benefit one individual or another depends on their home-ownership status. We therefore restrict our estimates across these two variables to renters and homeowners only. We assume that homeowners perceive aggregate increases in house prices to be net positive for their individual welfare. Given this, we invert the measure for house prices too.

Gas price and house price growth are also available in the MSC. We treat these in the same manner as we do in the SCE.

Table A.14 displays the summary statistics for the (not normalised) forecast errors shown in Figure 1.8. These are the summary statistics for all forecast errors discussed in section 1.4.

In the main body of the article in Figure 1.8 we display the forecast errors across all the variables discussed, but normalised by the population forecast error. All normalised forecast errors therefore have a mean of unity. This allows for a graphical comparison across forecast errors. In Figures A.5, A.6 and A.7 we plot these normalised forecasts by income group across the SCE, MSC, and ECB-CES respectively. For the sake of legibility, confidence intervals have been omitted.

Across the SCE, almost all variables show a declining profile in income, albeit to differing degrees. The only exception to this is both house price forecast errors, which increase in income. For this variable

[70] Specifically requiring that the respondent doesn't perceive an increase in government debt as positive for welfare. In reality, an increase in government borrowing may be associated with higher welfare for at least some households. For example, it could improve the household budget constraint directly (either through an increase in transfers or a decrease in taxes) or indirectly (through an increase in income or asset returns in general equilibrium, say). For individuals to be pessimistic about government debt increases, these effects must either be ignored or perceived as net negative for the respondent

[71] And are explicitly prompted to "think about home prices nationwide" before being asked the question, too.

Table A.14: Summary Statistics for Forecast Error Measures

	Mean	Median	St. Dev	N
SCE				
Inflation (Point)	3.08	1.77	6.20	145048
Inflation (Density)	1.57	0.90	6.10	156473
Unemployment	0.16	0.25	0.47	159337
Interest Rate (Sav. Acc.)	0.13	-0.02	0.55	159344
House Prices (Point, owners)	1.80	1.13	8.14	110065
House Prices (Density, owners)	2.80	2.24	8.12	101580
Rental Prices (nonowners)	4.60	2.33	6.84	32481
Gov. Debt	4.44	2.26	11.82	127354
Gas Prices	4.67	7.58	23.74	129191
Food Prices	3.75	2.76	5.94	128957
Medical Prices	7.88	6.09	8.27	127354
Stock Prices	0.42	0.50	0.45	158730
MSC				
Inflation	0.26	0.06	4.19	256908
Unemployment	0.16	0.00	0.88	296796
Interest Rates (Borrowing)	0.38	0.00	0.93	293925
House Prices (owners)	2.47	3.24	7.57	75989
Gas Prices	12.37	8.40	33.55	114190
ECB-CES				
Inflation (Point, CPI)	0.40	-0.72	6.20	693625
Inflation (Density, CPI)	2.49	0.55	7.97	677424
Unemployment	5.95	2.70	11.49	760100
Interest Rates (Mortgages)	0.79	0.32	2.36	466667
Real Consumption (Point)	1.12	-1.02	14.49	286167
Real Consumption (Density)	2.48	0.20	14.77	268741
Nominal Expenditure	1.57	0.00	6.41	186121

This table displays summary statistics for all of the forecast error variables calculated throughout section 1.4. The summary statistics are constructed using survey weights.

only is there evidence to suggest that richer households form more pessimistic expectations (i.e. expect lower house-price growth than is realised). It is possible that even home owners do not perceive house price increases as welfare positive, which may help to explain the cross-sectional differences (although not the sign of the forecast error).

In the MSC, again all variables show a declining profile in income, with the exception of gas inflation forecast error, which is hump-shaped. Inflation forecast errors display the most pronounced dependence on income, followed by unemployment expectations.

Finally, the ECB-CES shows a declining profile across all variables considered - here, unemployment forecast errors display the strongest correlation with income group. Here we employ the full sample of the ECB-CES dataset. As discussed in 1.4.1, it is unsurprising that the point inflation forecast error is negative given the prominence of the 2021-2022 inflationary shock in the sample.

Figure A.5: Normalised Forecast Errors by Income (SCE)

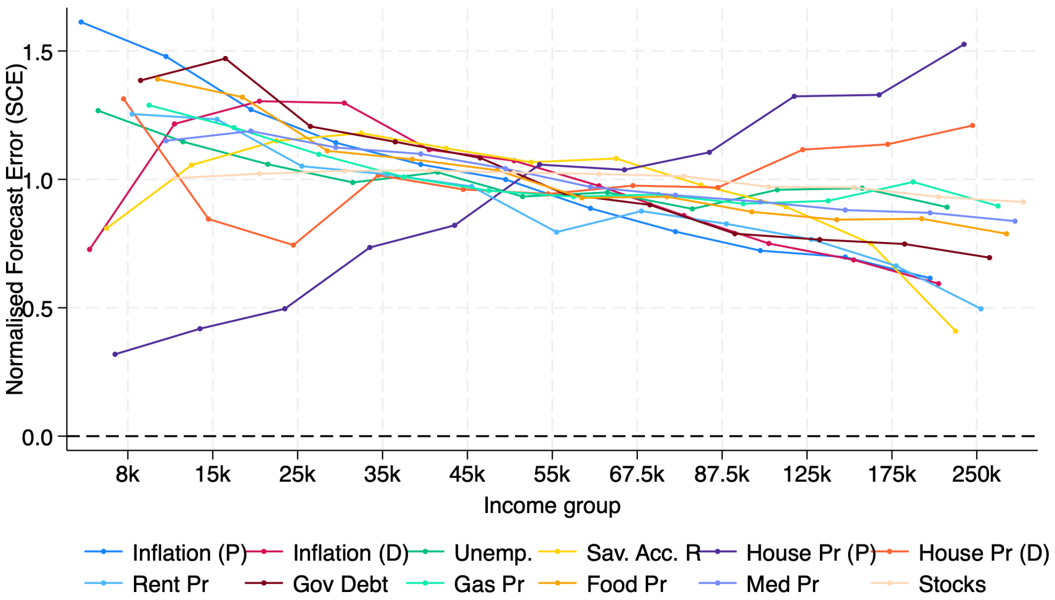

Here we plot the marginal effect of income on the normalised forecast errors, for the full range of expectations considered in the SCE. The regressions include time fixed effects, and the full set of demographic controls used in the SCE specifications. The two house price regressions are run over the sub-sample of homeowners only. Similarly, the rental price regression is run over the subsample of non-owners. (P) and (D) denote the point and density-implied forecast errors respectively.

Figure A.6: Normalised Forecast Errors by Income (MSC)

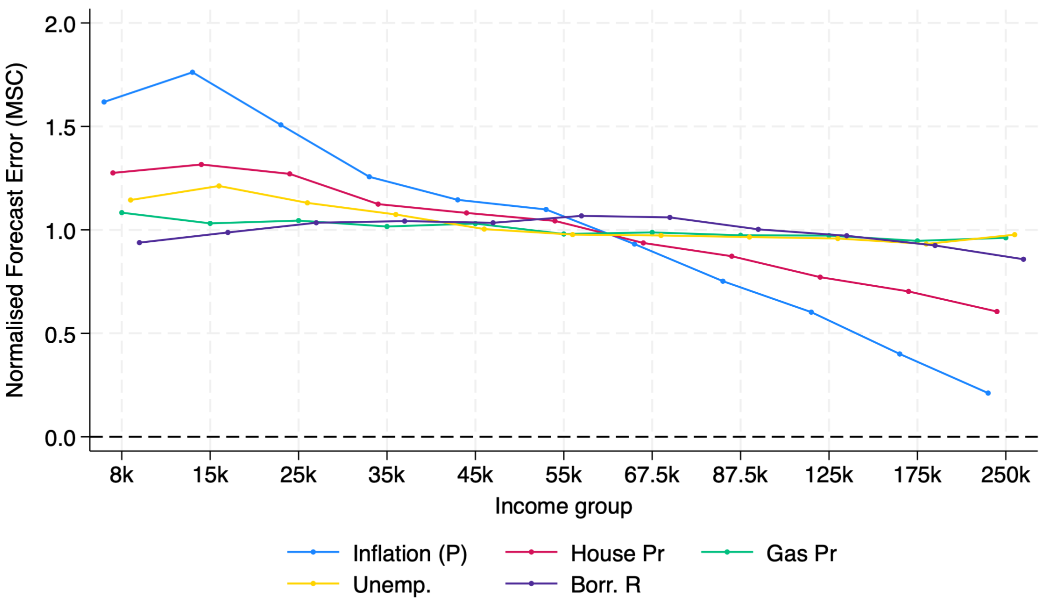

Here we plot the marginal effect of income on the normalised forecast errors, for the full range of expectations considered in the MSC. The regressions include time fixed effects, and the full set of demographic controls used in the MSC specifications. The house price regression is run over the sub-sample of homeowners only.

Figure A.7: Normalised Forecast Errors by Income (ECB-CES)

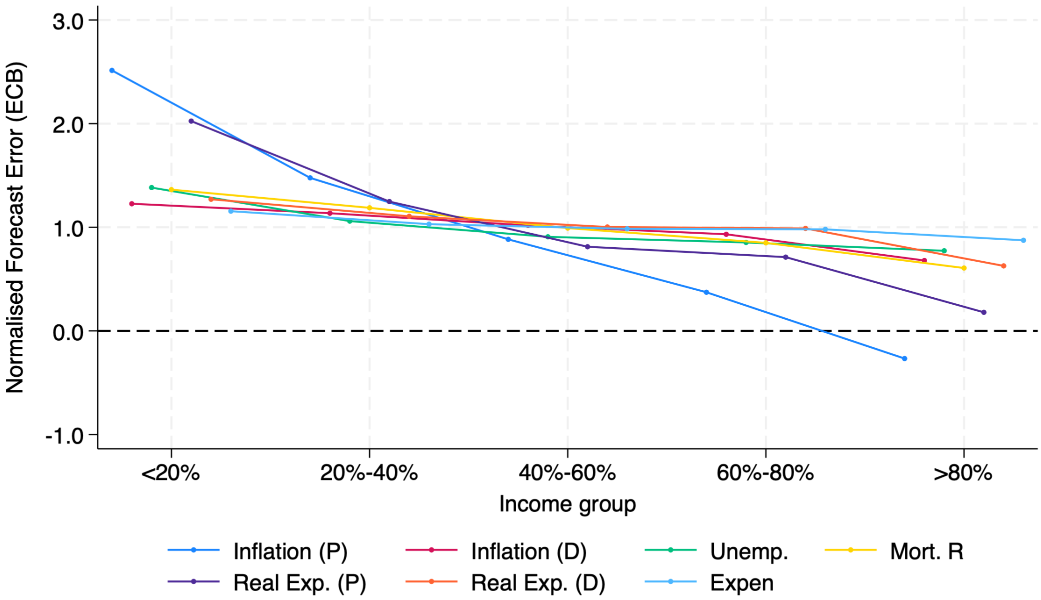

Here we plot the marginal effect of income on the normalised forecast errors, for the full range of expectations considered in the ECB-CES. The regressions include time and country fixed effects, and the full set of demographic controls used in the ECB-CES specifications. (P) and (D) denote the point and density-implied forecast errors respectively. 'Expen' denotes the nominal expenditure forecast error, discussed in the Section 1.4.2.

A.7 Difference in Difference Appendix

In this section, we provide additional details on the inflation forecast error difference-in-difference specification of Section 1.4. We use three alternative definitions of income shocks. These are:

Baseline Measure: For all respondents who respond twice to the MSC, we calculate their change in reported income. We convert this into 2015 dollars using the CPI index, and then annualise the percentage difference based on the time between surveys (almost always 6 months). Separately, we calculate their *expected* real income change using their income and inflation expectations one year ahead. Both of these measures are then trimmed at the 5-95% level. The difference between these two is the "raw" shock size $\text{shock}_{i,t}$. If this is greater than 2 or less than 0.5, we remove the observations. Using a threshold of 10% unexpected changes, we define our shocks as: $\text{shock}^{+}_{i,t} = \mathbb{I}\left(\text{shock}_{i,t} > 0.1\right)$ and $\text{shock}^{-}_{i,t} = \mathbb{I}\left(\text{shock}_{i,t} < -0.1\right)$ respectively.

Nominal Measure: We follow a similar procedure as the baseline measure but do not deflate the actual change or the expected change in income by CPI. We therefore take their annualised percentage change in nominal income and calculate the difference with their expected nominal income change using the one-year ahead income expectations. After trimming at the 5-95% level and discarding values greater than 2 and less than 0.5, we define the "raw" nominal shock as $\text{nominal shock}_{i,t}$ and define $\text{nominal shock}^{+}_{i,t} = \mathbb{I}\left(\text{nominal shock}_{i,t} > 0.1\right)$ and $\text{nominal shock}^{-}_{i,t} = \mathbb{I}\left(\text{nominal shock}_{i,t} < -0.1\right)$ accordingly.

Anticipated Measure: Our final measure eliminates the expectation component of $\text{shock}_{i,t}$. We calculate the annualized percentage change in real income as above, trimming at the 5-95% level and removing values outside the range [0.5, 2]. We then use this as our "raw" shock $\text{anticipated shock}_{i,t}$, and define: $\text{anticipated shock}^{+}_{i,t} = \mathbb{I}\left(\text{anticipated shock}_{i,t} > 0.1\right)$ and $\text{anticipated shock}^{-}_{i,t} = \mathbb{I}\left(\text{anticipated shock}_{i,t} < -0.1\right)$.

For both the baseline version and the nominal measure, we include a "continuous" version of the shock variable. That is, we also let the intensity of the treatment (i.e. the size of the shock) vary between individuals. In essence, we are interacting the magnitude of the shock with the treatment/no-treatment dummy variables. These are defined as:

$$\text{c-shock}^{+}_{i,t} = \begin{cases} z_{i,t} \text{ if } |z_{i,t}| > 0.1 \\ 0 \text{ otherwise.} \end{cases} \qquad \text{for } z_{i,t} \in \{\text{shock}_{i,t}\ ,\ \text{nominal shock}_{i,t}\}$$

These measures are also included in Table 1.7.

Table A.15: Income Shock Proxy Summary Statistics

	Mean	Median	St. Dev	N
Baseline				
-ve Shock	-0.26	-0.24	0.12	25822
No Shock	0.01	0.00	0.05	42280
+ve Shock	0.50	0.38	0.38	51248
Nominal				
-ve Shock	-0.28	-0.25	0.13	33116
No Shock	-0.01	0.00	0.05	51041
+ve Shock	0.53	0.40	0.40	54110
Anticipated				
-ve Shock	-0.26	-0.25	0.11	23399
No Shock	0.00	0.00	0.03	45879
+ve Shock	0.50	0.38	0.38	50072

This table shows the summary statistics for the "raw" income shock poxy employed in the MSC data. Here we have subsetted each "raw" shock over the intervals $\text{shock}_{i,t} < -0.1$, $-0.1 < \text{shock}_{i,t} < 0.1$ and $\text{shock}_{i,t} > 0.1$. For example, row one shows the summary statistics for the baseline raw shock, conditional on the raw shock being below the cut-off -0.1. Note that these are not the variables used in the regression (which, for the baseline specification, is the standard Diff-in-Diff treated-untreated indicator variable. Survey weights are employed.

In Table A.15 we display the summary statistics for each of these three definitions of the "raw" shocks. Given the large standard error of the conditional means of the "raw" shock series, we view the discrete versions as preferred.

Across Income Groups

We may be interested in whether or not the effect of an income shock on inflation expectation forecasting errors also vary by initial income. We investigate this by employing an alternative specification, where we interact the treatment and shock variables in the above specification with the household's initial real income tercile (within the calendar year). This allows us to investigate to what extent income shocks have heterogeneous effects on inflation forecasting errors. To this end we also specify the following regression:

$$\begin{aligned} \mathbb{E}_{i,t-12}(\pi_t) - \pi_t = & \sum_{g \in \{L,M,H\}} \left(\beta^+ + \beta_g^+ \cdot \mathbb{I}(i \in g)\right) \cdot \text{treat}_{i,t} \times \text{shock}_{i,t}^+ + \\ & + \sum_{g \in \{L,M,H\}} \left(\beta^- + \beta_g^- \cdot \mathbb{I}(i \in g)\right) \cdot \text{treat}_{i,t} \times \text{shock}_{i,t}^- + \lambda_i + \delta_t + \varepsilon_{i,t} \end{aligned}$$

Here, the ATT for a high-income group in receipt of a positive income shock is $\beta^+ + \beta_H^+$, if $i = L$ is chosen as our base group. Given that forecast errors are generally larger for lower income households, and marginal utility of consumption (and therefore income) is decreasing, we may expect the magnitude of the response of forecast errors to an income shock to be largest for the poorest households and declining in income. We plot the ATTs obtained from the above regression in Figure A.8.

Figure A.8: Marginal Effect of Income on Inflation Forecast Error's Response to Income Shock Proxy

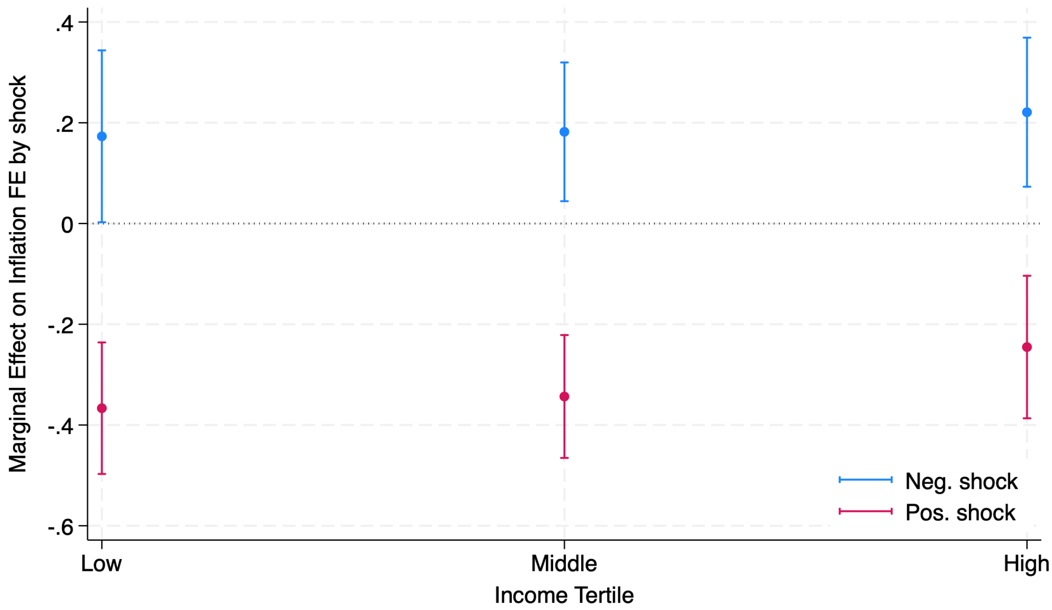

Plotted is the marginal effect (ATT) of an unanticipated income shock on inflation forecast error, across income tercile (in the initial period). Standard errors are clustered at the individual level, and the confidence intervals denote the 95% range.

We see some, albeit limited, evidence for heterogeneous effects in the specification. At the lower end of the income distribution, the effects of a negative income shock on inflation forecast errors are not statistically distinguishable from zero, but this effect increases and becomes significant as income increases. For positive income shocks, the expected shape of the point estimates holds, although the estimates are not precise enough to detect a declining magnitude as income changes. Naturally - given the fairly crude nature of our "shock", and limited data - our estimates being imprecise is unsurprising.

Chapter 2

Ambiguity Aversion, Asymmetric Signal Extraction, and Household Expectations

2.1 Introduction

Why do surveys of household inflation expectations consistently exhibit an upward bias? This paper examines a particular explanation, ambiguity aversion, and inquires whether this mechanism can account for the facts presented in Chapter 1. The seminal work of Gilboa and Schmeidler (1989) demonstrates that for households facing Knightian uncertainty - where households are unsure of the data-generating distribution itself - ambiguity-averse agents will behave under the worst-case set of beliefs. We explore the implications of such behaviour in the context of inflation expectation formation.

We use a simple model to define what ambiguity aversion is in the context of expectation formation. The model provides a guide on how to test the hypothesis that ambiguity aversion drives the persistent, positive forecast errors observed in the data. We test these two hypotheses against the data, linking individual-level surveys that contain information on both experimentally elicited ambiguity preferences and survey data. The data support these two hypotheses: individuals who are more ambiguity-averse display larger forecast errors, and expectations are convex in response to a signal shock.

Finally, the results of Chapter 1 detail that forecast errors are typically *more* positive for low-income households relative to high-income households. We test that households sort by their preferences for ambiguity in the income distribution. We find no concrete evidence to support this claim: it appears that low-income households behave as if they are more ambiguity-averse or possess different information sets regarding the nature of the data-generating process relative to higher-income households, rather than having different ambiguity preferences themselves.

Consumer inflation expectations are typically larger than realised inflation (Weber et al. (2022), Fofana et al. (2024)). One possible explanation - discussed somewhat in the literature already[1] is ambiguity aversion. Section 2.2 describes ambiguity in the context of expectations. Pessimistic forecast errors are consistent with several sources of ambiguity, including ambiguity over the true long-run mean, the persistence of the process, or the information content of signals regarding the data-generating process. Considering this, we derive a separating test: only ambiguity across the information content of future signals can generate both positive forecast errors and a convex expectation elasticity to signals. That is, we derive conditions under which the data is consistent with consumers forming subjective beliefs based on signals of the future, of which the information content is ambiguous[2].

With these hypotheses in hand, we perform two exercises. First, we link experimental data eliciting ambiguity preferences with economic forecast surveys. We combine experimental evidence on attitudes towards ambiguity with macroeconomic expectations data using RAND's American Life Panel (ALP). We link households across multiple surveys, enabling us to measure both survey-elicited preferences regarding ambiguity and inflation expectations. We demonstrate that among ambiguity-averse individuals, forecast errors increase with the degree of ambiguity aversion. An ambiguity-averse individual at the 10th percentile in our data, compared to one at the 90th percentile, exhibits an inflation forecast error that is 0.52 percentage points larger.

Secondly, we provide evidence to suggest that the response of consumers' expectations to public signals is asymmetric. Conditional on the ambiguity environment of Section 2.2, this feature is only consistent with ambiguity over the information content of signals[3]. Household expectations react strongly to negative public signals (as proxied by the Survey of Professional Forecasts) but respond in a much weaker manner to positive signals. This supports the hypothesis that at least one source of ambiguity for consumers is the information content of the signals they receive, and is not explained by Knightian

[1]See Ilut and Schneider (2014), Baqaee (2020), Ilut and Schneider (2023) and Bhandari et al. (2024) for recent treatments.

[2]We do not claim that the framework of ambiguity is the only possible explanation for positive forecast errors. For example, an alternative possibility is the work of Malmendier and Nagel (2016). Whether these possibilities also explain the empirical patterns in expectations data is beyond the scope of this paper. We leave it to other authors. Our results are simply conditional on the reader accepting the possibility of Knightian uncertainty.

[3]This is supporting evidence to Baqaee (2020).

uncertainty over the other parameters of the data-generating process alone. The intuition behind an agent's response to a signal is simple enough: upon receiving a good signal about the future state of the world, consumers treat the signal as if it has low information content. On the other hand, when faced with a bad signal of the future consumer, treat the signal as more informative.

The results of Chapter 1 highlight that forecast errors decline with income. In the context of ambiguity aversion, this implies that more affluent households behave as if they are less ambiguity-averse or possibly less subject to ambiguity. We find evidence to support this. In our test of expectation asymmetry, we find that poorer households appear to show a greater degree of asymmetry relative to their more affluent counterparts. Again, this is consistent with the notion that poorer households act *as if* they are more ambiguity averse.

Finally, we further explore the interaction between ambiguity aversion and income. Using the ALP-HIEP data set, we test the hypothesis that richer households indeed have different preferences: that the cross-sectional relationship between income and expectations is an artefact of sorting due to latent preference heterogeneity. However, despite considerable variation in preferences towards ambiguity in the data, we find that preferences are uncorrelated with measures of income or wealth. This lack of correlation suggests that differing preferences themselves do not drive the relationship between expectation formation and income, but nonetheless leads households to behave *as if* they are less ambiguity averse when richer: either epiphenomenally, or due to experiencing less Knightian uncertainty than poorer households. This finding supports the view that ambiguity is, at least in part, salient for the expectation formation process for consumers.

Literature Our paper is motivated by the now extensive body of microeconomic evidence on household expectations. Coibion, Georgarakos, et al. (2023), D'Acunto, Charalambakis, et al. (2024), and Chapter 1 of this thesis detail that households typically have positive forecast errors. Weber et al. (2022) and Fofana et al. (2024) both provide useful overviews of the demographic characteristics of inflation expectations. Worth highlighting is Masolo and Monti (2024), who finds that savers have higher expectations of inflation than borrowers. In this paper, we employ simple Gilboa and Schmeidler (1989) preferences, in which the utility of households across possible states of the world matters only insofar as the worst-case state. Our set-up of Section 2.2 would be consistent with their empirical results if the primary cost of inflation is the Fisher channel (Fisher (1933))[4]. Other work, such as that of Coibion, Gorodnichenko, and Kamdar (2018), details positive forecasting errors for firms too (albeit closer to both the expert forecasts

[4] In smooth ambiguity models, such as Klibanoff et al. (2005), Maccheroni et al. (2006) and Klibanoff et al. (2009) - of which Section 2.2 and Gilboa and Schmeidler (1989) are limiting cases - it is easier to generate such behaviour.

of the Survey of Professional Forecasters and the rational forecast of zero). Kamdar and Ray (2024), and again, Chapter 1 of this thesis, provide evidence that household forecast pessimism is not just limited to expectations of inflation, but also exists across a range of other economic outcomes.

Baqaee (2020) performs an exercise using the same framework as ours and shows that the response of inflation forecast errors is asymmetric on receipt of a signal. We provide additional detail to these results, and illustrate three extra results: first, that the asymmetry in expectations also exists across unemployment and interest rate expectations; second, that the *degree* of asymmetry (convexity) is decreasing in income; third, that more ambiguity averse individuals also display the same feature of their expectations. More generally, we contribute to the literature of microeconomic evidence of household expectations by investigating a discussed potential cause of such forecasting pessimism - ambiguity aversion - and outlining the types of ambiguity that are consistent with the data. Moreover, we also employ experimental survey evidence that allows us to link ambiguity preferences and forecasting errors for the first time.

In light of the empirical evidence, a separate body of literature attempts to explain the potential causes of these patterns in the data. One strand of this literature is the sticky-information Mankiw and Reis (2002) class of models, which posits that agents may intermittently update their information sets. Under certain conditions (such as downward-trending inflation), this framework can generate persistent forecast errors. An alternative explanation is that households engage in adaptive learning, one aspect of bounded rationality Sargent (1993). In the work of Malmendier and Nagel (2016), agents learn about the inflation process but may have biased priors, formed from their inflation experience. Here, agents use a subjective weighted average of past experiences to form their expectations, meaning those born during periods of high inflation typically have persistent positive forecast errors, which serves to generate a cohort effect[5].

This paper instead focuses on an alternative explanation: ambiguity aversion. This explanation has gained traction in macroeconomics due to its simplicity in directly accounting for pessimistic or worst-case thinking. Ellsberg (1961) first illustrated ambiguity with the Ellsberg paradox: people typically prefer lotteries with known odds, rather than unknown odds; they act as if the odds of the unknown lottery are worse. Dimmock et al. (2013)) asks respondents to perform an Ellsberg urn experiment

[5]Other mechanisms can generate deviations from Full Information Rational Expectations (FIRE). Rational inattention (see Maćkowiak et al. (2021) for a review) suggests that agents update their information set optimally. However, this theory can only generate positive forecast errors in the short run, as expectations are sluggish to update. An alternative line of work considers behavioural biases, such as diagnostic expectations (Bordalo, Gennaioli, and Shleifer (2022), Bianchi et al. (2024),Gabaix (2019) Here, agents overweight recent or representative information. Again, this can generate persistent forecast errors, but not over long periods. Both mechanisms also generate overreaction to signals. However, at least in the most standard presentations of these theories, this overreaction is symmetric.

and finds that ambiguity preferences vary across people but appear unrelated to most demographics, suggesting ambiguity aversion is a deep preference parameter, and not just a proxy for knowledge or economic circumstances. In Section 2.4, we link the same survey data to data on inflation forecast errors. Abramson et al. (2024) also utilises the same dataset but instead links respondents to a survey questioning mental health status. They show those who display more "pessimistic thinking", in their words - ambiguity aversion, in the original surveyors - are also typically suffering a greater degree of depression and anxiety. To our knowledge, our paper is the first work linking (at the individual level) experimental evidence on ambiguity preferences to expectations data.

Gilboa and Schmeidler (1989) formalises ambiguity preferences and shows that agents facing a set of possible probability distributions evaluate lotteries by using the worst-case distribution. We lean on this presentation in Section 2.2 to outline the hypotheses we test later in the paper. Other papers situate ambiguity in the context of general equilibrium, such as Ilut and Schneider (2014), which integrates ambiguity aversion into a New Keynesian model. Baqaee (2020) employs ambiguity in response to signal strength to motivate why wages exhibit downward rigidity. Huo et al. (2023) characterise the effects of ambiguity under dispersed information in equilibrium, and Masolo and Monti (2021) considers the impact of ambiguity over the Central Bank's inflation target[6]. We contribute to the literature on expectation formation by providing and testing hypotheses which refine the *types* of ambiguity aversion consistent with household expectation data. We therefore do not rule out other proposed mechanisms in the literature, but rather seek to refine exactly in what way one mechanism may or may not explain the data. Our conclusions, therefore, rest on the reader accepting the ambiguity-aversion narrative as at least possible.

Outline This paper is structured as follows: Section 2.2 outlines ambiguity and provides testable hypotheses for the possible sources. Section 2.3 describes the data sources used throughout the paper. Section 2.4 links the American Life Panel experimental evidence on ambiguity aversion preferences to inflation forecast errors. Section 2.5 tests whether ambiguity preferences and income are related. Section 2.5 explores the elasticity of expectations to public signals, using the Survey of Professional Forecasters. Section 2.7 concludes.

[6]See also Ilut and Schneider (2023) for a review of ambiguity-aversion in macroeconomics. A closely related modelling framework is that of concern for model misspecification, à la (Hansen and Sargent (2001)). Bhandari et al. (2024) uses this to evaluate the effect of time-varying pessimism in a representative agent general equilibrium model. Chapter 3 employs the robustness framework in the context of a heterogeneous agent (partial equilibrium) portfolio-choice model, calibrated to match the aggregate inflation expectations of US households.

2.2 Ambiguity and Forecasting Errors

An ambiguous environment is one in which agents do not know the true parameters of the data-generating distribution. Instead, they may only know that the true parameter lies in a particular range[7].

Why does such an environment generate pessimism? We focus on the example of agents forecasting inflation for the sake of exposition. However the following holds for any variable in which welfare is monotonic. Assume that the household in question is a saver, such that increases in inflation are associated with a lower future utility[8]. Further, assume inflation follows a first-order auto-regressive process:

$$\pi_{t+1} = (1-\rho)\mu + \rho\pi_t + \varepsilon_{t+1} \qquad \varepsilon \sim N(0,\ \sigma^2)$$

Here, π_t is inflation in period t, μ is the long-run mean of inflation and ρ governs the persistence of the process. It is assumed that households can observe current and past inflation rates. In order to form expectations about future inflation, households must form expectations over the future shock ε_{t+1}. Imagine households receive a signal in period t of shock ε_{t+1}. The signal they receive in time t is:

$$s_t^i = \varepsilon_{t+1} + \epsilon_{t+1}^i \qquad \epsilon_{t+1}^i \sim N(0,\ \sigma_s^2)$$

Here ϵ_{t+1}^i is i.i.d. signal noise, and is independent of both the shock to inflation ε_{t+1}. For notational convenience from here on we suppress dependence of s_{t+1}^i on i unless necessary. Based on this signal, the household updates their beliefs according to the Bayes rule. The rational forecast is therefore given by:

$$\mathbb{E}_t[\pi_{t+1}|\ s_t] = (1-\rho)\mu + \rho\pi_t + K(\phi)\cdot s_t$$

where $K(\phi) = \frac{\phi}{1+\phi}$ is the typical Kálmán gain, and $\phi \equiv \frac{\sigma^2}{\sigma_s^2}$ is the signal to noise ratio. We rewrite this term in terms of ϕ as the relative size of σ and σ_s is all that matters for what follows.

Gilboa and Schmeidler (1989) provide a representation theorem for the preferences of ambiguity averse agents when the true parameters governing the data-generating process are not known. In such a case, agents follow a Min-Max procedure, where they make decisions to maximise their utility in the worst-

[7]One easy way of thinking of this is considering an ambiguous environment to be one in which there is second order certainty.

[8]This can be easily be explicitly added to the model of this section, but adds little to the exposition.

case state of the world. As consumers are savers and so are hurt by inflation, the worst case state of the world chooses the parameter that maximises their inflation expectation[9].

2.2.1 What is the Source of Ambiguity?

We now explore the implications of ambiguity aversion across the three parameters (μ, ρ, ϕ) of our simple environment.

Ambiguity over μ. First, consider the case in which consumers are unsure of the true mean rate of inflation, and instead only know that the true mean lies in the interval $[\underline{\mu}, \bar{\mu}]$. The true parameters ρ and ϕ are known. Therefore the worst-case state of the world is the one in which the long-term mean of inflation is high ($\bar{\mu}$). The expectation of the ambiguity averse agent at time t, conditional on receipt of a signal s_t over the future state of the world, is given by:

$$\mathbb{E}_t^A [\pi_{t+1}| \, s_t] = (1-\rho)\bar{\mu} + \rho\pi_t + K(\phi) \cdot s_t$$

This is of course greater than the rational expectation without ambiguity. Define the forecast error on receipt of a signal of the ambiguity-averse agent as:

$$FE_{t+1}(s_t|\pi_{t+1}) \equiv \mathbb{E}_t^A [\pi_{t+1} \mid s_t] - \pi_{t+1}$$

This is the forecast error of the ambiguity averse agent in any given period. Over a long course of signals, or in the cross section, we can consider average forecast error. Define this as:

$$\overline{FE_{t+1}}(\pi_{t+1}) \equiv \int_s FE_{t+1}(s_t|\pi_{t+1}) \, dF(s_t)$$

Here, $F(\cdot)$ is the cumulative density function of a normal distribution with mean zero and variance $\sigma^2 + \sigma_s^2$. This is the empirical counterpart of the forecast errors we measure in Chapter 1: it is the average forecast error measured over a long history of signals, or across a population who receive uncorrelated signals (i.e. it is unconditional on s_t). When there is ambiguity towards the true long-run mean of inflation, μ, we have: $\overline{FE_{t+1}} = (1-\rho)(\bar{\mu}-\mu)$. This is strictly positive so long as $\rho \in [0,1]$ and $\bar{\mu} > \mu$ (i.e. there is at least some ambiguity).

[9]Here, the dependence of utility on the stochastic variable is trivial, but nonetheless, the expectations of the consumer are a function of future utility. When there are more stochastic variables that the consumer must form expectations over, these expectations will all co-vary with future utility across states of the world. Expectations over many variables will therefore all be predicted by the consumer's sentiment towards their own future welfare across states, whereas under FIRE this would of course not be the case.

Ambiguity over ρ. Of course, ambiguity over the mean rate of inflation is a natural choice of parameter to generate such a result. It is not the only possibility. When there is ambiguity over ρ - the agent only knows that the true ρ lies in the interval $[\underline{\rho}, \bar{\rho}]$ - we demonstrate in Appendix B.1 one can write the average forecast error as:

$$\overline{\text{FE}_{t+1}} = \sigma \cdot \frac{(\bar{\rho} - \underline{\rho})}{(1-\rho^2)^{\frac{1}{2}}} \cdot \left(\frac{2}{\pi}\right)^{\frac{1}{2}} \quad > \ 0 \quad \text{when} \quad \bar{\rho} > \underline{\rho}$$

This follows from the properties of the half-normal distribution. As $\bar{\rho} - \underline{\rho} \rightarrow 0$, the mean forecast error $\overline{\text{FE}_{t+1}} \rightarrow 0$: as the agent becomes sure of the data-generating process, the mean observed forecast error declines to the rational forecast.

Ambiguity over ϕ. A second possibility is that consumers are ambiguous over the information content of a given signal (ϕ). This may be generated by ambiguity regarding the noise component of any signals they receive (σ_s), the variance of the underlying process itself (σ), or potentially both. For the sake of our results, ambiguity over either generates the same qualitative predictions and in either case reduces to the agent being uncertain over the value of $K(\phi)$. Let the agent know that the true value of ϕ lies in the range $[\underline{\phi}, \bar{\phi}]$. In such an instance, the expected value of future inflation is given by:

$$\mathbb{E}_t^A\left[\pi_{t+1} | \ s_t\right] = \begin{cases} (1-\rho)\mu + \rho\pi_t + K(\underline{\phi}) \cdot s_t & \text{if} \quad s_t \ < \ 0 \\ (1-\rho)\mu + \rho\pi_t + K(\bar{\phi}) \cdot s_t & \text{if} \quad s_t \ \geq \ 0 \end{cases}$$

Importantly, expectations of future inflation depend on the agent's future utility across states of the world. If the signal is positive, they take this seriously; the worst-case scenario is one in which the signal of bad news is informative, and so they form their expectations as if this were the case. If the signal suggests that inflation will be low (i.e. good news), the worst-case scenario is that the signal is noisy, meaning that little weight should be placed on it. This is as in Baqaee (2020). Here, the unconditional forecasting error as measured of a long history of signals is given by:

$$\overline{FE_{t+1}} = \frac{1}{\sqrt{2\pi}} \cdot \frac{\bar{\phi} - \underline{\phi}}{(1+\bar{\phi})(1+\underline{\phi})} \cdot \left(\frac{\sigma^2}{K(\phi)}\right)^{\frac{1}{2}}$$

It is clear to see that as $\bar{\phi} - \underline{\phi} \rightarrow 0$ we have $\overline{FE_{t+1}} \rightarrow 0$. The first term is the mean of a half-normal distribution. The second term represents the breadth of ambiguity: in worst-case Kálmán gains, it is $K(\bar{\phi}) - K(\underline{\phi})$. The final term represents the scale of the uncertainty, and could be rewritten as the standard deviation of the signal: $(\sigma^2 + \sigma_s^2)^{\frac{1}{2}}$. This scales the mean forecast error: the noisier the signals are, the larger the mean forecast error is.

Distinguishing the source of ambiguity. All three of these possibilities would be consistent with positive empirical forecast errors, if we observe consumers making decisions over a long history of signals, or in the cross-section where signals across respondents are uncorrelated. To distinguish the source of ambiguity, we consider the response of the individual forecast error to a realisation of a given signal. Doing so gives rise to Proposition 2.2.1:

Proposition 2.2.1. *For an ambiguity-averse individual i forecasting a variable x_t under Gilboa and Schmeidler (1989) preferences, assume that the agent perceives increases in x_t to be detrimental to welfare. Let x_t follow an AR(1) process and the agent receive signals s_t at time t over future innovations in x_t. Define the forecast error on receipt of a signal as $FE^i_{t+1} = \mathbb{E}^i_t[x_{t+1}|s^i_t] - x_{t+1}$, and the average forecast error across realisations of the signal as $\overline{FE_{t+1}} = \mathbb{E}_i\left[FE^i_{t+1}\right]$. Let $\mathcal{P}$ denote the ambiguity set over each parameter. For $\kappa \in \{\mu, \rho, \phi\}$, assume the true value of the parameter adheres to $\inf \mathcal{P} < \kappa < \sup \mathcal{P}$.*

- *Ambiguity over the long-run mean (μ), persistence (ρ) and signal strength (ϕ) all generate positive mean forecast errors. That is: $\overline{FE_{t+1}} > 0$.*

- *Only ambiguity over ϕ generates an asymmetric response of $\mathbb{E}^i_t[x_{t+1}|s^i_t]$ to a given signal. Formally, the response $\frac{d\mathbb{E}^i_t[x_{t+1}|s^i_t]}{ds^i_t}$ is asymmetric if $\frac{d\left|\mathbb{E}^i_t[x_{t+1}|k]\right|}{ds^i_t} \neq \frac{d\left|\mathbb{E}^i_t[x_{t+1}|-k]\right|}{ds^i_t}$ for $k \in \mathbb{R}$. For ambiguity over the long-run mean μ and persistence ρ, the response of an individual forecast is symmetric. For ambiguity over signal strength ϕ, the sensitivity of the forecast to a given signal is asymmetric. It is given by:*

$$\frac{d\mathbb{E}^i_t[x_{t+1}|s^i_t]}{ds^i_t} = \begin{cases} K(\underline{\phi}) & \text{if} \quad s^i_t < 0 \\ K(\bar{\phi}) & \text{if} \quad s^i_t \geq 0 \end{cases} \qquad \textit{where:} \qquad K(\underline{\phi}) < K(\bar{\phi})$$

Proof: See Appendix B.1.

Ambiguous Forecasting and Income. How can we link the theory presented so far to the notion that these forecast errors decline with income, i.e. the results of Chapter 1? Assume that the *degree* of ambiguity varies by income and let the set of parameter values over which the individual is ambiguous towards vary by income. For low-income households this set is large, incorporating a wide range of possible values; for richer households this set is smaller. For any individual, maintain that the true value of the parameter remains within the ambiguity set in all cases[10]. Then, when there is ambiguity over ϕ, we can generate both: 1) forecast errors that are pessimistic but decline with income; and 2) an asymmetric responses of forecasts to signals, but where the asymmetry declines with income. This is stated more formally in Proposition 2.2.2:

Proposition 2.2.2. *For a Gilboa and Schmeidler (1989) agent, the degree of ambiguity the individual faces is captured by the diameter of the set of priors. Let $\mathcal{P}$ denote the set of possible values of $\kappa \in \{\mu, \rho, \phi\}$. Define the dimension of this set as $k := \text{diam}\, \mathcal{P} = \sup \mathcal{P} - \inf \mathcal{P}$. Assume that the true value of this parameter lies in the interior of $\mathcal{P}$.*

- *The average forecast error across individuals $\overline{FE_{t+1}}$ is increasing in k for $\kappa = \mu, \rho, \phi$.*
  Similarly: $\lim_{k \searrow 0} \overline{FE_{t+1}} = 0$

- *The asymmetry of the response expectations to a signal is increasing in k for $\kappa = \phi$. Also:*

$$\underset{k \searrow 0}{lim} \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s_t^i > 0} - \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s_t^i < 0} = 0$$

  For $\kappa = \mu,\ \rho$ this difference equals 0, irrespective of k.

Proof: See Appendix B.1.

It should be said that this is not the only way to generate such a phenomenon. Maccheroni et al. (2006) establish the formal link between *preferences* over ambiguity (as in the smooth ambiguity framework of Klibanoff et al. (2005)) and the *degree* of ambiguity - i.e. the diameter of the set of ambiguity - as in the framework here[11]. This is important for the work of Section 2.4, where we have a measure of ambiguity preferences (rather than the degree of ambiguity). It could therefore be the case that richer individuals typically have a greater *taste* for ambiguity, whereas poorer households display greater ambiguity aversion. In Section 2.5, we provide evidence that while households may differ in their preferences towards ambiguity, there is no evidence that this is correlated with income.

[10] One possible way of modelling this for ambiguity over signal information content ϕ: suppose that at the beginning of the period, before the signal has been received, households can exert effort to reduce the degree of ambiguity in ϕ. After making their effort choice, the household receives its signal of future inflation and adjusts their expectations accordingly. More concretely, let $\mathbb{P}(n_i)$ denote the set of priors household i has over the signal variance, if they exert effort n_i (i.e. post effort decision, but pre-realisation of the signal). We can define:

$$\mathbb{P}(n_i) = \left[\phi - \chi(n_i),\ \phi + \chi(n_i)\right] \equiv \left[\underline{\phi}(n_i),\ \bar{\phi}(n_i)\right]$$

where $\chi(x)$ is a decreasing function, with $\lim_{x\to\infty} \chi(x) = 0$ and $\lim_{x\to 0} \chi(x) = k < \phi$. In other words, the set of priors for household i is centred on the true value of the signal to noise ratio, and is a wider interval for low values of effort. As effort increases, the size of the interval decreases; in the limit, the interval is a singleton at ϕ and there is no ambiguity.

[11] They also formally establish the connections between Gilboa and Schmeidler (1989) preferences and Hansen and Sargent (2001) preferences. In Chapter 3 we use the latter to model forecast pessimism, due to its tractability.

2.2.2 Two Testable Predictions

Proposition 2.2.1 gives us our two testable empirical hypotheses. If we were to find evidence of forecast optimism (or indeed neutrality) across ambiguity-averse individuals, this would suggest that ambiguity aversion is not a possible cause of the persistent forecast errors observed in the empirical data. This first test, therefore, is a test of the plausibility of ambiguity aversion driving forecast pessimism. Section 2.4 tests this hypothesis.

For our second test, we measure the response of expectation towards x on receipt of a signal of its future path. If we find evidence of an asymmetric response to a given signal, this is inconsistent with theories which use either ambiguity over the long-run mean μ or persistence ρ to explain forecast pessimism. Conditional on there being evidence that ambiguity aversion may be a cause of forecast pessimism, this test therefore refines the possible ambiguity environments. If expectations are asymmetric to a signal, then if ambiguity aversion causes forecast pessimism across households, they must at the very least experience ambiguity concerning the strength of the signals they receive. Of course, this does not rule out that they also face ambiguity over the long-run mean μ or persistence ρ as well. Section 2.6 tests this second hypothesis.

2.3 Data

In this Section, we detail the data sources used throughout the subsequent sections. In Section 2.4 we test our first hypothesis: that more ambiguity-averse individuals display larger forecasting errors. We employ the RAND American Life Panel and its associated modules to investigate this question.

In Section 2.6 we investigate our second hypothesis: that forecast errors respond asymmetrically to public signals. We employ the Michigan Survey of Consumers (MSC) in conjunction with the Survey of Professional Forecasters (SPF).

2.3.1 Ambiguity Aversion and Forecast Errors Data

To investigate the relationship between ambiguity aversion and expectation formation, we use the RAND American Life Panel (ALP), a representative survey of U.S. adults. The American Life Panel (ALP) is a survey administered by the RAND Corporation. The ALP consists of several thousand households that respond to a variety of surveys ("modules") on an ad hoc basis. The majority of surveys are conducted online, but a small subset of households who lack internet access are given a laptop to answer, to limit selection bias. The survey is designed to be representative of the U.S. adult population, and survey weights are provided to achieve this goal. When respondents are first admitted to the survey, basic demographic information is collected. These data are updated infrequently, but other modules can be used to gain more timely information on respondent characteristics if required. For more detail on this, please see Appendix B.2. Details on the modules we employ from this survey follow.

Ambiguity Aversion Preferences. We employ a survey conducted between March and April 2012. Respondents were presented with the Ellsberg urn problem Ellsberg (1961). We refer to this as the Ellsberg module. The module was initially used by its surveyors in two papers, Dimmock et al. (2013), and Dimmock et al. (2015)[12]. The survey elicits the probability p that makes the respondent indifferent between a lottery $\mathbb{K}$ with a win probability of p and a lottery $\mathbb{U}$, which has the same outcomes but unknown probabilities (i.e., is an ambiguous lottery). The respondents were incentivised with a \$15 payment when performing the Ellsberg experiment[13]. We also use several control variables from this module. For more detail on these, please see Appendix B.2.

Inflation Expectations. The Ellsberg module does not contain information on economic expectations. We therefore employ a second set of modules conducted by the New York Federal Reserve. These "mini-modules" formed part of the Household Inflation Expectations Project (HIEP), consisting of surveys fielded every six weeks from May 2008 to November 2012. These modules were a precursor to the New York Fed's Survey of Consumer Expectations (SCE) and elicited inflation expectations similarly, more precisely, they contain information on point- and density-based forecasts for year-ahead inflation. The benefit of using these surveys is that they can be matched at the individual level to the Ellsberg module, allowing us to link a measure of ambiguity aversion to measures of economic expectations. Of the 3,268 individuals who completed the Ellsberg module, 1,651 also completed at least one of the NY Fed mini-modules. For more detail on these surveys, and the HIEP more generally, please see Bruine de Bruin, Vanderklaauw, et al. (2010) or Sections 1 or 2 of Armantier, Bruine de Bruin, Potter, et al. (2013). For further detail on how we employ these modules, please see Appendix B.2.

Wealth Data. In Section 2.5, we require more detailed income and wealth data than provided in any of the two previous modules or the baseline ALP surveys. We therefore supplement the above data with information from the Health and Retirement Study (HRS) in the ALP. The ALP HRS is a user-friendly version of the U.S. Health and Retirement Study, harmonised and cleaned by the RAND Corporation for ease of use by researchers. It is a longitudinal panel survey that tracks a representative sample of

[12] They establish that respondents with a greater degree of ambiguity aversion have lower stock market participation, and lower portfolio shares of risky assets. These findings are supporting evidence for the model discussed in Chapter 3.

[13] For more detail on the survey, as well as the elicitation algorithm used, please refer to the appendix of Dimmock et al. (2015).

Americans aged 50 and above. For our purposes, it contains more detailed information on income and wealth. All respondents from the previous two sets of modules are included in the HRS data. See Appendix B.2 for further detail.

2.3.2 Asymmetric Expectation Responses Data

In Section 2.6, we utilise data from two widely cited surveys in the literature. We employ the Michigan Survey of Consumers, using its expectations data as our left-hand side variable. We also use the Survey of Professional Forecasters as a proxy for changes in news of the future. In more detail:

Expectations Data. The Michigan Survey of Consumers (MSC) is a long-running, high-frequency survey conducted by the University of Michigan that collects data on the economic expectations and sentiment of U.S. households. It has been fielded monthly since the late 1940s. The survey underwent a major revision in 1978, with additional questions added sporadically over time. For our purposes, we use the post-1978 sample. The survey typically receives around 500 respondents per month. We use this survey as our primary measure of inflation, unemployment, and interest rate expectations; it is more suitable than, for example, the SCE for our purposes, due to its larger sample. It contains demographic information covering the respondents, which we use as controls. These are: gender, education, census region and age. The survey is designed to be representative of the U.S. population, and survey weights are provided. For more detail on the wording of the expectations questions, and the precise definitions of the control variables, please see Appendix A.1.

Constructing a Signals Proxy. The Survey of Professional Forecasters (SPF) is a quarterly panel of U.S. macroeconomic forecasts that began in 1968, which surveys professional forecasters across a range of macroeconomic variables. For each survey round, it reports the cross-sectional distribution of professional forecasts. For our purposes, we focus on the forecasts for one-year ahead inflation, the national unemployment rate and the one-year ahead 3-month rate, which were first available in 1981Q3. We employ the median professional forecast for each variable at time t. We present further detail on this in Section 2.6.

Having outlined the data sources used throughout the remaining sections, we now examine whether more ambiguity-averse individuals exhibit a greater degree of forecast pessimism.

2.4 Experimental Measure of Ambiguity

In this section, we aim to demonstrate that these three patterns are identical to those observed in ambiguity aversion. That is, the expectations of more ambiguity-averse individuals show similar properties to those of low-income individuals in three popularly used consumer surveys. Importantly, all of the results of this section hold when we also control for income.

As in Dimmock et al. (2013), we use the Ellsberg module to construct a measure of ambiguity aversion AA using the indifference probability m, where $AA = 50\% - m$. A positive measure is consistent with ambiguity aversion, while a negative measure is consistent with ambiguity-seeking behaviour. Within the Ellsberg module, 53% of respondents are ambiguity averse, 12% are ambiguity neutral, and 36% are ambiguity seeking.

We take our inflation expectations data from the HIEP modules. We construct a measure of forecast errors by subtracting the realised BLS CPI inflation rate from the 12-month ahead forecast of inflation. We therefore define our inflation forecast error variable as $\text{FE}_{it} \equiv \pi^e_{i,t+12|t} - \pi_{t+12}$. Forecast errors over other prices are constructed similarly. For unemployment, the survey asks respondents for the "*percentage chance that 12 months from now the unemployment rate in the US will be higher than it is now?*". We construct a measure of individual i's unemployment forecast error using a variable which takes the value one if the BLS national unemployment rate is indeed higher than 12 months in the future, as our measure of the outcome. The question and therefore procedure for interest rate expectations is the same. For detail on the inflation elicitation in the HIEP, please see Appendix B.2.

Empirical Findings. We consider how inflation forecast errors vary across ambiguity-averse individuals (i.e. those with $AA > 0$). As discussed in Section 2.2, individuals with greater ambiguity aversion should exhibit higher inflation expectations on average if they perceive inflation to be associated with lower future welfare. They would therefore also be expected to display larger (positive) forecast errors, given that they form more pessimistic expectations. For each observation, we construct a measure of the one-year ahead inflation forecast error as we did in Section 1.4. We then regress the inflation forecast error on the Ellsberg measure of ambiguity aversion and a selection of controls.

Here, we use the same set of demographic controls as we did in Chapter 1[14]. Above and beyond demographics, we may be concerned that individuals either misunderstand or quickly become fatigued

[14]They are: educational attainment, race, marital status, gender, labour force status, age, squared age, survey tenure and geographic region. The only difference between this set of controls and the set used in Chapter 1 Section 1.4 is that we can now only control for geography at the regional rather than state level.

by the Ellsberg Urn problem. To address this concern, we also include an additional set of controls provided by the Ellsberg module. The module includes questions designed to elicit financial literacy, as well as risk aversion and agreeableness[15]. We include these as controls. The module also asked a follow-up question to the Ellsberg question, designed to test for the consistency of the elicited ambiguity preferences. We therefore also include whether or not the check question was answered as expected, as a control, too. To help control for survey engagement, we also control for: the time spent answering the survey, whether the respondent found the questions to be clear, and whether the respondent was asked about their risk preferences first. We include time-fixed effects, as before[16]. For more detail on the ALP data, please see Appendix B.2.

Finally, and most crucially, we also control for income: we are interested in the effects of ambiguity-aversion on expectation properties *separate* from the effects outlined in Chapter 1. In our baseline set of regressions, we use income as provided by the ALP, which is infrequently updated[17]. However, in Section 2.5 we employ alternative measures of income, and our choice of control here is unimportant for the results.

We therefore specify the following regression:

$$\text{FE}_{it} = \beta_0 + \beta_{AA} AA_i + \gamma \boldsymbol{X}_{it} + d_t + \varepsilon_{it} \tag{2.1}$$

The left-hand side is our measure of the individual's forecast error, realised at time t. Our coefficient of interest is β_{AA}, which measures the marginal effect of an increase in ambiguity aversion on the inflation forecast error. We include time-fixed effects across all specifications to account for aggregate factors that influence inflation expectations. Unfortunately, since our measure AA_i does not exhibit time variation, we cannot employ individual fixed effects.

Table 2.1 shows how inflation forecast errors vary with ambiguity aversion, where each column corresponds to the inclusion of different controls. First, we observe that the coefficient on the degree of ambiguity aversion and inflation forecast error is positive and statistically significant in all specifications; that is, individuals who are more ambiguity averse also exhibit larger inflation forecast errors. Secondly,

[15]Financial literacy helps the researcher proxy for question understanding. Controlling for risk aversion ensures that the measure captures something separate from risk appetite. Finally, survey respondents may believe they are being deceived by the Researcher when performing the Ellsberg experiment, and so controlling for agreeableness helps to mitigate this potentially confounding effect.

[16]As all individuals in the Ellsberg module complete the NY Fed questionnaire in the same month, the inclusion of time-fixed effects is important. We are therefore implicitly assuming that AA captures fundamental (i.e. time-invariant) preferences.

[17]This measure of income was elicited in 2006, several years before the Ellsberg and NY-FED surveys were conducted. However, it is available to all respondents simultaneously. We therefore view this as a proxy for household economic status.

the magnitude of the coefficients is sizeable. Taking the baseline specification for the inflation point forecast, the marginal effect of an increase in ambiguity aversion is 1.2. Ceteris paribus, an ambiguity-averse individual at the 10th percentile versus an individual at the 90th percentile of AA (0.03 vs 0.47) has an inflation forecast error that is 0.52 percentage points larger. Finally, these results are true for both the point and density-based forecasts[18]

Table 2.1: Marginal Effect of Ambiguity Aversion on Inflation Forecast Errors

	Inflation (Point)			Inflation (Density)		
AA	0.919**	1.411***	1.188**	0.844**	1.441***	1.946***
	(2.26)	(2.91)	(2.21)	(2.44)	(3.46)	(4.34)
Obs	5,282	5,281	5,233	5,040	5,040	5,009
Adj. R^2	0.207	0.241	0.321	0.249	0.287	0.384
Mean dep var	1.57	1.57	1.56	1.13	1.13	1.12
Survey Controls		✓	✓	✓	✓	✓
Dem. Controls			✓			✓

Here we show the results of regressing inflation forecasting errors on ambiguity aversion. Columns 1) and 4) include only time fixed effects, whereas Columns 2) and 5) also include the set of survey controls. These are: financial literacy, checking for consistent Ellsberg response, asking questions about risk first, risk aversion index, finding questions clear, agreeableness, and indicators for whether the survey was completed in under 10 minutes or over 24 hours. Columns 3) and 6) also include demographic controls and controls for real income. These controls are: educational attainment, race, marital status, gender, labour force status, age, squared age, survey tenure and geographic region. The first three columns use the point forecast implied forecast error as the dependent variable, and the final three columns use the density-implied forecast error. We cluster standard errors at the individual level, and report the t-statistic in parentheses. All results employ (Ellsberg module) survey weights. One, two or three stars denote significance at the 10%, 5% and 1% level, respectively.

Across other variables. Is this unique to inflation expectations, or a feature of expectations more generally? We replace the right-hand side of regression (2.1) with forecast errors constructed over other aggregate variables. In particular, the NY Fed also elicits one year-ahead (point) forecasts over gas price inflation, food inflation, U.S. government debt growth, national house price growth and unemployment rate changes[19]. The only caveat to add is that the regression for house price forecast error is conducted over the sub-sample of homeowners only[20]. We also invert this measure, so that a positive error is consistent with a notion of expectation' pessimism': an increase in the left-hand side variable for house prices is therefore associated with homeowners' expected lower house price increase than actually occurs.

Again, we can see that the magnitude of the forecast error is increasing for all variables considered, and

[18]See Chapter 1 Section 1.4.1 for more.

[19]We construct this forecast error as in Chapter 1: we assign a value of 1, 0 and -1 to "increase", "stay the same" and "decrease" for unemployment expectations. We then measure the outcome of changes in unemployment in a similar manner, where one year ahead unemployment is recorded as having "stayed the same" if it is within one standard deviation of current unemployment

[20]This avoids the problem of non-house owners perceiving increases in house prices to be welfare-negative, unlike homeowners who are likely to perceive it as positive for their net wealth. This is consistent with Chapter 1.

Table 2.2: Marginal Effect of Ambiguity Aversion on Other Forecast Errors

	Unemp.	Gas Pr.	Food Pr.	G. Debt Growth	House P.
AA	0.170***	5.510***	3.808***	4.593***	6.186***
	(3.22)	(3.83)	(4.49)	(3.41)	(4.66)
Obs	5,517	3,519	3,598	4,174	2,367
Adj. R^2	0.385	0.819	0.303	0.156	0.660
Mean dep var	.0448	-.866	2.04	.321	.865
Survey Controls	✓	✓	✓	✓	✓
Dem. Controls	✓	✓	✓	✓	✓

Here we show the same regression as Table 2.1, but now using forecast errors across different outcomes. All results employ survey weights. One, two, or three stars denote significance at the 10%, 5% and 1% level, respectively, and t-statistics are shown in parentheses. House price forecast errors are calculated over the sub-sample of homeowners. This measure has been inverted, as in section 1.4, so as to be consistent with the notion of "pessimism".

this effect is significant at the 5% level. Quantitatively, the effect of ambiguity aversion is significant for all variables. For example, a mildly ambiguity-averse respondent at the 10th percentile has a food price forecast error roughly 2.5% higher than a strongly ambiguity-averse respondent at the 90th percentile.

In summary, the combined ALP-NYFed data appears to lend credence to the possibility that ambiguity aversion is, at least in part, a possible explanation for why consumers may form pessimistic expectations. More ambiguity-averse individuals exhibit higher forecast errors across not only inflation expectations but also multiple other economic variables.

2.5 Income and Ambiguity Aversion

Across differing ambiguity-averse preferences, we witness the same set of patterns as we do across income groups in the consumer survey data. This then raises the question: are the results of Section 1.4 driven by preferences over ambiguity? That is, are low-income households more likely to be ambiguity-averse (for some reason), and are differences in expectations across the income distribution driven by heterogeneity in preferences? Or is it simply that lower-income households behave *as if* they are more ambiguity-averse, but underlying preferences are uncorrelated with income? We consider these possibilities in this section.

In Section 2.4, we controlled for income using the ALP income data, which was measured several years before either the Ellsberg module or NYFed modules were conducted. Therefore Section 2.4 identifies the effect of ambiguity-aversion on forecast errors independent of income. However, if a significant relationship exists in the experimental data between ambiguity aversion and income, this raises the possibility that it is ambiguity aversion (rather than possibly because caused by a households economic condition itself) that drives the empirical relationship between income and forecasting errors[21]. That is, if we find evidence that ambiguity preferences and income co-vary, it may be preference heterogeneity driving the relationship between income and pessimistic expectations, rather than a feature of a household economic conditions itself. We combine our ALP-HIEP dataset with a module of the RAND Health and Retirement Study (HRS) conducted on the ALP sample, as this survey contains more detailed questions covering both income and wealth for all respondents within the sample. While the measures of income and wealth taken from the HRS are observed much closer to the time of the Ellsberg module (2013 versus 2012) than the income measure provided by the ALP (2004), they are nonetheless still not concurrent.

Since we have no variation over time for each individual in our income measure or AA_i, we do not perform a panel regression, as in the previous section. Instead, we perform simple OLS. We therefore consider the following regression:

$$AA_i = \beta_0 + \beta_{\text{inc}} \log(\text{income}_i) + \gamma \boldsymbol{X}_i + \varepsilon_i \tag{2.2}$$

Unlike in the previous section, where we measured the effect of ambiguity aversion on forecast pessimism, we now attempt to use household income to explain the difference in AA_i. Our coefficient of interest is therefore β_{inc}, which measures the marginal effect of income on ambiguity-aversion. We include $\boldsymbol{X}_{it}$ as the vector of controls.

Table 2.3: Marginal Effect of Income on Ambiguity Aversion

	(1)	(2)	(3)
log(income)	-1.964***	-0.893	-0.667
	(-3.40)	(-1.30)	(-1.03)
Obs	1,499	1,383	1,381
Adj. R^2	0.015	0.041	0.220
Mean dep var	17.2	17	17
Survey Controls			✓
Dem. Controls		✓	✓

Here we show our measure of ambiguity aversion (AA) regressed on (log) real income. Column 1 contains only time-fixed effects, Column 2 adds demographic controls, and Column 3 adds survey controls. One, two, or three stars denote significance at the 10%, 5% and 1% level, respectively. T-statistics are shown in parentheses. All regressions employ survey weights.

A significant benefit of the HRS data is that it provides us with a continuous measure of income. We first regress our measure AA on (log) personal income as measured in the RAND HRS module. The results are displayed in Table 2.3. The data display a significant, albeit small, correlation unconditionally.

[21] As discussed in Chapter 1.

A 10% increase in real income is associated with ambiguity aversion falling by 0.19 (0.014% of the standard deviation of AA). When we control for survey-based and demographic variables, however, this correlation fades away. In Appendix B.3 we perform a similar exercise using the categorical ALP income data, and find an almost identical result. In the Ellsberg data, the fact that the (linear) marginal effect of income on ambiguity-aversion is zero is not particular to the RHS measure.

Naturally, we may be concerned that the relationship between income and AA_i as measured is possibly non-linear. We therefore turn to a non-parametric method to test whether this is the case. In Figure 2.1, we use a bin-scatter to illustrate the relationship between income and ambiguity aversion non-parametrically. The black crosses (with a white error band) denote the result of the bin-scatter without controlling for any covariates other than time fixed-effects. The blue crosses (with blue confidence bands) illustrate the same bin-scatter, but with the full range of demographic and survey controls included, as in Section 2.4. Visually, the binned scatter does not display quantitatively meaningful covariance between ambiguity preferences and income. This provides reassurance that the significant (but small) result of specification (1) in Table 2.3 is not an issue of small sample size.

Figure 2.1: Bin-scatter of Ambiguity Aversion Across Income

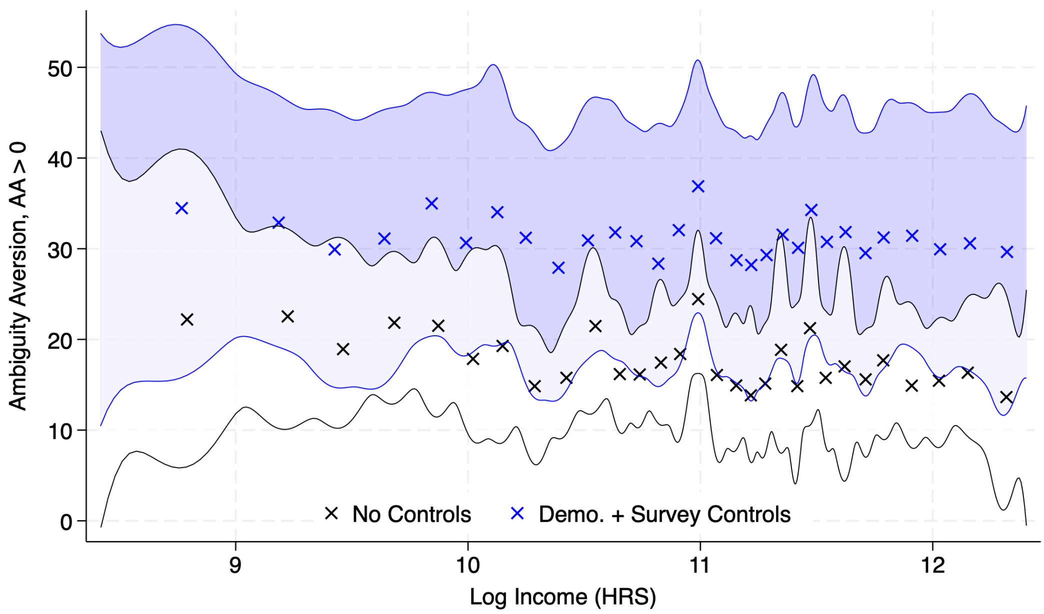

Here, we plot two binned scatter plots of (log) real income on our measure of ambiguity aversion. Here, the income variable used is taken from the RAND-HRS module, denominated in 2015 dollars. The black crosses denote a specification where we control for time fixed effects only. The blue crosses include the same bin-scatter, but include the full set of demographic and survey controls. Each specification has been put into 30 bins. The white-shaded area denotes a non-parametrically estimated 95% confidence interval for the specification with only time fixed effects, and the blue band represents a 95% confidence interval for the specification including controls. These confidence bands are calculated following Cattaneo et al. (2024) using standard errors clustered at the individual level in the first-stage regression. In Appendix B.3 we display a similar graph, but over *all* individuals in the Ellsberg module.

The finding that ambiguity preferences are not correlated with income is consistent with the evidence presented in Dimmock et al. (2015), who find that the (full sample) measure of ambiguity aversion correlates very little with demographic or economic variables. This may reassure the reader that the survey is measuring deep underlying preferences, rather than picking up on an individual's circumstances. In Appendix B.3 we take this analysis further and test several other measures of income and wealth in the specification of (2.2), as well as over sub-samples of the data[22]. We find that none of these measures appears to show a significant or sizeable relationship between income and ambiguity aversion.

We draw from this the following conclusion: that while ambiguity-averse individuals and low-income households display similar expectation properties in the data, it is likely *not* due to differing underlying preferences. That is to say, the world in which households have different preferences, and those with higher ambiguity aversion find themselves in lower income circumstances by some mechanism, appears not to be the case. Were it so, and the estimates of ambiguity preferences were reliable, then we would expect income (or wealth, perhaps) to be correlated with our measure. Instead, this evidence suggests that low-income households *behave as if* they are more ambiguity averse than more affluent households.

None of the results of the previous two sections point towards the *source* of ambiguity that consumers may face. In the next section, we therefore explore the response of consumer expectations to public signals.

2.6 Expectation Asymmetry

As discussed in Section 2.2, if consumers receive ambiguous signals, we should expect the response of expectations to be larger for bad public signals (which is taken to be informative) than for good public signals (taken to be uninformative), for the consumer. This asymmetry should be stronger for poorer consumers and weaker for richer consumers if the effective degree of ambiguity varies by income group. In this section we examine whether or not consumer expectations as measured in surveys do indeed appear to respond asymmetrically to signals.

2.6.1 Survey Data

To identify the effect of public signals, we require a long time dimension. We therefore use the MSC in this section. In order to test this for this asymmetry, we perform a similar exercise to Baqaee (2020). In

[22]Our alternative measures of income and wealth are grouped personal income, total household income, net wealth, net financial wealth, gross assets and net housing wealth. We test the significance of these variables, controlling for survey and demographic factors, over four subsamples of the data. These are the ambiguity averse, the full sample, and then both of these, excluding all respondents with less than $500 in net wealth.

our toy model, agents solve a simple linear filtering problem and so combine information contained in the signal they receive with past information. Therefore consider the following reduced form model of expectation of variable y:

$$\mathbb{E}_t^i[y_{t+12}] = \text{past info}_t + \beta_+ \cdot \text{signal}_t^+ + \beta_- \cdot \text{signal}_t^- + C_{i,t}$$

where positive signal and negative signal both represent new information as to the future direction of y, and $C_{i,t}$ is a set of controls. In order to generate a proxy for public signals, the Survey of Professional Forecasters (SPF) is used[23]. For each variable y, we consider the public signal to be positive if the SPF median forecast $\mathbb{E}_t^{\text{SPF}}[y_{t+12}]$ is greater than the current value y_t, and negative is the reverse is true. That is, our proxy for public signals is:

$$\text{signal}_t = \mathbb{E}_t^{\text{SPF}}[y_{t+12}] - y_t$$

Appendix B.4 provides more detail on the construction of our proxy for public signals and provides plots[24].

As Baqaee (2020) discusses, using professional forecasts is a reasonable candidate for the public signal received by households. As long as professional forecasts perform better than households at forecasting on average [25], households will find our measure signal_t to be informative, on average. We do not claim that households are aware of the median professional forecast in any given period, only that our SPF measure will reflect other possible sources of information of the future, which households are privy too. To this extent, our measure serves only as a proxy for the public signals received by households.

Figure 2.2 presents bin-scatters of MSC[26] expectations of inflation, interest rates changes and unemployment changes[27]. All three variables illustrate an asymmetric response to our proxy for public signals, with the response of expectations flatter in times when professional forecasters expect decreases in y_t

[23]This is a quarterly survey of professional US forecasters conducted by the Federal Reserve Bank of Pennsylvania, covering a wide range of variables. It is available since 1981.

[24]An alternative specification might be to use SPF forecast revisions instead, to possibly reflect "news". In our regression specification (2.3) we control for past information, including the lagged median forecast of the SPF. We use this specification as it allows for direct comparison to Baqaee (2020).

[25]Ang et al. (2006) finds the median of the SPF to outperform financial market implied forecasts, as well as a range of statistical models up to the three year horizon.

[26]For this section we use the MSC as it is a much longer sample than the SCE. Over the SCE's sample (2013-) there is very little variation in public signals over future inflation, unemployment or interest rate changes to be able to accurately estimate the regression model of this section.

[27]Inflation is a point forecast, whereas for interest rates (here, borrowing) and unemployment consumers report whether they believe the variable will increase stay the same, or decrease. Consumers are asked whether they think borrowing rates will change. As the SPF does not ask forecasters for their expectations for a particular consumer borrowing rate, we use their forecast of the 3-month Treasury bill as a proxy.

than increases. In the specification employed in Figure 2.2, we use the current value of the outcome variable to control for the current information set. These results are confirmed more formally in Appendix B.4, where we perform hypothesis tests on the marginal effect of positive and negative public signals. We find that these slopes are all statistically different to one another at (at least) the 5% level.

Figure 2.2: Asymmetric Response of Expectations to the Proxy for Public Signals

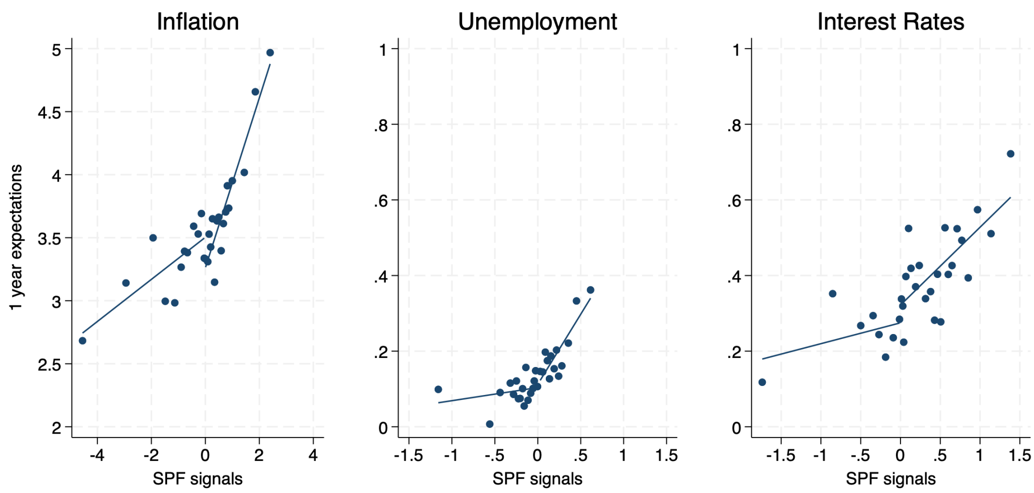

Here we show binned scatter plots of inflation expectations, unemployment rate change expectations and interest rate change expectations in the MSC against our proxy measure for public signals. Our inflation measure continues to be trimmed at the 5-95% level, as before. These expectations are measured at a quarterly frequency, where we aggregate using the mean expectation for a given respondent over the quarter. Each panel shows expectations binned into 30 groups, and controls for past information (using the current value of the appropriate variable) as well as individual fixed-effects. We fit two separate regression lines: one for positive values of public signals, and one for negative.

Across income groups. The results of Chapter 1, amongst others, suggests that forecast pessimism declines as income increases. Therefore if ambiguity aversion is a possible driven of persistently positive forecast errors, we should see that this asymmetry declines with income (as in Section 2.5 and discussed in Proposition 2.2.2). We now turn to investigating whether or not this asymmetry changes between income groups.

To do so, we run the following regression.

$$\mathbb{E}_t^i[x_{t+12}] = \sum_k \text{signal}_t^+ \cdot \left(\beta_+ + \beta_+^k \cdot \text{income}_{it}^k\right) + \text{signal}_t^- \cdot \left(\beta_- + \beta_-^k \cdot \text{income}_{it}^k\right) + \gamma_p \boldsymbol{X}_t^{\text{past}} + \gamma_c \boldsymbol{C}_{it} + \varepsilon_{it} \tag{2.3}$$

Here, we define $\text{signal}_t^+ = \text{signal}_t \cdot \mathbb{I}(\text{signal} \geq 0)$, that is a variable which takes the value of signal_t when the public signal is positive, and 0 otherwise. We define signal_t^- similarly. income_{it}^k is a indicator variable that takes the value 1 if individual i is in income group k. We measure income groups by self-reported income quintile by year, in 2015 dollars. We include the current value of the variable (x_t), the lagged value of the variable (x_{t-1}) and lagged SPF forecast $\mathbb{E}^{SPF}_{t-3}[x_{t+9}]$ to control for past information[28]. Our vector of controls includes the following demographic characteristics: age, squared age, marital status, education level, geographic region. Our coefficients of interest are the β's. The expression $\beta_+ + \beta^k_+$ captures the response of expectations to positive public signals for an individual in the k-th income quintile. Similarly, $\beta_- + \beta^k_-$ captures the response to negative public signals. The asymmetry of response is given by the difference between these two expressions.

If ambiguity over the information content of signals *is* the underlying cause of consumer forecast errors, and these decline in income, then it would be the case that this asymmetry also declines in income. If this hypothesis is correct, we should expect: 1) $\beta_+ + \beta^k_+$ to be positive but decreasing in k; 2) $\beta_- + \beta^k_-$ to be positive but increasing in k; 3) the difference between these to be positive and declining in k. A difference of 0 suggests that consumer expectations behave symmetrically to positive and negative public signals. In Figure 2.3 we plot these three expressions for each of the three variables. Appendix B.4 displays the raw regression results. Inflation expectations display the starkest results. The response of expectations to positive (negative) signals is strongly declining (increasing) in income. Those in the bottom 20% of the income distribution display a strong asymmetry to inflation, whereas those in the top 20% react only slightly more to positive signals than negative. For interest rate expectations and unemployment expectations, the pattern is similar, although less pronounced. Unemployment expectations in times of increasing-unemployment public signals do not change across income groups, although the response of expectations in times of decreasing signals is strongly increasing. This is sufficient to generate a similar asymmetry in unemployment expectations, where the top two quintiles of individuals by income display no asymmetry. For the top 20% of respondents by income, the expectation response to public signals is asymmetric in the opposite direction to what was expected, although mild. The response of interest rate expectations in times of increasing-interest rate signals is mildly increasing in income, but strongly increasing for times of decreasing-interest rate signals. Again, this is sufficient to generate a declining asymmetry of expectations as income increases.

We may be concerned that the asymmetry of expectation response is driven by an asymmetry in our proxy for public signals. Appendix B.4 provides formal tests for skewness. We cannot reject the null of zero skewness in the distribution of our public signals proxies at the 5% level. It appears that the asymmetry of the response of expectations, across all three variables, is driven by a factor determined

[28] As interest rates and unemployment expectations are measured as differences, whereas inflation is a point-forecast, we also include the current inflation rate π_t in our regression for inflation. Appendix B.1 includes a specification where we include time-fixed effects instead. The results are qualitatively the same.

Figure 2.3: Marginal Effect of Public Signals Proxy on Expectations, By Sign

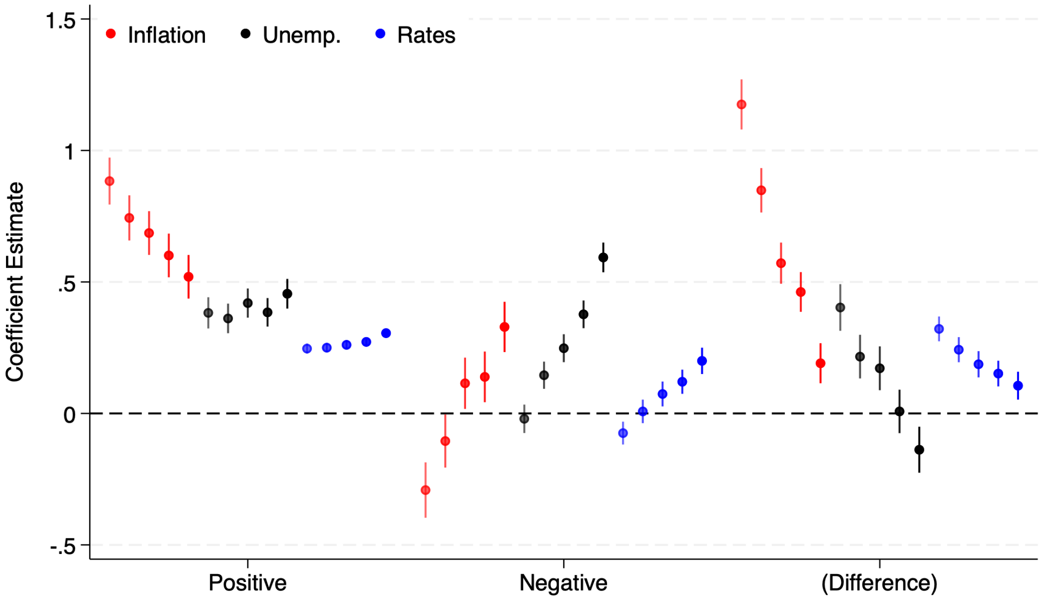

This graph reports combine coefficient estimates for our regression specification across the three variables. The 'Positive column reports $\beta_+ + \beta_+^k$ across k, where the bottom income quintile is leftmost, and the top income quintile is rightmost. the Negative column reports $\beta_- + \beta_-^k$ for each k, and the Difference column reports $(\beta_+ + \beta_+^k) - (\beta_- + \beta_-^k)$ for each k. The regressions control for past information and demographic variables and include survey weights. The lines represent the 95% confidence interval of the hypothesis test that the combined coefficients equal 0 (i.e. that $\beta_+ + \beta_+^k = 0$). These are calculated using standard errors clustered at the individual level. For example, we can be 95% confident that unemployment expectations respond positively on receipt of (proxied) inflationary public signals. For the 60th-80th percentile by income however, we fail to reject the null hypothesis that expectations respond asymmetrically to our proxy measure of public signals - i.e. we cannot reject that $\beta_+ + \beta_+^{80-100} = \beta_- + \beta_-^{80-100}$.

at the individual level.

2.6.2 Expectation Asymmetry in the ALP Data

We now seek to demonstrate that this final set of results—specifically, that consumer survey expectations of inflation, changes in the unemployment rate, and changes in interest rates display an asymmetric response to public signals—also holds for ambiguity-averse individuals. For ambiguity neutral individuals, the magnitude of the response of expectations to signals should be independent of the sign of the signal. However in an ambiguous environment as in Section 2.2, household expectations should respond more to signals which indicate welfare may be lower in the future than signals which indicate welfare may be greater. Importantly, this discrepancy increases with ambiguity aversion.

In order to establish to test this hypothesis, we perform an exercise similar to that of Section 2.6. Here, we focus on inflation expectations alone. As the NY-FED modules were completed over a relatively short time frame, only our SPF public signals proxy for inflation contains a reasonable amount of variance[29]. Otherwise, inflationary and dis-inflationary signals are defined as before. We use a similar reduced form-model for inflation expectations and test whether more ambiguity-averse consumers respond more asymmetrically to signals than less ambiguity-averse consumers.

$$\mathbb{E}_t^i[\pi_{t+12}] = \text{signal}_t^+ \cdot \left(\beta_+ + \beta_+^{\text{AA}} \cdot \text{AA}_i\right) + \text{signal}_t^- \cdot \left(\beta_- + \beta_-^{\text{AA}} \cdot \text{AA}_i\right) + \gamma_p \boldsymbol{X}_t^{\text{past}} + \gamma_c \boldsymbol{C}_{it} + \varepsilon_{it}$$

We use a similar specification for past-information as before. However within our set of controls $\boldsymbol{C}_{it}$ we add the additional survey-design controls discussed in the previous section, alongside income. In a second specification, we instead use time fixed effects to control for the past information set instead. Please refer to Appendix B.4 for the regression output.

In a baseline specification households revise their expectations more strongly to positive signals than to negative signals. In this data, the effect of negative signals is statistically indistinguishable from zero. Testing the null hypothesis that $\beta_+ = \beta_-$ is rejected in both specifications sets of specifications at the 1% level. Moreover, the signs of β_+^{AA} and β_-^{AA} are in agreement with our hypothesis that households that are more ambiguity averse should show more asymmetric responses. We see that more ambiguity averse households have higher responses to inflationary public signals, and less of a response to disinflationary signals. Testing the null hypothesis that $(\beta_+ + \beta_+^{\text{AA}}) = (\beta_- + \beta_-^{\text{AA}})$ is also rejected at the 1% level in the baseline specification.

In Figure 2.4, we plot the difference between the response of expectations to inflationary signals and dis-inflationary signals against ambiguity aversion. As expected more ambiguity averse individuals respond more asymmetrically to public signals. For those individuals with low ambiguity aversion, signals of inflation increasing by 1% in the future increases expectations by 0.5% more than similarly sized dis-inflationary signals. For those with strong ambiguity aversion, this discrepancy in expectation response is roughly three times as large.

Although data limitations restrict us from performing similar exercises for expectations of unemployment changes and interest rate changes, this section provides further supporting evidence that low-income households as survey in consumer surveys have expectations with properties we would expect to see of ambiguity-averse individuals.

[29] Again, we might be concerned that for the period in question the proxy for inflationary signals is asymmetric in and of itself. Tests of both the mean and the skewness of the distribution of inflationary signals in the sample period cannot reject the null that either are zero.

Figure 2.4: Marginal Effect of Ambiguity Aversion on Inflation Expectation Asymmetry

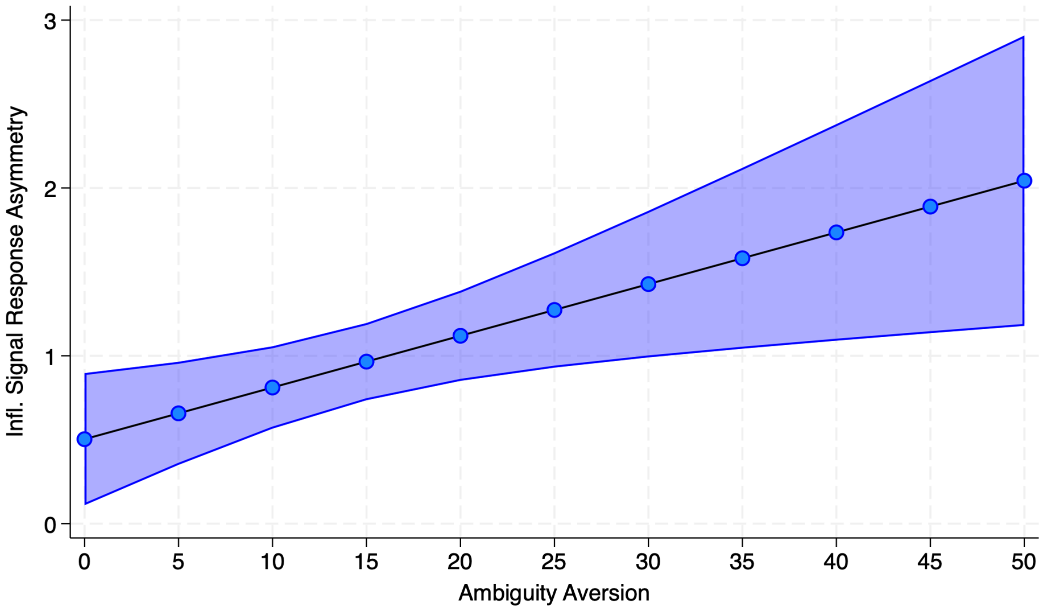

Here we plot $(\beta_+ + \beta_+^{AA} \cdot AA_i) - (\beta_- + \beta_-^{AA} \cdot AA_i)$, where the β's are taken from the baseline specification. The 95% confidence interval is shown, calculated using standard errors clustered at the indivdual level. The regression employs survey weights. Our measure AA has been multiplied by 100 for ease of interpretation.

2.7 Conclusion

This paper set out to determine whether households' persistent, upward-biased inflation forecasts can be rationalised by ambiguity aversion, and, if so, which parameters of the data-generating process households perceive as ambiguous. We derived two separating tests: (i) more ambiguity-averse agents should exhibit larger unconditional forecast errors, and (ii) only ambiguity about the information content of signals (ϕ) generates an *asymmetric* response of expectations to good versus bad public signals. Using the RAND American Life Panel we linked experimentally-elicited Ellsberg preferences to individual forecasts and showed that forecast errors rise monotonically with measured ambiguity aversion across inflation, unemployment, interest rates, and several price indices. With MSC data and an SPF-based public signals proxy, we documented a pronounced asymmetry in the response of expectations to signals of the future, a relationship that declines sharply with income. This mirrors the pattern predicted by ambiguity over ϕ when poorer households face a wider prior set. Crucially, we found no systematic correlation between income (or wealth) and the Ellsberg measure itself, implying that low-income households behave *as if* they are more ambiguity averse without actually possessing different deep preferences.

The marginal contribution is twofold. Empirically, we provide the first micro-level evidence connecting experimentally measured ambiguity preferences to survey forecast errors, thereby validating ambiguity aversion as a behavioural channel behind household expectation biases. Conceptually, we refine the ambiguity narrative by showing that only uncertainty about signal precision, and not about long-run means or persistence, can simultaneously explain both the level of forecast pessimism and the documented asymmetries in expectation updating. This sharpened identification helps reconcile disparate findings in the expectations literature and offers a disciplined way to incorporate Knightian uncertainty into macro-behavioural models without invoking ad-hoc belief distortions.

Several avenues remain open. First, experimental work that treats the precision of perceived signals, could provide a causal test of the asymmetric update channel identified here[30]. Second, dedicated panel-data tracking deep preferences and expectations at a higher frequency would allow us to test the drivers of forecast pessimism further. Third, papers exploring the relative plausibility of competing explanations of forecast pessimism would enable us to further test the validity of the ambiguity aversion hypothesis. Fourth, embedding the ϕ-ambiguity mechanism in a heterogeneous agent general equilibrium model would allow us to investigate and quantify how the cross-sectional distribution of expectations, and the *response* of expectations, impacts monetary transmission and welfare. This is especially pertinent if poorer households, who act *as if* they are more ambiguity-averse, bear disproportionate welfare losses. Finally, extending the analysis to firm expectations would of course be of value.

[30] Recent work using information treatments to study inflation expectations include Coibion, Georgarakos, et al. (2023).

B Supplementary Material for Chapter 2

B.1 Ambiguity Characteristics

For reference, we quickly outline again the environment. We have:

- AR(1) process: $x_{t+1} = (1-\rho)\mu + \rho x_t + \varepsilon_{t+1}$, where $\varepsilon \sim N(0, \sigma^2)$ and $\rho \in [0, 1]$.
- Signals $s_t^i = \varepsilon_{t+1} + \epsilon_{t+1}^i$, where $\epsilon_{t+1}^i \sim N(0, \sigma_s^2)$. We assume that ε_t and ϵ_t^i are uncorrelated i.i.d. noise.

Defining $\phi \equiv \frac{\sigma^2}{\sigma_s^2}$, then the rational forecast under Bayes rule is given by:

$$\mathbb{E}_t[x_{t+1}|\ s_t] = (1-\rho)\mu + \rho x_t + K(\phi) \cdot s_t \quad \text{where} \quad K(\phi) = \frac{\phi}{1+\phi} \ \text{ for } \ \phi > 0$$

Assume that welfare is monotonic and declining in x_t (as in the main text, for $x = \pi$). The Gilboa and Schmeidler (1989) max-min decision maker indexed by A picks the value of the ambiguous parameter $\kappa \in \{\mu, \rho, \phi\}$ that makes the expectation as large as possible (conditional on the signal s_t^i). That is:

$$\mathbb{E}_t^A[x_{t+1}|s_t] = \max_{\kappa} \ \mathbb{E}[x_{t+1}|s_t]$$

Define the forecast error at $t+1$ and the mean forecast error as in the main text:

$$FE_{t+1}^i \equiv \mathbb{E}_t^A \left[x_{t+1} \mid s_t^i\right] - x_{t+1} \ , \qquad \overline{FE_{t+1}}(x_{t+1}) \equiv \int_s FE_{t+1}(s_t^i|x_{t+1})\, dF(s_t^i)$$

$F(\cdot)$ is the cumulative density function of a normal distribution with mean zero and variance $\sigma^2 + \sigma_s^2$.

Proof of Proposition 2.2.1

We first restate Proposition 2.2.1 for reference.

Proposition. *For an ambiguity-averse individual i forecasting a variable x_t under Gilboa and Schmeidler (1989) preferences, assume that the agent perceives increases in x_t to be detrimental to welfare. Let x_t follow an AR(1) process and the agent receive signals s_t at time t over future innovations in x_t. Define the forecast error on receipt of a signal as $FE_{t+1}^i = \mathbb{E}_t^i[x_{t+1}|s_t^i] - x_{t+1}$, and the average forecast error across realisations of the signal as $\overline{FE_{t+1}} = \mathbb{E}_i\left[FE_{t+1}^i\right]$. Let $\mathcal{P}$ denote the ambiguity set over each parameter. For $\kappa \in \{\mu, \rho, \phi\}$, assume the true value of the parameter adheres to $\inf \mathcal{P} < \kappa < \sup \mathcal{P}$.*

- *Ambiguity over the long-run mean (μ), persistence (ρ) and signal strength (ϕ) all generate positive mean forecast errors. That is: $\overline{FE_{t+1}} > 0$.*

- *Only ambiguity over ϕ generates an asymmetric response of $\mathbb{E}_t^i[x_{t+1}|s_t^i]$ to a given signal. Formally, the response $\frac{d\mathbb{E}_t^i[x_{t+1}|s_t^i]}{ds_t^i}$ is asymmetric if $\frac{d\left|\mathbb{E}_t^i[x_{t+1}|k]\right|}{ds_t^i} \neq \frac{d\left|\mathbb{E}_t^i[x_{t+1}|-k]\right|}{ds_t^i}$ for $k \in \mathbb{R}$. For ambiguity over the long-run mean μ and persistence ρ, the response of an individual forecast is symmetric. For ambiguity over signal strength ϕ, the sensitivity of the forecast to a given signal is asymmetric. It is given by:*

$$\frac{d\mathbb{E}_t^i[x_{t+1}|s_t^i]}{ds_t^i} = \begin{cases} K(\underline{\phi}) & if \quad s_t^i \ < \ 0 \\ K(\bar{\phi}) & if \quad s_t^i \ \geq \ 0 \end{cases} \qquad where: \qquad K(\underline{\phi}) < K(\bar{\phi})$$

Proof. We take ambiguity over each of μ, ρ and ϕ in turn. For each case, we show $\overline{FE_{t+1}} > 0$. For μ and ρ we show that $\frac{d\mathbb{E}_t[x_{t+1}|s_t^i]}{ds_t^i}$ is linear; for ϕ we show that $\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i}$ is piecewise linear, with differing gradients across pieces. Let $\mathcal{P}$ denote the set of possible values of $\kappa \in \{\mu, \rho, \phi\}$. Define $\underline{\kappa} := \inf \mathcal{P}$ and $\bar{\kappa} := \sup \mathcal{P}$ for each of μ, ρ, and ϕ.

Ambiguity over μ. The ambiguity averse individual picks $\mu \in \mathcal{P}$ to maximise $\mathbb{E}_t[x_{t+1}|s_t^i]$. This is equivalent to solving $\max_{\mu \in \{\underline{\mu}, \bar{\mu}\}} \ (1-\rho)\mu + \rho x_t$. As $\bar{\mu} > \underline{\mu}$, we have:

$$\mathbb{E}_t^A[x_{t+1}|s_t] = (1-\rho)\bar{\mu} + \rho x_t + K(\phi) \cdot s_t$$

Therefore the forecast error is given by:

$$\text{FE}_{t+1}^i = (1-\rho)(\bar{\mu} - \mu) + K(\phi) \cdot s_t^i - \varepsilon_{t+1}$$

As $K(\phi) \cdot s_t - \varepsilon_{t+1}$ is mean zero, the mean forecast error is therefore given by: $\overline{\text{FE}}_{t+1} = (1-\rho)(\bar{\mu} - \mu) \geq 0 \ \ \forall \rho \in [0,1]$, as $\bar{\mu} \geq \mu$. This holds with strict inequality when $\bar{\mu} > \mu$.

Response to a signal. Finally, differentiating $\mathbb{E}_t^i[x_{t+1}|s_t^i]$ w.r.t. s_t^i yields:

$$\frac{d\mathbb{E}_t[x_{t+1}|s_t^i]}{ds_t^i} = K(\phi) \in [0,1]$$

Here, $K(\phi)$ is constant. Therefore FE_{t+1} is linear.

Ambiguity over ρ. The Gilboa and Schmeidler (1989) max-min decision maker picks the value of ρ that makes the expectation as large as possible. This is equivalent to solving $\max_{\rho \in \{\underline{\rho}, \bar{\rho}\}} (1-\rho)\mu + \rho x_t$. The worst-case expectation conditional on the noisy signal s_t is therefore:

$$\mathbb{E}_t^A[x_{t+1} \mid s_t] = \begin{cases} (1-\bar{\rho})\mu + \bar{\rho} x_t + K(\phi) \cdot s_t & \text{if } x_t > \mu \\ (1-\underline{\rho})\mu + \underline{\rho} x_t + K(\phi) \cdot s_t & \text{if } x_t \leq \mu \end{cases}$$

Using this alongside the data-generating process we can therefore write the forecast error as:

$$FE_{t+1} = \begin{cases} (\bar{\rho} - \rho)(x_t - \mu) + K(\phi) \cdot s_t - \varepsilon_{t+1} & \text{if } x_t > \mu \\ (\underline{\rho} - \rho)(x_t - \mu) + K(\phi) \cdot s_t - \varepsilon_{t+1} & \text{if } x_t \leq \mu \end{cases}$$

Next, note that $(x_t - \mu) \sim N\left(0, \frac{\sigma^2}{1-\rho^2}\right)$. It is also independent of the term $K(\phi) \cdot s_t - \varepsilon_{t+1}$, which is normally distributed. By Lemma B.1 we can write the forecast error unconditional on a particular realisation of a signal ($\overline{\text{FE}_{t+1}}$), as:

$$\overline{\text{FE}_{t+1}} = \sigma \cdot \frac{(\bar{\rho} - \underline{\rho})}{(1-\rho^2)^{\frac{1}{2}}} \cdot \left(\frac{2}{x}\right)^{\frac{1}{2}}$$

It is plain to see that $\overline{\text{FE}_{t+1}} > 0$ as long as $\bar{\rho} > \underline{\rho}$. Furthermore, if $\bar{\rho} - \underline{\rho} \to 0$, we have $\overline{\text{FE}_{t+1}} \to 0$.

Response to a signal. How do expectations $\mathbb{E}_t^A[x_{t+1} \mid s_t]$ depend on the realisation of a signal? The only term that directly depends on the realisation of the signal s_t is the last one. Differentiating with respect to the signal we get:

$$\frac{d\mathbb{E}_t^A[x_{t+1} | s_t^i]}{ds_t^i} = K(\phi) \in (0,1).$$

A more inflationary signal (higher s_t) lowers the subsequent forecast error because the household adjusts its forecast upward, albeit only partially. Note that this slope depends solely on the known signal to noise ratio ϕ; ambiguity about persistence affects the intercept of the forecast but not its responsiveness to the signal. That is, the response of expectations to a given realisation of a signal is linear and therefore symmetric.

Ambiguity over ϕ. Let the agent know that the true value of ϕ lies in $\mathcal{P} \subset \mathbb{R}_+$. The signal to noise ratio only affects the forecast of x_{t+1} through the Kálmán gain term $K(\phi)$. Note that $K(\phi) \in [0,1]$ and $K'(\phi) \geq 0 \;\; \forall \phi \geq 0$. Maximising the rational expectation is equivalent to maximising $K(\phi)$ if $s_t^i > 0$ and minimising $K(\phi)$ if $s_t^i < 0$. Therefore:

$$\mathbb{E}_t^A\left[x_{t+1}|\ s_t^i\right] = \begin{cases} (1-\rho)\mu + \rho x_t + K(\underline{\phi}) \cdot s_t^i & \text{if} \quad s_t^i \quad < \quad 0 \\ (1-\rho)\mu + \rho x_t + K(\bar{\phi}) \cdot s_t^i & \text{if} \quad s_t^i \quad \geq \quad 0 \end{cases}$$

$$\implies \text{FE}_{t+1}^i = \begin{cases} (1-\rho)\mu + \rho x_t + K(\underline{\phi}) \cdot s_t^i - \varepsilon_{t+1} & \text{if} \quad s_t^i \quad < \quad 0 \\ (1-\rho)\mu + \rho x_t + K(\bar{\phi}) \cdot s_t^i - \varepsilon_{t+1} & \text{if} \quad s_t^i \quad \geq \quad 0 \end{cases}$$

By Lemma B.1, after some short algebra one can show that:

$$\overline{FE_{t+1}} = \frac{1}{\sqrt{2\pi}} \cdot \frac{\bar{\phi} - \underline{\phi}}{(1+\bar{\phi})(1+\underline{\phi})} \cdot \left(\frac{\sigma^2}{K(\phi)}\right)^{\frac{1}{2}}$$

Note that as $\overline{\phi} - \underline{\phi} \longrightarrow 0$ we have $\overline{FE_{t+1}} \longrightarrow 0$.

Response to a signal. Differentiating $\mathbb{E}_t^A[x_{t+1}|s_t^i]$ w.r.t. s_t^i yields:

$$\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} = \begin{cases} K(\underline{\phi}) & \text{if} \quad s_t^i \quad < \quad 0 \\ K(\bar{\phi}) & \text{if} \quad s_t^i \quad \geq \quad 0 \end{cases} \qquad \text{where:} \qquad K(\underline{\phi}) < K(\bar{\phi})$$

Therefore as long as their is ambiguity (i.e. $K(\underline{\phi}) < K(\bar{\phi})$ the response of the forecast to a signal is asymmetric. Positive signals have a larger effect on the forecast error than negative ones.

This proves all parts of the proposition. □

Proof of Proposition 2.2.2

We restate Proposition 2.2.2.

Proposition. *For a Gilboa and Schmeidler (1989) agent, the degree of ambiguity the individual faces is captured by the diameter of the set of priors. Let $\mathcal{P}$ denote the set of possible values of $\kappa \in \{\mu, \rho, \phi\}$. Define the dimension of this set as $k := \text{diam}\,\mathcal{P} = \sup \mathcal{P} - \inf \mathcal{P}$. Assume that the true value of this parameter lies in the interior of $\mathcal{P}$.*

- *The average forecast error across individuals $\overline{FE_{t+1}}$ is increasing in k for $\kappa = \mu, \rho, \phi$.*
  Similarly: $\lim_{k \searrow 0} \overline{FE_{t+1}} = 0$

- *The asymmetry of the response expectations to a signal is increasing in k for $\kappa = \phi$. Also:*

$$\lim_{k \searrow 0} \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s_t^i > 0} - \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s_t^i < 0} = 0$$

  For $\kappa = \mu,\ \rho$ this difference equals 0, irrespective of k.

Proof. Let $\mathcal{P}$ denote the set of possible values of $\kappa \in \{\mu, \rho, \phi\}$. Define the dimension of this set as $k := \operatorname{diam} \mathcal{P}$. Let $\underline{\kappa} := \inf \mathcal{P}$ and $\bar{\kappa} := \sup \mathcal{P}$. Therefore, $k = \sup \mathcal{P} - \inf \mathcal{P} = \bar{\kappa} - \underline{\kappa}$. Assume that $\underline{\kappa} < \kappa < \bar{\kappa}$, where κ denotes the true value of the parameter.

- **Mean forecast errors.**

  - $\kappa = \mu$. We have $\overline{FE}_{t+1} = (1-\rho)(\bar{\mu} - \mu)$, which is linear in $\bar{\mu}$. With $\bar{\mu} = \underline{\mu} + k$ and $\mu \in [\underline{\mu}, \bar{\mu}]$ we have $\overline{FE}_{t+1} = (1-\rho)(\underline{\mu} + k - \mu)$. This is increasing in k, and vanishes as $k \longrightarrow 0$.

  - $\kappa = \rho$. We have $\overline{FE}_{t+1} = \sigma \sqrt{\frac{2}{\pi}} \frac{k}{(1-\rho^2)^{\frac{1}{2}}}$. This is increasing in k, and vanishes as $k \longrightarrow 0$.

  - $\kappa = \phi$. We have $\overline{\mathrm{FE}}_{t+1} = \frac{1}{\sqrt{2\pi}} \frac{k}{(1+\bar{\phi})(1+\underline{\phi})} \left(\frac{\sigma^2}{K(\phi)}\right)^{\frac{1}{2}}$. This is increasing in k, and if $k \longrightarrow 0$ then $\overline{\mathrm{FE}}_{t+1} \longrightarrow 0$.

  Hence $\overline{\mathrm{FE}}_{t+1}$ is monotone in the diameter of the ambiguity set for all three parameters and converges to zero as $k \searrow 0$.

- **Asymmetric signal response.** The 'Slope gap' satisfies:

$$\Delta\Big(\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i}\Big) := \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s>0} - \left. \frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} \right|_{s<0}.$$

  - $\kappa = \mu$ or $\kappa = \rho$. From the proof of Proposition 2.2.1, $\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} = K(\phi)$ everywhere, so $\Delta(\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i}) = 0$ for every k.

  - $\kappa = \phi$. We have $\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i} = K(\bar{\phi})$ for $s > 0$ and $K(\underline{\phi})$ for $s < 0$; thus:

$$\Delta\Big(\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i}\Big) = K(\bar{\phi}) - K(\underline{\phi}) = \frac{k}{(1+\bar{\phi})(1+\underline{\phi})}$$

    This is strictly increasing in k and satisfies $\lim_{k \searrow 0} \Delta(\frac{d\mathbb{E}_t^A[x_{t+1}|s_t^i]}{ds_t^i}) = 0$.

□

Lemma B.1

We state and prove a useful Lemma used in B.1.

Lemma B.1 (Moment of a piecewise normal)**.** *Let* $X \sim \mathcal{N}\left(0, \sigma_X^2\right)$*, and* $f_X(x)$ *be its associated probability density function. For constants* $c_+, c_- \in \mathbb{R}$*, define the piecewise function* $g(X) = c_+ \mathbf{1}_{\{X \geq 0\}} + c_- \mathbf{1}_{\{X<0\}}$*. Then:*

$$\mathbb{E}[g(X)\, X] = (c_+ - c_-)\, \frac{\sigma_X}{\sqrt{2\pi}}.$$

Proof. Because X has a symmetric density $f_X(x) = f_X(-x)$,

$$\begin{aligned} \mathbb{E}[g(X)\, X] &= \int_{-\infty}^{\infty} g(x)\, x\, f_X(x)\, dx \\ &= \int_{0}^{\infty} c_+\, x\, f_X(x)\, dx + \int_{-\infty}^{0} c_-\, x\, f_X(x)\, dx \\ . &= (c_+ - c_-) \int_{0}^{\infty} x\, f_X(x)\, dx \end{aligned}$$

Note that $\int_0^\infty x\, f_X(x)\, dx$ is the mean of a half-normal distribution with variance σ_X. Therefore $\int_0^\infty x\, f_X(x)\, dx = \frac{1}{\sqrt{2\pi}}\, \sigma_X$. Substituting back,

$$\mathbb{E}[g(X)\, X] = (c_+ - c_-) \frac{\sigma_X}{\sqrt{2\pi}},$$

□

B.2 ALP Data Appendix

The American Life Panel (ALP) is a survey administered by the RAND Corporation. The ALP consists of several thousand households who respond to a variety of surveys ("modules") on an ad-hoc basis. The majority of surveys are conducted online, but a small subset of households who lack internet access are given a laptop to answer, to limit selection bias. The survey is designed to be representative of the US adult population, and survey weights are provided to this end.

When survey respondents first join the ALP sample, they are questioned on various demographic characteristics. Using these questions, we construct our set of demographic controls. These are:

Age: Age and age squared are included (calculated by the ALP).

Educational Attainment: We encode the respondent's highest level of education into four categories: high school diploma or less, Associate degree or less, Bachelors degree or less, Postgraduate degree.

Race: Encodes whether a respondent identifies as: White, Black, American Indian/Alaskan Native, Asian or Pacific Islander, Hispanic/Latino or Other.

Gender: Takes the value 1 if male and 2 if female.

Marital Status: We encode the question over the respondents living situation into three categories: never married, previously married, married or living with a partner.

Labour Force Status: Encodes the respondent's answer to "current employment situation" as either: Working, Not Working or Retired.

Household Income: In Section 2.4 we control for income using the ALP's income measure. Household income is provided categorically, in 17 bins. We group these into three categories: $\$50k$, $\$50k - \$100k$ and $> \$100k$. These are the more coarse bins used in the regressions of Appendix A.3.

Our final demographic control comes from the Health and Retirement Study module, where respondents list their W2 census division of residence.

Ellsberg Module

We use three separate groups of modules throughout the paper. The first is the ambiguity aversion survey administered by Dimmock et al. (2013), which provides the basis for the experimental evidence on ambiguity aversion. From this module we use a number of variables. These are:

Ambiguity Preferences: First and foremost, we use the group of questions designed to elicit preferences over ambiguity. Here, respondents choose between lotteries with known probabilities and lotteries with unknown probability iteratively, to hone in on the known-lottery win probability that induces indifference. To construct the measure AA in the paper, we follow Dimmock et al. (2013), and form the same measure exactly.

Consistency Checks: To reduce measurement error of ambiguity preferences, Dimmock et al. (2013) include two consistency check questions. These test consumers preferences between urns at lottery values close to the indifferent point reached earlier in the survey. A consumer answered the check questions correctly if they chose the ambiguous box $\mathbb{U}$ after the probability for box $\mathbb{K}$ has been lowered, and box $\mathbb{K}$ when then probability is higher than their point of indifference. Our index is the number of consistent answers to these two questions.

Risk Aversion: Preferences over risk are also elicited by the authors, through a process similar to the questions eliciting ambiguity preferences. Respondents repeatedly pick between certain outcomes with particular payoffs, and lotteries with know probabilities. After a process of iteration, a point of indifference is found. From there, the authors back out the respondent's risk aversion. We follow the original paper exactly to construct the same measure.

Risk Question asked first: As in the original paper, we also include an indicator for each respondent which records whether or not they were asked the set of risk aversion questions before or after the set of ambiguity questions (the order of these questions was randomised across the sample). We include this variable as a proxy for survey fatigue.

Financial Literacy: Again, this is the same measure as in the parent paper, which itself follows the method implemented in the Health and Retirement Study. The survey includes three questions concerning basic financial literacy, of which the index is formed by taking the number of correct responses.

Agreeableness: Respondents were asked: *"Generally speaking, would you say that most people can be trusted, or that you can't be too careful in dealing with people? Please indicate on a score of 0 to 5."*. Our measure is the response to this question.

Question Clarity: Respondents were asked to grade whether or not they found the survey questions to be clear. We also include their response as a control for possible measurement errors[31].

Survey time: We calculate the difference between the survey start time and the survey end time. We generate two variables from this. The first measures whether the survey was completed in under 10 minutes, and the second whether it took longer than 24 hours to complete. We view both of these measures as proxies for survey engagement.

The first of these variables forms our dependent variable for Section 2.4, and the rest we employ as controls. In Table B.1 we display summary statistics for each of these variables.

Household Inflation Expectations Project

This set of modules of the ALP data was the precursor to the New York Federal Reserve's SCE, and was fielded after preliminary interviews with respondents in the MSC eliciting information regarding questioning alternatives. The data-set we use, is the third stage of the project. The sample population consists of MSC respondents who agreed to take part in further surveys after completing the Michigan questionnaire. The survey was launched in November 2007 and continued until May 2012. For more detail on the HIEP, please see Bruine de Bruin, Vanderklaauw, et al. (2010) or sections 1 or 2 of Armantier, Bruine de Bruin, Potter, et al. (2013).

The questions largely reflect the same format as the SCE. The primary exception to this is the point inflation forecast. Respondents are asked for a point forecast (with"rate of inflation wording"), a density forecast (as in the SCE), but are also asked the "prices in general" point forecast question (as in the MSC, for the sake of comparison). For both point forecasts, the elicitation process is more involved than in either of the MSC or the SCE. Here the questioning worked as follows:

- Respondents asked for their 12 month inflation expectation and can give a best guess *or* a range.
- If they only give a range, they are then prompted again for a best guess
- They are then asked to confirm their best guess of the rate of inflation if they responded with a best guess (or the range included) a value greater than 5%.
- If they then answer no to this question, they are given the opportunity to respond again with their best guess

We track the respondent's answers through this question tree, and take their final response to be their

[31] Dimmock et al. (2013) find that ambiguity aversion is significantly higher (by 5.4%) among the sub-sample of respondents who answered both check questions correctly.

inflation expectation. For those that do not respond with a best guess but rather a (valid) range, we take the midpoint of this range to be their expected rate of inflation.

For unemployment changes, house prices, gas prices, food prices and government debt growth, the elicitation process is the same as in the full SCE. We construct the forecast errors as in 1.4, trimming expectations at the 5-95% level by date of survey.

Health and Retirement Study

The data of this module originate from the RAND's Health and Retirement Study (HRS), although we specifically use the "User-Friendly data file of Financial Behaviour" derived from Version K of the HRS, produced by the Centre for Economic and Social Research at the University of Southern California[32]. This module provides the most detailed income and wealth data for ALP respondents. All respondents in the previous two sets of modules are surveyed within the HRS. We therefore employ this data in Section 2.5 and Appendix B.3.

From this data, we use the following:

Total Household Income (continuous): We deflate self-reported Total Household Income by CPI, to generate a continuous measure of real income. These data are trimmed at the 5-95% level. This is the variable used in Table 2.3.

Total household Income (categorical): Respondents also respond to a categorical question eliciting household income. These are grouped into 17, fine bins. We create a continuous measure by taking the midpoint of each bin. We use this variable as an alternative definition of income in case of the continuous measure of input error.

Wealth: We deflate total wealth (less IRA) by CPI. This is a net wealth measure. Please see the HRS codebook for the exact definition of total wealth. These data is trimmed at the 5-95% level.

Financial Wealth: We deflate total non-housing financial wealth by CPI. This is also a net measure. Please see the HRS codebook for the exact definition of non-housing financial wealth. These data are trimmed at the 5-95% level.

Housing Wealth: We deflate the Net Value of Primary Residence by CPI. These data are trimmed at the 5-95% level.

[32] See https://alpdata.rand.org/index.php?page=specialstudies or https://www.rand.org/well-being/social-and-behavioral-policy/portfolios/aging-longevity/dataprod/hrs-data.html for more on the User Friendly file and the HRS respectively.

Assets: We construct a gross asset position as the sum of the following: Businesses Owned, Stocks, Bonds, Checking/Savings Accounts, CDs/Bonds/T-bills, IRA, Primary Home, Secondary Home(s), Vehicles, Other. This is then deflated by CPI, and trimmed at the 5-95% level.

Total Household Income (continuous) is used in the main-body of the paper, and the others are used in robustness exercise of B.3. As mentioned earlier, we also use the W2 Census Region data included here as an added demographic control.

Summary Statistics

Summary statistics for the data used in this Section 2.4, Section 2.5 and Section 2.6 are shown in Table B.1.

Table B.1: ALP Data Summary Statistics

	Mean	Median	St. Dev	N
Ellsberg Module				
Ambiguity Aversion	15.9	12	13.6	5430
Financial Literacy	2.35	2	.674	5430
Consistency Check	3.12	4	1.01	5430
Risk Q. First	.558	1	.497	5430
Risk Aversion	1.4	1	.57	5429
Found Q.s Clear	3.59	4	.694	5430
Low Survey Time	.126	0	.332	5430
High Survey Time	.0163	0	.126	5430
Trust Index	2.92	3	1.32	5430
ALP				
Age	54.5	55	8.75	5425
Education	2.51	2	1.05	5419
Race	1.27	1	.945	5419
Gender	1.48	1	.5	5421
Marital Status	1.67	1	1.29	5430
LF Status	1.65	1	.844	5430
Income	13	14	3.04	5419
Region	5.12	5	2.51	5382
HRS				
log(Income)	11	11.1	.812	1474
log(Family Income)	10.8	10.9	.559	1223
$\sinh^{-1}$(Wealth)	3.47	4.63	3.09	1488
$\sinh^{-1}$(Financial Wealth)	1.05	.105	3.14	1533
$\sinh^{-1}$(Housing Wealth)	2.76	3.3	2.66	1523
log(Assets)	4.46	5.03	1.97	1423
HEIP				
Inflation (Dens) FE	1.04	.5	2.96	5033
Inflation (Point) FE	1.51	.915	3.3	5430
Unemployment FE	.0202	0	.397	5281
Gas Prices FE	.68	2.72	17.1	3477
Food Prices FE	1.83	1.22	4.18	3570
Gov. Debt FE	.086	-1.89	8.56	3962
House Prices FE	.803	.717	7.03	3947

Here we display summary statistics from the four different modules in the ALP data used throughout the paper. For the Ellsberg module, ALP base module and HEIP module, the data is the subset for respondents of which we have *AA* and inflation forecast error measurements (i.e. the sample used in the panel regressions). For HRS, this is the sample used in the longitudinal OLS regression. All summary statistics employ the ALP weights. For specific definitions of each variable, please see the previous sub-sections.

B.3 Ambiguity and Income in the Ellsberg module

In this section we provide additional results and robustness to Section 2.5. We first provide robustness to the result displayed in Table 2.3, and show that the insignificant coefficient on income is not particular to the HRS measure along. In figure B.1 we display the marginal effect of the ALP (categorical) income measure on ambiguity-aversion as elicited in the Ellsberg module. The results come from a regression that includes the set of demographic and survey controls. In the main body of the paper we purposefully avoid using the non ambiguity-averse respondents after an appeal to Occam's razor: we would posit it is hard to rule out the respondents misunderstanding the question (rather than truly displaying ambiguity neutral or loving preferences). However here, we show both the ambiguity-averse respondents, as well as all individuals surveyed in the Ellsberg survey. While there may be some small instances of particular pairwise bins of income display average ambiguity-preferences that are statistically different from one another, quantitatively there is little to no linear impact of income group on ambiguity-aversion.

Figure B.1: Marginal Effect of Income Groups on Ambiguity Preferences

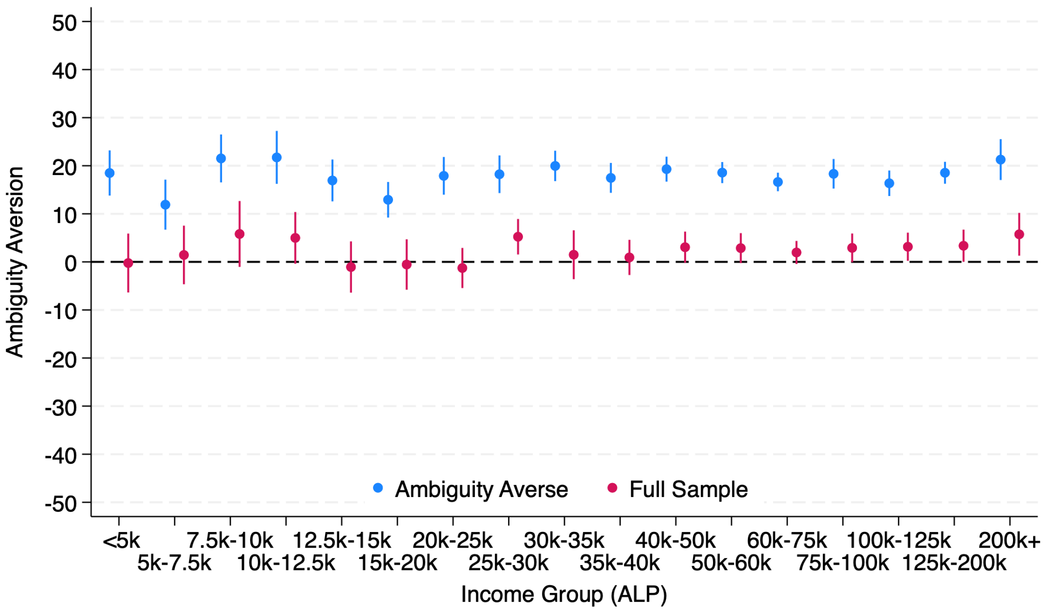

This graph plots the marginal effect of the (categorical) income measure from the ALP data on ambiguity-preferences, taken from a regression included the full set of demographic and survey controls and included time-fixed effects. The confidence intervals represent the 5-95% certainty range that the estimated marginal effect is 0. The standard errors are clustered at the individual level and (Ellsberg module) survey weights are employed. The blue points represent the regression specified over-ambiguity averse respondents, and the red points over all individuals who answered the Ellsberg module.

We now try to further rule out the possibility of the relationship between income and ambiguity preferences being (possibly complicatedly) non-linear. Figure B.2 performs a similar exercise as in Figure 2.1, and non-parametrically plots the relationship between (log) income and ambiguity preferences.

Figure B.2: Bin-scatter of Ambiguity Preferences by Income (Full Sample)

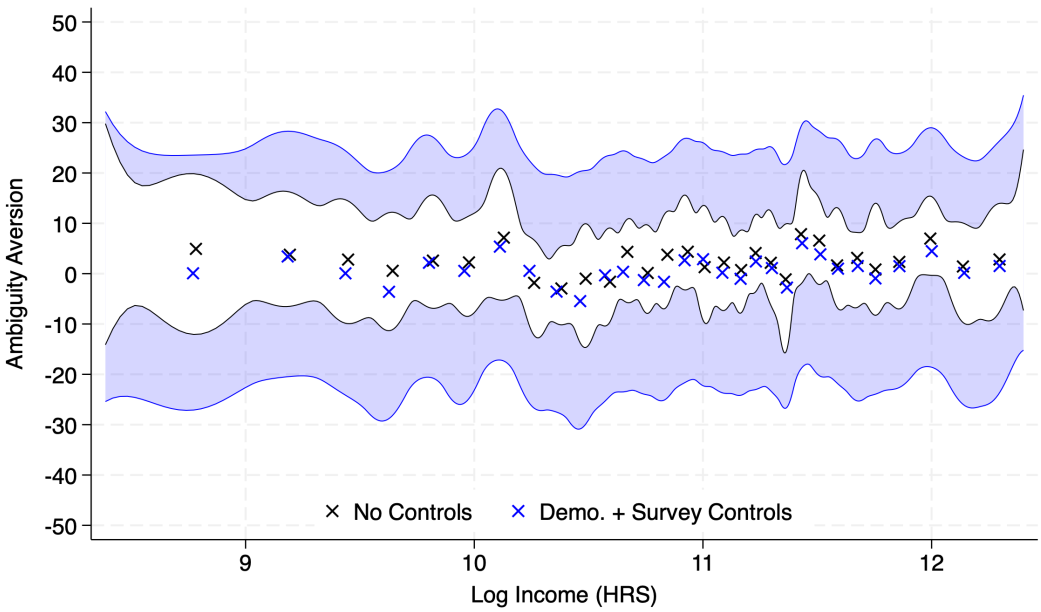

Here, we plot two binned scatters of (log) real income on our measure of ambiguity aversion. The difference between this and Figure 2.1 is that here we use the full sample of respondents to the Ellsberg module, and not just those who are ambiguity-averse. The income variable used is taken from the RAND-HRS module, denominated in 2015 dollars. The black crosses denote a specification where we control for time fixed effects only. The blue crosses include the same bin-scatter, but including the full set of demographic and survey controls. Each specification has been put into 30 bins. The white shaded area denotes a non-parametrically estimated 95% confidence interval for the specification with only time fixed effects, and the blue band a 95% confidence interval for the specification including controls. These confidence bands are calculated following Cattaneo et al. (2024) using standard errors clustered at the individual level in the first stage regression.

The bin-scatter reassures us that there is nothing particular about the sub-set of those who completed the Ellsberg module when it comes to income. There appears to be little relationship between ambiguity preferences *at all* and income in this survey, whether we control for demographic and other survey controls or not.

Finally, we turn to several other measures of income and wealth available to us in the HRS data, to provide further robustness to the result that there is no significant relationship between income and ambiguity-aversion in the Ellsberg module. We use the same specification as in 2.5, but now use alternative measures of the economic status of a respondent. These alternative measures are a continuous measure of household income, a categorical measure of household income, total net wealth, total net financial wealth, total net housing wealth and gross assets. In figure B.3 we plot the marginal effect of these measures on ambiguity preferences, across different sub-samples of the data. In blue, the sample of ambiguity-averse respondents is shown; in green, the full set of Ellsberg module respondents is shown; in red, the sub-set of ambiguity averse individuals with net wealth greater than $500; in yellow, the sub-set of all Ellsberg module respondents with net wealth greater than $500. We use the full sample of Ellsberg respondents to illustrate that is nothing particular about ambiguity-averse consumers, and we include only those with positive wealth to alleviate concerns that the lack of correlation is not driven by measurement error within the measure of wealth used.

Figure B.3: Marginal Effects of Alternative Income/Wealth Measures on Ambiguity Aversion

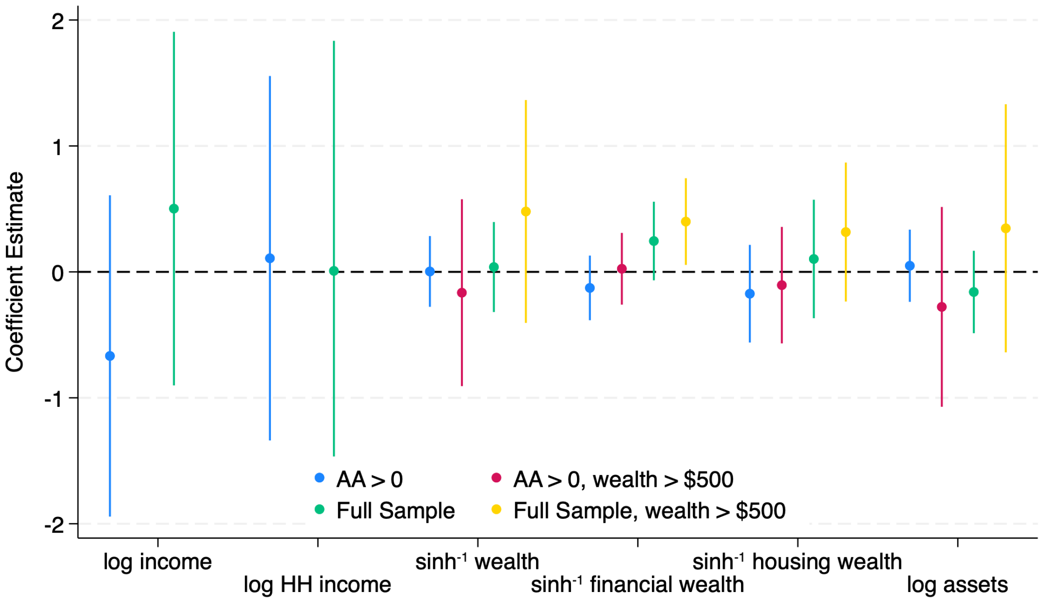

Here we plot the coefficients on the x-axis variables estimated in a regression of ambiguity-preferences, controlling for demographic and survey-controls, with time fixed effects. Column 1 uses the continuous measure of (log) income. Column 2 uses a (fine) categorical recording of income, where the (log) mid-point of each range was taken to be the respondent's income. Columns 3-6 instead employ measures of a respondents wealth. Column 3 uses net total wealth (less IRA), Column 4 uses net financial wealth, Column 5 uses net housing wealth and Column 6 uses (log) gross assets. All of these variables are taken from the HRS data. For Columns 3-5 an inverse hyperbolic sine transformation was used rather than a log-transformation, as net wealth positions can (and do) take negative values. All of these variables are deflated by CPI, and trimmed at 5-95% level. These regressions are ran across different sub-sets of the data, here denoted by the color of the estimate. These are: in blue, the sample of ambiguity-averse respondents is shown; in green, the full set of Ellsberg module respondents is shown; in red, the sub-set of ambiguity averse individuals with net wealth greater than $500; in yellow, the sub-set of all Ellsberg module respondents with net wealth greater than $500. The confidence interval represents the 5-95% interval, and are calculated using standard errors clustered at the individual level. Sample weights from the Ellsberg module are employed.

Across Figure B.3 almost all coefficient estimates are statistically indistinguishable from 0 (with the exception of net financial wealth, within the sub-set of all respondents to the Ellsberg module with net wealth greater than $500). However, quantitatively, all of these estimates are almost negligible. Taking the largest estimate in absolute magnitude (the very first point shown), a coefficient of -0.7 implies that for a 10% increase in income, ambiguity-aversion drops by 0.07 (roughly 0.0055% of the standard deviation of AA)[33]. Even if our estimates are imprecise (which they likely are given our numerous controls and

[33] For the inverse hyperbolic sine transformation the interpretation of the coefficient is not quite as clean. At small values the transformation is roughly linear, but at large values it is roughly log. Here, our measures of wealth are denominated in 1,000s of dollars, so to take the significant estimate of 0.3 as an example: for a individual with $5,000 of net wealth, the effect of a 10% increase in wealth on ambiguity preferences would be an increase in ambiguity-aversion of $\frac{5500-5000}{1000} \times 0.3 = 0.15$. For an individual with $500,000 of net wealth, the effect of a 10% increase would be $\left(\frac{550000}{500000} - 1\right) \times 0.3 = 0.03$. Both of these effects are vanishingly small.

relatively small sample size), the point estimates for the marginal effect of various measures of income or wealth on ambiguity-aversion are negligible.

B.4 Survey Expectation Asymmetry

This section provides additional detail to Section 2.6 in the main body of the paper. We first describe the SPF data used in this section.

Public Signals Proxy

To construct our proxy for public signals (signal_t) we take the median professional forecast for each quarter, and remove the current value of the forecast variable. This is as in Baqaee (2020). Therefore our proxy if positive when professional forecasts expect an increase in, say, inflation over the coming 12-months, and negative when they expect a decrease. For one-year ahead inflation and the national unemployment rate, the choice of forecast variable is self-evident. For interest rates, we choose the 3-month T-Bill forecast as our proxy for secular interest-rate movements. The outcome variables we use for inflation, unemployment and interest-rate changes are headline CPI inflation, the BLS national unemployment rate and the 3-month T-Bill rate. In Figure B.4 we plot our proxies for public signals flow across inflation, unemployment and interest rate changes: the blue line is constructed as $\max\{\text{signal}_t, 0\}$ and the red line as $\min\{\text{signal}_t, 0\}$. We use data from 1985-2019 inclusive, to maximise the time-dimension of our data but to avoid large shocks at the beginning of the 1980s, the covid-period, and 2021-2022 inflation episode overly impacting our coefficient estimates.

Figure B.4: Signals Proxy Measures Over Time

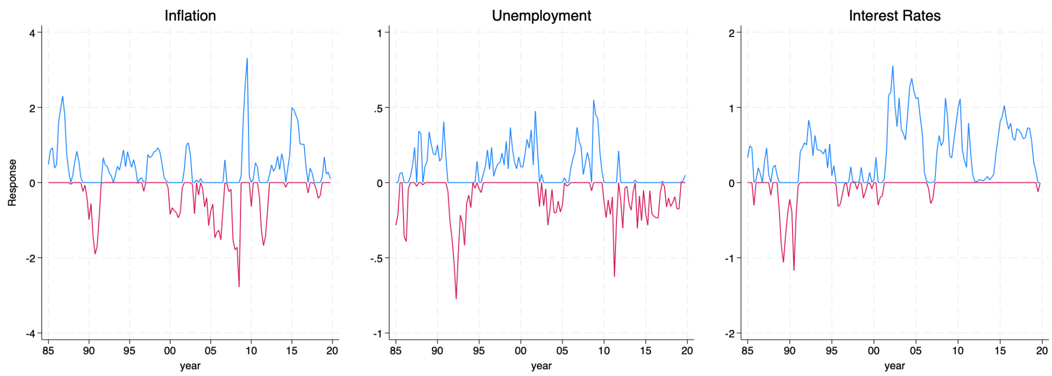

Here, we show the proxies derived from the SPF for public signals across one year ahead inflation, national unemployment and the 3-month treasury bill interest rate. The data is quarterly. Each signal variable is constructed as the median SPF forecast in a given quarter, minus the current value of the variable in question. We have split the signal variable into two parts: $\text{signal}_t^+ = \max\{\text{signal}_t, 0\}$, shown in blue, and $\text{signal}_t^- = \min\{\text{signal}_t, 0\}$, shown in red. Here, a value of 3.3 in 2009Q3 implies that the median SPF forecast for inflation in 2010Q3 was 3.3% above the 2009Q3 value.

We may be concerned that when estimating the marginal effect of public signals on expectations, our coefficients will be affected by the distribution of the signals themselves. In particular, if the signal proxy we use is skewed towards positive (or negative) values. In Table B.2 we report tests of skewness and kurtosis of signal_t for each of inflation, unemployment and interest rate changes. We fail to reject the null hypothesis of zero skew in any of the measures (although our measure of inflation signals appears to exhibit fat-tails, but this is less of a concern).

Table B.2: Skewness and Kurtosis Test Results For Signals Proxies

	Obs.	p(skew)	p(kurt)	p(joint)
Inflation News	140	0.921	0.029	0.089
Unemployment News	140	0.141	0.240	0.165
Interest Rate News	140	0.670	0.336	0.570

This table displays the results of a finite-sample corrected test for skewness across each of the proxies for inflation, unemployment and interest rate public signals. Columns 1 and 2 report p-values for tests of skewness and kurtosis. Column 3 reports the p-value for the joint test. The null hypotheses are skew = 0, kurtosis = 3 and normality respectively. See D'Agostino and Belanger (1990) for more.

Asymmetry Regressions

In Figure 2.2 we display a bin-scatter of our proxy for signals against the three main outcome variables. We test for this formally in this section, using a regression. For outcome variable (x_t) - inflation, unemployment and interest rates - we perform the following regression:

$$\mathbb{E}_t^i[x_{t+12}] = \boldsymbol{X}_t^{\text{past}}\,\boldsymbol{\beta} + \beta_+ \cdot \text{signal}_t^+ + \beta_- \cdot \text{signal}_t^- + \delta_i + \varepsilon_{i,t}$$

Where $\boldsymbol{X}_t^{\text{past}}$ is a vector of time-varying controls, which control for the current aggregate state of information, $\text{signal}_t^+ = \max\{\text{signal}_t, 0\}$ is our proxy for signals of a coming increase in y_t, and $\text{signal}_t^- = \min\{\text{signal}_t, 0\}$ for a coming decrease. As our coefficients of interest β_+ and β_- use time-variation, we can include individual fixed effects (δ_i) in our regression - something we did not have the luxury of doing in Section 1.4.

For our independent variables (x_t) - inflation, unemployment and interest rate expectations[34] - we use the same variables as we have throughout the paper to construct forecasting errors. For 12 month ahead inflation changes, this is simple enough. For unemployment and interest rate changes, consumers provide a categorical response on whether or not they expect an increase, decrease, or for the variable to remain unchanged. We therefore code each of these variables as taking the value 1, -1 and 0 for each of these possibilities. As the SPF data is quarterly, we assign an individual as making their expectation

[34] As discussed earlier, the MSC questions respondents about their expectations over *borrowing* rates specifically. The SPF does not specifically produce forecasts for consumer borrowing rates. By using the forecast of the 3-month T-Bill are therefore implicitly assuming that the relationship between government borrowing rates and consumer borrowing rates does not change drastically over time.

in quarter q if the survey took place within that quarter[35].

In Aggregate. Figure B.3 displays the results of this regression. In columns 1, 3, and 5, our vector $\boldsymbol{X}_t^{\text{past}}$ controlling for the current information set contains the current value of x_t. In columns 2, 4 and 6 we expand this set of controls, and also include the previous quarter's value (x_{t-1}) as we as the lagged value of the SPF forecast $\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$[36]

Table B.3: Regression Output of Expectations on Both Positive and Negative Signals

	Inflation		Unemployment		Interest Rates	
signal_t^+	0.809***	0.873***	0.249***	0.228***	0.269***	0.300***
	(8.56)	(8.76)	(12.03)	(8.16)	(15.24)	(16.51)
signal_t^-	0.337***	0.475***	0.153***	0.196***	0.0110	-0.0152
	(3.44)	(4.42)	(5.90)	(5.76)	(0.36)	(-0.50)
x_t	0.818***	1.065***	-0.0109	0.151***	0.141***	0.536***
	(9.88)	(9.91)	(-1.08)	(2.66)	(17.21)	(16.49)
x_{t-1}		-0.141**		-0.203***		-0.334***
		(-2.15)		(-3.24)		(-8.93)
$\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$		-0.297***		0.0433		-0.0977***
		(-3.17)		(1.54)		(-6.87)
$\text{signal}_t^+ - \text{signal}_t^-$	0.472***	0.397***	0.095**	0.032	0.258***	0.315***
	(6.78)	(5.64)	(2.57)	(0.77)	(7.14)	(8.73)
Obs	172,118	172,118	207,847	207,847	206,149	206,149
Adj. R^2	0.403	0.403	0.370	0.370	0.299	0.308
Mean dep var	3.1	3.1	.154	.154	.414	.414
First obs	1985q1	1985q1	1985q1	1985q1	1985q1	1985q1
Individual FE	✓	✓	✓	✓	✓	✓

This table displays regression results for the regression specification above, conducted for one-year ahead inflation, unemployment and interest rates expectations. In Columns 1, 3 and 5, we control for the current value of x_t. In Columns 2, 4 and 6 we expand this set of controls, and also include the previous quarter's value (x_{t-1}) as we as the lagged value of the SPF forecast $\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$. All specifications include individual fixed effects. Standard errors are clustered at the individual level, and MSC survey weights are employed. We also include a row that reports the difference between β^+ and β^-, and the t-statistics for the hypothesis test that $\beta^+ = \beta^-$ are reported. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

The regression display similar results to Figure 2.2. For all specifications, the responses of expectations to positive public signals is greater than the response to negative signals: furthermore, all coefficients except negative signals in the simple specification for unemployment, these estimates are significantly different from zero. Using additional controls in the vector $\boldsymbol{X}_t^{\text{past}}$ - our control for the current set of information - affects the quantitative estimates but not the qualitative result. The first row after the regression estimates tests the hypothesis that $\beta^+ = \beta^-$: these differences are positive and significant at

[35] Aggregating like this leaves open the possibilities of within-period timing becoming a concern - particularly in reference to the SPF forecast. Individuals who answered the survey in July (say), would not yet have seen the professional forecasts the SPF lists as Q3. However, we view our measure as a proxy for signals, rather than the signal itself. Furthermore, since we are considering expectations for one year in advance, it is unlikely to be important either way.

[36] This is the same set of controls for the current information set as used in Baqaee (2020).

the 5% level, with the exception of unemployment. In aggregate, consumers' responses to our proxy for public signals display asymmetry depending on the direction of the signal. However, the average response of consumer expectations as elicited by surveys is convex in signals of the future.

Across Income groups. Given the cross-sectional results discussed in 1.4, a natural question is whether this asymmetry varies by income group. If ambiguity over the information content of the signals *is* the underlying cause of consumer forecast errors and these decline in income, then it would be the case that this asymmetry also declines in income.

To test this hypothesis, we are interested in exploring how β_+ varies by income. We therefore let our marginal effects of a public signal shock depend on income, and decompose our coefficient as $\beta_+ \approx \beta_+ + \beta_+^k \cdot \text{income}_{it}^k$ for individual i belonging to income group k. Here we define income groups as income quintile by year, in 2015 dollars[37]. Doing so leads us to the following regression specification (reproduced from the main body of the text):

$$\mathbb{E}_t^i[x_{t+12}] = \sum_k \text{signal}_t^+ \cdot \left(\beta_+ + \beta_+^k \cdot \text{income}_{it}^k\right) + \text{signal}_t^- \cdot \left(\beta_- + \beta_-^k \cdot \text{income}_{it}^k\right) + \gamma_p \boldsymbol{X}_t^{\text{past}} + \gamma_c \boldsymbol{C}_{it} + \varepsilon_{it}$$

We use two different specifications. Our baseline specification is above, where we include our vector of controls for the current information set, $\boldsymbol{X}_t^{\text{past}}$ [38] and the set of demographic controls available in the MSC[39]. These are included as we are interested in the cross income-group effect, and so cannot employ individual fixed effects. Our second specification simply includes time-fixed effects only in $\boldsymbol{X}_t^{\text{past}}$. The baseline regression shows the coefficient estimates shown in Figure 2.3. The specification included time-fixed effects displays much the same pattern as the baseline specification.

---

[37] In previous sections we grouped individuals in the MSC according to the scheme < \$50k, \$50k − \$100k and > \$100k. We used quintiles here for the sake of visualisation; the choice of income groups does not affect the results

[38] Containing the current value of the outcome variable x_t, the lagged value of the outcome variable x_{t-1}, and $\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$.

[39] see Appendix A.1 for more.

Table B.4: Regression of Expectation Asymmetry in Response to Public Signal, by Income

	Inflation		Unemployment		Rates	
signal_t^+	0.884***		0.382***		0.247***	
	(19.39)		(12.62)		(25.26)	
signal_t^-	-0.291***		-0.0205		-0.0748***	
	(-5.42)		(-0.74)		(-3.38)	
20-40pct x signal_t^+	-0.140***	-0.149***	-0.0212	-0.0310	0.00346	0.00979
	(-4.18)	(-4.49)	(-0.59)	(-0.88)	(0.34)	(0.97)
40-60pct x signal_t^+	-0.198***	-0.208***	0.0373	0.0253	0.0143	0.0111
	(-5.97)	(-6.35)	(1.02)	(0.70)	(1.37)	(1.07)
60-80pct x signal_t^+	-0.283***	-0.287***	0.00208	-0.00351	0.0255**	0.0333***
	(-8.60)	(-8.79)	(0.06)	(-0.10)	(2.40)	(3.20)
80-100pct x signal_t^+	-0.364***	-0.376***	0.0727*	0.0474	0.0587***	0.0669***
	(-10.81)	(-11.26)	(1.92)	(1.27)	(5.44)	(6.30)
20-40pct x signal_t^-	0.186***	0.203***	0.166***	0.172***	0.0825***	0.0823***
	(4.00)	(4.40)	(5.04)	(5.27)	(2.78)	(2.80)
40-60pct x signal_t^-	0.406***	0.422***	0.268***	0.243***	0.149***	0.149***
	(8.91)	(9.41)	(7.87)	(7.17)	(4.82)	(4.86)
60-80pct x signal_t^-	0.430***	0.456***	0.397***	0.380***	0.195***	0.193***
	(9.50)	(10.22)	(11.54)	(11.11)	(6.40)	(6.39)
80-100pct x signal_t^-	0.620***	0.632***	0.614***	0.609***	0.275***	0.270***
	(13.72)	(14.18)	(16.93)	(16.96)	(8.52)	(8.47)
x_t	0.760***		0.597***		0.734***	
	(17.97)		(28.80)		(56.78)	
x_{t-1}	-0.0268		-0.415***		-0.506***	
	(-0.88)		(-16.99)		(-32.41)	
$\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$	-0.486***		-0.170***		-0.237***	
	(-12.76)		(-19.04)		(-42.04)	
Obs	160,764	160,764	192,229	192,229	190,879	190,879
Adj. R^2	0.035	0.068	0.030	0.049	0.056	0.093
Mean dep var	3.1	3.1	.151	.151	.416	.416
First obs	1985q1	1985q1	1985q1	1985q1	1985q1	1985q1
Demographic Controls	✓	✓	✓	✓	✓	✓
Time FE		✓		✓		✓

This table displays the regression output of the asymmetric public signals by income-group specification. Columns 1, 3 and 5 display the baseline regression, where we control for current information set using the current value of the outcome variable x_t, the lagged value of the outcome variable x_{t-1} and the lagged SPF forecast $\mathbb{E}_{t-3}^{\text{SPF}}[x_{t+9}]$. Columns 2, 4, and 6 replace these with time fixed effects. Both specifications include demographic controls. The standard errors are clustered at the individual level, and survey weights employed. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

ALP data Expectation Asymmetry

Here we supplement Section 2.6.2 with additional detail. Using the ALP data, we regress inflation expectations as measured in the HIEP on the measure of inflation signal proxy discussed in Section 2.6 interacted with the measure of ambiguity aversion ambiguity-aversion taken from the Ellsberg module[40]

Our baseline specification employs the same set of controls for the current information set ($\boldsymbol{X}_t^{\text{past}}$) as before: current inflation, π_t, lagged inflation (π_{t-1}) and the lagged value of the SPF forecast $\mathbb{E}_{t-3}^{\text{SPF}}[\pi_{t+9}]$. We also control for the set of demographic controls and survey-based controls available in the combined ALP data. Finally, we also control for income - the reasoning is as in Section 2.5: we are interested in the effect of ambiguity aversion *separate* to the possible effects of income on expectation formation. For robustness, we include an alternative specification where we control for the current information set using time-fixed effects[41]. Table B.5 shows the regression output.

Table B.5: Response of Expectations to Positive and Negative Signals, by Ambiguity Aversion

	Baseline	Time FEs
signal_t^+	1.232***	
	(3.15)	
signal_t^-	0.729	
	(1.58)	
AA x signal_t^+	0.0122**	0.0120**
	(2.29)	(2.23)
AA x signal_t^-	-0.0186**	-0.0185**
	(-2.36)	(-2.36)
$(\beta_+ + \beta_+^{AA}) - (\beta_- + \beta_-^{AA})$	0.534***	N/A
	(2.79)	
Obs	3,622	3,621
Adj. R^2	0.175	0.179
Mean dep var	3.49	3.49
Survey Controls	✓	✓
Dem. Controls	✓	✓
Time FEs		✓

This table displays the regression output behind Figure 2.4. In our baseline specification, we include our control for the current information set ($\boldsymbol{X}_t^{\text{past}}$), which contains current inflation, π_t, lagged inflation (π_{t-1}) and the lagged value of the SPF forecast $\mathbb{E}_{t-3}^{\text{SPF}}[\pi_{t+9}]$. We also include the full set of demographic and survey controls, as detailed in Appendix B.4. The second column replaces ($\boldsymbol{X}_t^{\text{past}}$) with time-fixed effects, leaving only the interaction terms identified. The row below the coefficients shows the hypothesis test that $(\beta_+ + \beta_+^{AA}) = (\beta_- + \beta_-^{AA})$. All standard errors are clustered at the individual level, and Ellberg module survey weights are employed.

[40] As the ALP data is over a short time period, there is insufficient variation in the unemployment and interest rate public signals to perform similar exercises.

[41] In theory, it would be possible to estimate this regression by doing away with the demographic and survey controls and employing individual fixed effects as well. However, there is a very short panel (only half of individuals are exposed to more than one value of the signal proxy), and relatively few observations. It is therefore infeasible in practice.

As discussed in the main body of the paper, the results for ambiguity averse individuals mimic the results we found for income. Individuals respond positively to public signals. In this sample, the marginal effect of negative signals about inflation is statistically insignificant from zero. As shown in the table and shown in Figure 2.4, the effect of ambiguity on the degree of asymmetry is significant at the level 1%.

Chapter 3

Pessimistic Inflation Expectations, Portfolio Choice and Welfare

3.1 Introduction

As measured by surveys of expectations, households expect more inflation than materialises. Moreover, this varies across households: poorer households typically have larger forecast errors than richer counterparts[1] (Angelico and Giacomo (2024), Chapter 1 of this thesis). To what extent do these misperceptions matter for a household's portfolio decisions? Moreover, if poorer households typically have larger forecast errors, to what extent does this drive differences in investment behaviour? Finally, as these effects accumulate over a households lifetime, how much do these misperceptions matter for welfare? This paper answers these questions. We first illustrate using household survey data that inflation forecast errors are indeed correlated with investment decisions. We find that inflation forecast errors monotonically decline as stock market investments increase, among a sample of U.S. households. We then write a model of household portfolio choice which can capture the empirical correlation between inflation forecast errors and investment decisions, and show that forecast errors of the magnitude displayed in consumer surveys can have significant effects on portfolio choice when agents face nominal risk. In the cross-section, the effect of persistent forecast errors on deviations from the rational, expected utility maximising benchmark, is largest for the poorest households. When we then turn to the welfare implications of these forecast errors, we find them to be large: even seemingly small misperceptions of

[1] This is true not just for the extremes. See Chapter 1 for evidence that this relationship is monotonic.

the real rate of return can generate large welfare costs, as the effect of missed excess returns accumulates quickly.

Inflation expectations affect a household's dynamic decisions: whether to consume today versus tomorrow, whether to invest for the future, and if so, by what means. In Section 3.2, we use the Michigan Survey of Consumers and detail a novel fact: the degree to which households over-expect inflation declines as stock market holdings increases. This result is true controlling for income.

In order to explain this fact, we turn to a model of portfolio choice with nominal risk. This is presented in Section 3.3. In order to allow agents within the model to have expectations of the future that differ from the rational benchmark, we employ the robustness framework of Hansen and Sargent (2001). We distinguish ourselves from the rest of the robust portfolio choice literature in two key respects. First, rather than employ the homothetic transformation introduced by Maenhout (2004), and subsequently much of the work following, we instead use the non-homothetic specification. This is central to being able to match the cross-sectional profile of inflation expectations: when robustness is homothetic, the worst-case distortion is independent of the agent's state. Secondly, we specifically consider robust portfolio choice within an environment in which there is nominal risk. This allows us to take a novel approach to quantifying the model: we choose parameters to match the average inflation expectation of US households in consumer survey data, rather than calibrating parameters according to a test of statistical similarity.

In our baseline model, agents can invest in either a nominal bond or a risky-asset with potentially non-zero correlation with the inflation process. Nonetheless, as we present a model of robust portfolio choice, the main mechanism present in Maenhout (2004) is still present: the investor underinvests in risky assets as they perceive the returns of such assets to be lower. However the two novelties outlined above generate different mechanisms within the model.

First, motivated by survey data, we use a non-homothetic specification, which ensures that poorer households act as if they were more pessimistic than richer households. These households are more pessimistic due to their high marginal utility: the costs of their mental model being wrong is higher, and so they act as if in a state of the world that is worse. They expect both higher inflation and lower asset returns. They therefore choose optimal portfolio shares lower than the expected utility benchmark, and lower than their higher income or higher wealth counterparts.

Secondly, as poorer households have larger inflation forecasts than richer households in the model, even seemingly small differences in inflation forecast errors across groups can have large differences for their investment behaviour. This is due to the covariance structure of asset-returns and the inflation process. When the correlation between risky asset returns and inflation is negative - the empirically relevant case - this leads poorer households to under-invest in stocks above and beyond that which concern for misspecification of the returns process would generate alone. These two effects therefore reinforce one another in our calibrated model, and amplify the effects of small differences in beliefs. In Section 3.4, we calibrate the model to match the aggregate inflation forecast error and aggregate risky-asset portfolio share in the US, and quantify the magnitude of this mechanism. We find that the differences between the expected-utility benchmark and the model with empirically consistent expectations are large, and aggregate investment behaviour changes dramatically: among the poorest households, the maximum decline in the optimal risky-asset share is 90%.

Given these effects, in Section 3.5 we quantify just how important these mechanisms are. We calculate consumption equivalent measures of welfare and find the costs of robustness in this environment to be substantial. As even moderately pessimistic expectations have an outsized effect on portfolio allocation, over the course of a life-time the unrealised excess returns amounts to a large effect on wealth accumulation, and therefore consumption. We employ the tools of Achdou et al. (2022) and solve the model quantitatively using a finite difference scheme. This has the particular benefit of allowing us to find the ergodic distribution of asset holdings within the model. Using this, we calculate that within the calibrated model the potential welfare effects of an expected utility maximising agent being endowed with the pessimism of the robust agents to be large. We also find that the welfare effects of such pessimism, perhaps unsurprisingly, falls most on the poorest consumers with the largest forecast errors.

Finally, in Section 3.6 we investigate an extension of the model, in which portfolio adjustment is subject to a fixed-transaction cost. Fixed costs deliver the a simple micro-foundation for the lumpy, infrequent trading that is ubiquitous in household data (Vissing-Jørgensen (2002); Calvet et al. (2009)). Including a fixed portfolio-adjustment cost therefore allows us to study the interaction between state dependent pessimistic expectations and the investment inertia of low-wealth investors. We investigate the effects of including robustness in such a model. We show that robustness narrows the inaction region of agents in the model, as a greater concern for robustness implies the expected utility loss of being away from the frictionless portfolio rule is greater. This implies that an economy in which households have an average inflation forecast error of 2% rather than 1% should see a greater number of portfolio adjustments. We perform a welfare exercise here too: we show that an increasing concern for robustness among agents increases the welfare loss generated by the fixed cost. That is to say, policies that reduce transaction frictions have higher values in economies in which agents distrust the model more (i.e., they have higher forecast errors).

Literature Our empirical motivation builds on the literature investigating cross-sectional consumer expectation surveys, of which an early contribution was Mankiw, Reis, and Wolfers (2003). Papers such as Coibion, Georgarakos, et al. (2023) and D'Acunto, Charalambakis, et al. (2024) detail that US households typically report inflation expectations that are greater than realised. Other papers, such as Malmendier and Nagel (2016), D'Acunto, Hoang, et al. (2023) and D'Acunto, Malmendier, et al. (2021) explore how inflation expectations as measured in surveys varies across demographic characteristics - cohort, IQ and gender respectively[2]. Other papers such as Angelico and Giacomo (2024), and Chapter 1 of this thesis explore how economic status (specifically income) is a strong predictor of household inflation expectation errors. Masolo and Monti (2024) find that savers have higher inflation expectations than borrowers, using the NY Fed's Survey of Consumer Expectations. Closest in spirit to this work is Nord (2022), which links household wealth data to expectation surveys, and illustrates that across Danish households wealthier households display more accurate inflation forecasts (i.e. lower mean forecast errors). We contribute to the stock of empirical facts regarding household inflation expectations with a new fact: US households with larger investments in the stock-market have lower, but positive, inflation forecast errors, relative to their low investment counterparts. The findings of all these papers are given greater credence by the results presented in Armantier, Bruine de Bruin, Topa, et al. (2015) Bachmann et al. (2015) and Crump et al. (2022), which find that surveys of inflation expectations are indeed informative for predicting household's dynamic consumption choices.

We write a model of continuous-time portfolio allocation, following a long literature beginning with Merton (1971). For a recent survey of this literature, see Gomes (2020). We focus on the decision of investors in the presence of nominal risk, first explored by Fama and Schwert (1977). Subsequent seminal contributions include Campbell and Viceira (2001) and Brennan and Xia (2002). The latter of these papers examines an environment akin to our own, wherein agents have access solely to nominal assets and are precluded from short-selling. They instead focus on the maturity structure of the optimal portfolio in a model without concern for misspecification. More recently, Li (2019) study the welfare effects of bond-indexation across varying structures in a continuous time portfolio problem. Zhang et al. (2019) provide analytic solutions to a life-cycle portfolio problem where agents have hyperbolic absolute risk aversion facing nominal risk. Relative to this literature, we focus on the implications of cross-sectional heterogeneity in the agent's beliefs on portfolio choice.

In order to model this cross-sectional heterogeneity in beliefs - i.e. the positive forecasting errors

[2] Also see Cocco et al. (2019), Weber et al. (2022) or Fofana et al. (2024) for reviews.

that respondents in consumer surveys of expectations display - we employ the influential robustness framework of Hansen and Sargent (2001), in which agents seek to make optimal decisions in the face of possible model misspecification. This is the first of many papers including fear of model misspecification; Hansen and Sargent (2010b) provides an overview. A stated initial goal of this approach was to provide a general framework for formalising the empirical findings of Ellsberg (1961), who details that consumers display aversion towards lotteries of unknown probabilities. Other than the robustness framework, a related but distinct way to formalise such behaviour begins with Gilboa and Schmeidler (1989), in which agents face uncertainty towards the parameters of the data-generating process, termed "ambiguity". This "max-min" approach of this contribution was then generalised by Maccheroni et al. (2006), who allow for "smooth" preferences towards ambiguity, and Klibanoff et al. (2009) who provide a dynamic version of these preferences. Both the robustness approach and the ambiguity approach have received extensive investigation within the context of macroeconomics: Cagetti et al. (2002), Hansen and Sargent (2010a) and Bhandari et al. (2024) all use robustness within the context of macroeconomics. The last of these papers is conceptually aligned to this paper: the authors calibrate a representative agent business cycle model to match aggregate household inflation and unemployment expectations. Within the ambiguity tradition, Ju and Miao (2012), Ilut and Schneider (2014), Ilut and Schneider (2023) consider the effects of worst-case thinking on business-cycle dynamics. We supplement such work by showing that the potential welfare effects of such thinking - albeit in partial equilibrium - are significant.

Other papers have considered the effect of robustness or ambiguity on portfolio allocation. Early contributions employing the ambiguity approach include Chen and Epstein (2002), Garlappi et al. (2007) and Epstein and Schneider (2008). Within the strand of research employing robustness, Uppal and Wang (2003), Maenhout (2004) and Maenhout (2006) are seminal papers. The last two of these most closely resemble the model presented in this paper, although they employ a homothetic transformation to the original Hansen and Sargent (2001) framework to solve the investor's problem analytically. By defining the concern for robustness (in the right way) to be a function of the state-variables itself, one can eliminate terms depending on the state from the optimal worst-case drift of the agent. Such an approach implies that the robust investor acts in the same manner as an investor with a higher degree of risk-aversion. However this approach is inappropriate in our context - the empirical profile of inflation forecast errors is declining in stock market holdings, which is explicitly ruled out when employing the homothetic robustness transformation of Maenhout (2004). For more recent applications of robust portfolio choice within the context of asset pricing, see Drechsler (2013), Brenner and Izhakian (2018) or Aït-Sahalia et al. (2021). To our knowledge, only Yang (2020) solve a robust portfolio problem in the context of nominal risk. This contribution focuses on providing an analytic solution to the problem using a homothetic robustness specification. Relative to this literature, we make two contributions: First, we consider the (dynamic) welfare effect of robust portfolio-choice in the cross-section. Secondly, we take a different approach to the literature to calibrate robustness preferences.

In our model, it is important that households cannot hedge the risk of inflation: if they could do so, they would have no reason to treat high inflation states as a worst-case. The work of Doepke and Schneider (2006), and later, Auclert (2019) measures the inflation exposure of household balance sheets (i.e. their Net Nominal Position (NNP)). They find that households balance sheets are not inflation hedged: they find that in 1989, poor households gain substantially (roughly 45% of average wealth). On the other-hand, the losses accrued by richer, older households, are in the range of 6-12% of GDP in present value terms. Bahaj et al. (2023) study the market for inflation swaps in the UK with transaction level data. They find the market to be segmented, with pension funds primarily holding positive long maturity positions. This is suggestive that household pensions are likely to be inflation-hedged at long-horizons. Asset managers - with whom a retail trader could hold an account - comprise a small proportion of the total market for inflation swaps. Together, these results suggest that very few households hedge inflation risk at short horizons, and only household pensions are potentially hedged, and even then only at long horizons. As the model presented in Section 3.3 is one of active investment, with short-term nominal risk, we view these results to provide empirical justification for our assumption that households cannot hedge inflation risk.

Reis (2021b) finds that households surveys consistently express large discrepancies in inflation expectations relative to market-based measured. Therefore for our calibration to accurately represent household expectations, it is important we match consumer surveys of expectations, rather than other measures of inflation expectations. In fact, using robustness to match these surveys is central to the calibration of our model[3]. However, this represents a departure from the literature so far. In Anderson et al. (2003), the authors introduce "detection-error probabilities" as a way to calibrate the preference for robustness of agents in the model. Here, the researcher calibrates the robustness parameter such that the agent's benchmark model and the worst-case model are difficult to empirically distinguish - typically error-detection probabilities in the range of 10-20% are employed. This would be consistent with agents in the model fearing model misspecifications that the econometrician would only detect 10% of the time. Later papers almost always follow this method, such as Maenhout (2006) and Hansen and Sargent (2010a). An exception to this is Ilut and Schneider (2014), which instead use the measured volatility of unknown stochastic process to provide bounds on the possible degree of ambiguity agents face.

[3]A secondary result of our model is demonstrating that the non-homothetic robustness framework matches the cross-sectional profile of expectations remarkably well, even when the cross-sectional moments are not targeted.

Many authors have considered the welfare costs of inflation. Notable early contributions include Cooley and Hansen (1998) and Lucas (2000). Papers in the neoclassical tradition typically find the welfare costs of inflation to be negligible. In business-cycle models with nominal rigidities, inflation (more specifically goods price dispersion) typically generates more meaningful costs of inflation of the order of 1-3% (see Cavallo et al. (2023) for a recent discussion). Other papers, such as Doepke and Schneider (2006), Wolff (2023), Del Canto et al. (2023) and Ferreira et al. (2024), consider the cross-sectional incidence of inflation. These papers typically find that poorer and younger households fare better after unexpected inflationary shocks, as they typically hold nominal debt, as opposed to richer, older households who typically hold nominal assets (cash). The last of these papers finds that the 2021 inflation episode reduced savings capacity of Spanish residents older than 65 by roughly 7% on average. We contribute to this literature in the following respect: rather than considering the welfare effects of inflation per se, we consider instead the welfare effects of the distribution of persistent, positive, inflation forecast errors. Montag (2024) is the only other paper we know of that is interested in a similar question. Here the authors find a relatively small (of the order 0.2%) effect of divergent beliefs over inflation in France. However their paper lacks the risky-asset investment choice and asset accumulation that our paper does, which serves to (strongly) amplify welfare costs: persistent over-prediction of inflation can generate welfare losses an order of magnitude larger than the direct effect of inflation as traditionally measured, as compounded missed excess-returns dramatically reduces wealth accumulation.

Outline Section 3.2 presents our empirical motivation, illustrating that the biggest investors in stocks have the lowest inflation forecast errors in the Michigan Survey of Consumers. Section 3.3 outlines our baseline model, and discusses the properties of inflation forecasts within the model. Section 3.4 details the quantitative results of the model and Section 3.5 explores the welfare costs of robustness in the baseline model. Section 3.6 extends the baseline model to also include a fixed-cost of portfolio adjustment. Section 3.7 concludes.

3.2 Inflation Forecast Errors Decline with Stock Holdings

As discussed, a common finding in the literature is that surveys of consumer expectations, respondents typically display positive forecast errors. However, relatively little work has been conducted on the relationship between wealth and inflation expectations[4]. While reliable data linking household wealth to economic expectations is absent for the United States, we use the Michigan Survey of Consumers (MSC) to illustrate that a similar relationship between a households stock-holdings and inflation expectations

[4]The key exception is Nord (2022), highlighted earlier. He finds that wealthier households typically display lower mean absolute errors than less wealthy households in the Netherlands data.

holds in the US. We then use this to motivate our model of Section 3.3.

We employ the Michigan Survey of Consumers (MSC), the most commonly referenced and longest running survey of consumer expectations in the United States[5]. The information within the MSC on the asset and liability position of its respondents is limited. Nonetheless it does contain some information: namely, since 1990 respondents have been asked whether or not a member of the household has investments in the stock market, including within retirement accounts[6]. In line with previous findings of consumer participation in the stock market[7], 61% of respondents report having at least some wealth invested in stocks. Secondly, if they respond to this first question with "yes", they are asked to estimate the market value of their total holdings.

Naturally, data such as these are likely to contain some measurement error. For example, respondents may be unaware of the asset class allocation of managed retirement accounts they own or may not have a good grasp of the current market value of their holdings. Given these possibilities, we attempt to minimise the impact of such concerns by grouping households based on their stock holdings. We group households into quintiles of self-reported holdings of stocks within a given survey wave.

We are interested in how consumer inflation forecasting errors vary across households by stock-holdings. We therefore first construct the households inflation forecast error, in the same manner as in Chapter 1. We define $FE_{t+12|t,i} \equiv \mathbb{E}_t^i[\pi_{t+12}] - \pi_{t+12}$, where $\mathbb{E}_t^i[\pi_{t+12}]$ is the one-year ahead inflation expectation of respondent i, and π_{t+12} is the realised year-on-year percentage change of the Bureau of Labour Statistic's Consumer Price Index measure. To avoid the possibility of measurement error driving our results, we trim this measure at the 5-95% level, by survey wave.

We then regress our measure of inflation forecast error on consumer stock-holdings groups. We employ the following regression specification:

$$FE_{t+12|t,i} = \gamma_0 + \sum_g \beta_g \cdot \mathcal{I}\left(i \in g\right) + \gamma \boldsymbol{X}_{i,t} + \delta_t + \varepsilon_{i,t}$$

Where $\mathcal{I}\left(i \in g\right)$ is an indicator variable that takes the value 1 if household i is within stock-holdings group g. Our six stock holding groups are the five quintiles of stock-holdings and non-participants. We also include a vector of (possible) controls, $\boldsymbol{X}_{i,t}$. Finally, we include time-fixed effects δ_t so as to control for aggregate factors which may influence consumer expectations. Our coefficients of interest

[5] For more information regarding the dataset, please see Appendix A.1.

[6] Respondents are prompted specifically to consider possible stocks held within 401(K)s, IRAs or Keogh accounts.

[7] See Cocco et al. (2019) for a review.

are therefore the marginal effect of belonging to a particular group, β_g. We utilise the survey weights provided by the MSC.

In Figure 3.1, we plot the average forecast error per stock-holdings group. Specifically, we plot: $\widehat{FE}_g$ for g either a non-participant (NP), or stock-holdings quantile $q = \{1, 5\}$. Here, $\widehat{FE}_g$ is the predicted forecast error from the regression of a respondent within group g, calculated as:

$$\widehat{FE}_g = \frac{1}{\sum_{j \in g} w_j} \sum_{j \in g} w_j \cdot \widehat{FE}_j$$

where w_j denotes the survey weight of observation j [8]. In words: Figure 3.1 shows the average predicted inflation forecast error within the sample of stock-holdings group g, measured against the predicted inflation forecast error for the base group of non-participants. For the underlying estimates, see Appendix C.1.

Figure 3.1: Inflation Forecast Errors by Stock Holdings, MSC

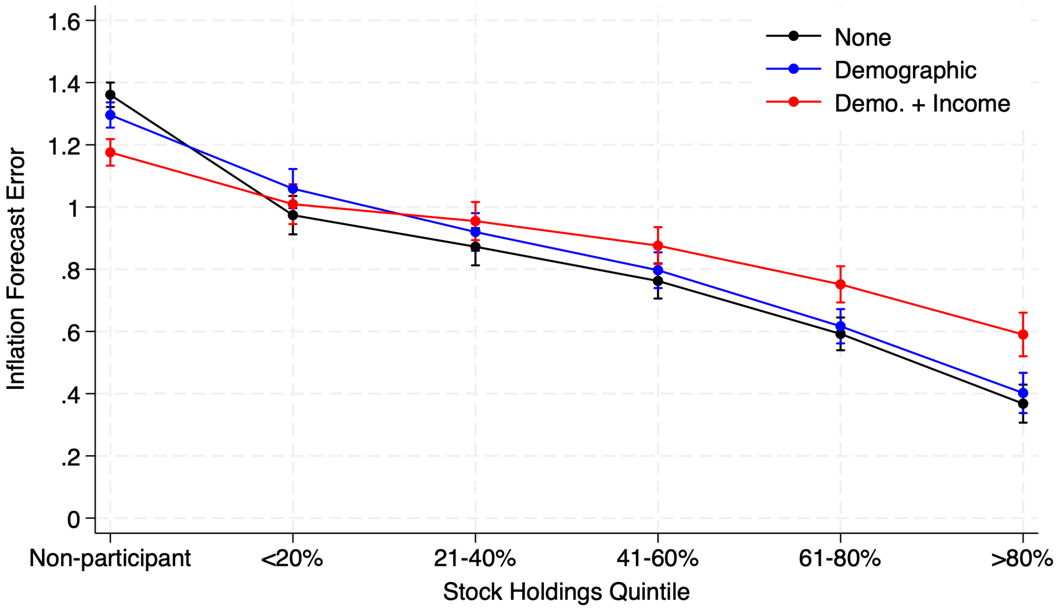

This graph plots the predicted inflation forecast error by stock-holdings group within the MSC. The black line shows the results of regression including time-fixed effects only. In blue, we also control for education, marital status, age, age squared and census region. In red, income group fixed effects are additionally included in the vector of controls. All standard errors are clustered at the individual level. The confidence bars show the 95% certainty range.

We show results from three differing specifications. The first, shown in black, is one in which we do not control for household demographic characteristics (i.e. the vector $\boldsymbol{X}_{i,t}$ is empty). The blue line shows a

[8] This is simply the *margins* command available in Stata, as well as other statistical software programmes.

specification in which we control for household demographics[9]. The red line additionally controls for income [10]. Across all specifications, the average forecast error of a respondent declines from roughly 1.3% for non-participants to around 0.5% for those in the top quantile of stock-holdings. Notably, this relationship is unchanged when we control for demographics. Moreover, even after controlling for income, the qualitative relationship between inflation forecast errors and stock holdings survives. Quantitatively, the difference between non-participants and the top quantile reduces from around 0.8% to 0.6%.

To what extent do stock holdings represent a good proxy for net wealth? For this measure to be a good proxy for household wealth requires two things. First, that the *level* of stock-market investments is non-decreasing with wealth. This would seem uncontroversial. Secondly, it requires that there are limits to the heterogeneity of stock-market investment preferences of consumers. This is harder to assess: the presence of respondents who would *never* invest in the stock-market even if they became extremely wealthy would render stock market holdings a poor proxy. However, there is evidence to suggest that transaction costs, rather than preference or risk-appetite heterogeneity explains a greater proportion of the puzzle of stock-market participation[11]. This would suggest that non-participants do not participate because they have low income or wealth, rather than preference heterogeneity. In the MSC data there is evidence to support this: the average income of a stock-market non-participant is approximately $44,000 versus $105,000 for those who participate, in 2015 dollars. If we are to interpret this measure as a proxy for net household wealth, we therefore must assume that any preference heterogeneity is small enough so as to keep households within one quintile of the level of stock holdings.

It is this result by which we motivate exploring a model in which households simultaneously possess inflation expectations which are typified by positive forecast errors, alongside a portfolio choice decision.

3.2.1 Modelling Pessimistic Expectations

In order to allow agents' expectations to diverge from rational expectations, we employ the robustness environment of Hansen and Sargent (2001). We outline in a general environment how such a specification can generate positive forecast errors. Take an agent with state variable x_t and choice variable c_t. The state variable is stochastic: innovations are determined by $dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dW_t$, where dW_t

[9]These are: Education, Marital Status, Age, Age Squared, Census Region. These are constructed in the same manner as in Chapter 1. Please therefore refer to Appendix A.1 for more detail.

[10]We group respondents by self-reported income (measured in 2015 dollars) into 11 bins, and include indicator variables in the vector $\boldsymbol{X}_{i,t}$. We group households to: 1) attempt to reduce the effect of measurement error on the regression; 2) to flexibly allow for the marginal effects of income on inflation forecast errors to be non-linear.

[11]Vissing-Jørgensen (2002) find that a modest transaction cost of $50 can explain almost half of stock-market non-participation.

is standard Brownian motion. Hansen and Sargent (2001) show that a consumer who seeks robustness (i.e. is unsure of the underlying data-generating process governing dx_t) solves the following Hamilton-Jacobi-Bellman (HJB) equation:

$$V(x) = \max_c \min_h u(c,x) + \frac{1}{2\theta}h^2 + \big(\mu(c,x) + \sigma(c,x)\cdot h\big)\cdot \partial_x V(x) + \frac{1}{2}\sigma(c,x)^2 \partial_{xx} V(x)$$

Where θ represents the agent's preference for robustness, and the choice variable h represents the distortion to the true data-generating process the agent employs in order to obtain robustness. For further detail on why the consumer's problem can be represented like so please see Appendix C.2.

Such an environment allows us to parsimoniously capture "pessimistic" expectations in a dynamic setting. If $\theta \to 0$, and the consumer is an expected-utility maximiser, then their expectation of movements in the state variable at time t is simply given by $\mathbb{E}_t[dx_t] = \mu(c_t, x_t)$. The agent's expectation of dx_t are therefore consistent with the rational expectation of the evolution of the state variable. However, when $\theta > 0$ their time-t conditional expectation[12] is: $\mathbb{E}_t[dx_t] = \mu(c_t, x_t) + \sigma(c,x)\cdot h^*$, where h^* denotes the h that obeys the first-order condition of the HJB. These are

$$h^* = -\theta \cdot \sigma(c,x) \cdot \partial_x V$$

Using this in the equation for $\mathbb{E}_t[x_t]$ yields:

$$\mathbb{E}_t[dx_t] = \mu(c_t, x_t) - \theta \cdot \sigma(c,x)^2 \cdot \partial_x V(x)$$

This equation highlights that the difference between the utility maximising agent's expectations and the robust agent's expectations rely on three terms. First, the sign of the difference depends on the marginal utility with respect to the state. For state variables in which utility is increasing in x_t (i.e. $\partial_x V(x) > 0$), the robust agent will have an expectation below (or more "pessimistic") than the expected utility maximiser. Secondly, the expectational difference is increasing in the agent's preference for ambiguity (θ) - those agents which seek a greater degree of robustness to misspecification will also be associated with larger deviations from the rational expectation. Finally, the discrepancy is increasing in the degree of risk (i.e. the variance of x_t, $\sigma(c_t, x_t)^2$) which the possibly misspecified state represents to the consumer. Our empirical motivation highlighted that inflation forecast errors across respondents in the MSC are positive. For this to be consistent with this model, it would therefore be the case that agents perceive

[12] As measured using their ex-ante assessment of probabilities of differing states of the world. Or, in the language of the robustness literature, the agent's reference model measure q_0.

inflation to be negative for welfare.

However, our empirical results highlight that inflation forecast errors decline as stock holdings rise. Interpreted through the lens of the model to come, we therefore require a non-homothetic specification to ensure that the worst-case distortion h^* can decline as our state-variable (household wealth) increases[13]. It is here that we make a departure from the robust portfolio-choice literature, which typically employs "homothetic robustness" and scales the preference for robustness parameter (i.e., θ) with the inverse of the value function by defining $\theta(V) = \frac{\theta}{V(x)}$ (Maenhout (2004)). This is, of course, useful for finding closed-form solutions to portfolio allocation problem and without such an assumption closed-form solutions do not exist. Under homothetic robustness, the worst-case distortion h^* is no longer a function of state-variables. We therefore focus on the case in which we maintain the standard Hansen and Sargent (2001) framework, in which concern for robustness is non-homothetic. This comes with a benefit and a cost - while it allows us to capture the extent to which pessimistic beliefs change as income changes, it means that closed-form solutions to the agents' HJB will be unobtainable[14].

3.3 Basic Model

We detail a continuous-time dynamic portfolio allocation problem, in which the consumer maximises expected utility from consumption of a single good, and has access to two financial assets: a real risky asset, and a nominal bond. The objective function the consumer maximises is:

$$U_t = \int_0^\infty e^{-\rho t} u(c_t)\, dt$$

The economy is subject to inflation risk. Specifically, the log-price of the consumption good evolves according to a standard Brownian motion:

$$d\log(P_t) = \bar{\pi} dt + \sigma_p dW_t^p$$

[13]To take a simple - yet common - example, let $\sigma(c, x) = \sigma \cdot x$. When we have homothetic robustness, $h^* = -\theta \cdot \sigma \cdot \frac{\partial_x V(x)}{V(x)}$. With homotheticity $\sigma \cdot x \cdot \frac{\partial_x V(x)}{V(x)} \neq f(x)$, and h^* is constant for all x.

[14]Beyond being able to fit the data more closely, we may have deeper reasons as to prefer the non-homothetic specification. As discussed extensively in Pathak (2002) there are potentially some concerns with the normalisation introduced by Maenhout (2004). Hansen and Sargent (2005) motivate the "multiplier" preferences detailed above as the Lagrangian multiplier on the constraint $\mathcal{R}(q) \leq \tau$, where τ represents the total lifetime entropy available to the decision maker to adjudicate alternative models by. This serves as a recursive formulation that captures the same consideration as Gilboa and Schmeidler (1989). Therefore making θ state dependent moves us away from this world, with its preference representation of ambiguity-aversion, and towards something else.

where $\bar{\pi}$ is the long-run average inflation rate. Here, dW_t^p is a standard Wiener process, and so the variance of $\log(P_t)$ is given by σ_p^2. The (nominal) price of the two assets evolves as:

$$dS_t = r_s S_t dt + \sigma_s S_t dW_t^s$$
$$dB_t = rB_t dt$$

where r_s and r denotes the instantaneous return on the risky-asset and nominal bond. In nominal terms, only the price of the risky asset is stochastic, as dW_t^s is a standard Wiener process. The risky asset therefore has variance σ_s^2.

Finally, we allow for correlation between inflation and the risky-asset return, by setting $\mathbb{E}[dW_t^p \cdot dW_t^s] = \rho_{sp}\ dt$. We can rewrite the process for the risky-asset price as follows: $dZ_t = \rho_{sp}\ dW_t^p + \left(1 - \rho_{sp}^2\right)^{\frac{1}{2}}\ dW_t^z$, where dW_t^z is a Wiener process independent of dW_t^p.

In nominal terms, the budget constraint of the consumer can be written as:

$$\dot{B}_t + \dot{S}_t = Y_t + rB_t + r_s S_t - P_t c_t$$

Let A_t denote total financial assets such that $A_t \equiv S_t + B_t$, and use lower-case letters to denote variables in real-terms (i.e. $x_t \equiv \frac{X_t}{P_t}$). Define $\omega_t \equiv \frac{s_t}{a_t}$ as the portfolio weight of risky assets. An application of Itô's Lemma yields[15] the consumers flow budget constraint denominated in real terms:

$$da_t = \left[y_t + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2 + (r_s - r + \sigma_s\sigma_p\rho_{sp}) \cdot \omega_t\right) \cdot a_t - c_t\right] dt +$$
$$+(\sigma_s\rho_{sp}\omega_t - \sigma_p) \cdot a_t\ dW_t^p + \sigma_s(1 - \rho_{sp}^2)^{\frac{1}{2}}\omega_t a_t dW_t^z$$

Note that the term $\sigma_s\sigma_p\rho_{sp}$ appears in the excess return expression. If $\rho_{sp} > 0$, then the risky asset is a hedge against inflationary shocks, and so earns a premium. We assume that the equity premium $r_s - r + \sigma_s\sigma_p\rho_{sp} > 0$.

Finally, alongside the "aggregate" risk that dW_t^z and dW_t^p represent, the agent is exposed to idiosyncratic risk as well. Their real income y_t follows a Poisson process, such that their income jumps from y_i to y_j with intensity λ_{ij}.

In order to be able to match the empirical profile of inflation expectations to the model, we allow the consumer to seek robustness to misspecification. They are concerned that their model of aggregate

[15]See Appendix C.3 for the full derivation.

fluctuations is incorrect, and so seek robustness towards the process dS_t and dP_t. We can therefore use Girsanov's theorem to rewrite da_t using the distorted Brownian motions $\hat{W}_t^p = W_t^p + \int_0^t h_\tau^p d\tau$ and $\hat{W}_t^z = W_t^z + \int_0^t h_\tau^z d\tau$. For the sake of reducing the state space, we combine these distorted Brownian motions into an aggregate (distorted) process for da_t. This does not affect the agent's choice of distortion, as we have written the problem in terms of independent Brownian motions dW_t^p and dW_t^z.

$$da_t = \Bigg[y_t + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2 + (r_s - r + \sigma_s\sigma_p\rho_{sp})\cdot\omega_t\right)\cdot a_t - c_t + \\ + (\sigma_s\rho_{sp}\omega_t - \sigma_p)\cdot a_t \cdot h_t^p + \sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}}\cdot\omega_t a_t \cdot h_t^z\Bigg]dt + \sigma_a(a_t,\omega_t)d\hat{W}_t^a$$

where we have:

$$d\hat{W}_t^a = \frac{(\sigma_s\rho_{sp}\cdot\omega_t - \sigma_p)}{\sigma_a(a_t,\omega_t)}\, d\hat{W}_t^p + \frac{\sigma_s(1-\rho_{sp}^2)\cdot\omega_t}{\sigma_a(a_t,\omega_t)}\, d\hat{W}_t^z$$
$$\sigma_a(a_t,\omega_t) \equiv \left(\sigma_s^2\cdot\omega_t^2 - 2\sigma_s\sigma_p\rho_{sp}\cdot\omega_t + \sigma_p^2\right)^{\frac{1}{2}}\cdot a_t$$

Finally, we impose a borrowing constraint and a no short-selling constraint on the consumer. Specifically, we impose that $B_t \geq -\underline{a}$ and $S_t = \omega_t \cdot a_t \in [0, a_t]$. That is, the consumer can borrow in the safe asset, but cannot hold a short position in the risky-asset. In terms of total financial wealth, these constraints can be written as $a_t \geq -\underline{a}$ and $\omega_t \in \left[0, 1 + \frac{\underline{a}}{a_t}\right]$ if $a_t > 0$ and $\omega_t = 0$ if $a_t \leq 0$[16]. As in many portfolio choice models of households (Gomes (2020)), our restrictions on the portfolio shares ω will mean that the optimal share may be at a boundary solution, as happens in our baseline calibration of Section 3.4. We view this a consequence of writing the model to describe households, rather than investors more generally.

In such an environment, we can write the consumer recursive problem (HJB) as:

$$\rho V(a,y) = \max_{c,\,\omega}\ \min_{h^z,\,h^p}\quad u(c) + \frac{1}{2}\left(\frac{1}{\theta_z}(h^z)^2 + \frac{1}{\theta_p}(h^p)^2\right) + \\ + \Bigg[y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2 + (r_s - r + \sigma_s\sigma_p\rho_{sp})\cdot\omega\right)\cdot a - c \\ + (\sigma_s\rho_{sp}\omega - \sigma_p)\cdot a\cdot h^p + \sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}}\cdot\omega\cdot a\cdot h^z\Bigg]\cdot\partial_a V(a,y) + \\ + \frac{1}{2}\left(\sigma_s^2\omega^2 - 2\sigma_s\sigma_p\rho_{sp}\cdot\omega + \sigma_p^2\right)\cdot a^2\cdot\partial_{aa}V(a,y) + \sum_{y\neq y'}\lambda_{y\to y'}(V(a,y') - V(a,y)) \tag{3.1}$$

where we have

$$a \geq -\underline{a} \quad , \quad \omega \in \begin{cases} \left[0, 1 + \frac{\underline{a}}{a}\right] & \text{if} \quad a > 0 \\ \{0\} & \text{if} \quad a \leq 0 \end{cases}$$

Here we have allowed the consumer to have potentially differing preferences over robustness towards misspecification of the risky-asset process and misspecification of the inflation process, as in Yang (2020). This will allow us to target two separate moments of the data when we calibrate the model.

Expected Utility Benchmark. To ground the results that follow, we first present the solution to the consumer's problem in benchmark case in which agents act like expected utility maximisers. This problem is the $\theta_p, \theta_z \to 0$ limit of equation (3.1).

The first order conditions are then given by:

$$u'(c^*) = \partial_a V(a, y)$$

$$\omega^* = \begin{cases} 0 & \text{if} \quad \omega' \leq 0 \quad\quad \text{or} \quad a \leq 0 \\ \omega' & \text{if} \quad \omega' \in \left(0, \frac{\underline{a}}{a}\right) \ \text{and} \ a > 0 \\ \frac{\underline{a}}{a} & \text{otherwise} \end{cases} \quad , \quad \omega' = \frac{-\partial_a V(a, y)}{a \cdot \partial_{aa} V(a, y)} \cdot \left(\frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s^2}\right) + \frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2}$$

The equation for the interior solution of the portfolio share (ω') is the standard Merton (1971) solution in the presence of a (potentially) hedging asset. The first term is the myopic demand for the risky-asset. The term in brackets denotes the risk premium the asset provides, which when divided by the variance of the risky returns yields the Sharpe-ratio. The term premultiplying this governs the risk tolerance of the household. When utility is CRRA, $\frac{-\partial_a V(a,y)}{a \cdot \partial_{aa} V(a,y)}$ is simply the inverse of the coefficient of relative risk-aversion, by the first-order condition on consumption[17]. The second term denotes the hedging demand for the risky-asset. If ρ_{sp} is positive, then the risky asset provides a hedge against inflation and the household demands more of the asset.

[16]Even though these constraints are easier to write if the household's choice variable is the outright risky-asset holdings (rather than the share), we keep the household's problem specified in terms of the portfolio share for the sake of comparison to previous work. When a strict no-borrowing constraint is used, $\underline{a} = 0$, this constraint collapses to the more familiar $\omega \in [0, 1]$. However, it rules out the possibility of the household issuing debt while also investing in risky assets. Such a possibility would mean that households would be able to engage in (relatively) sophisticated trading strategies. We assert that only a very small subset of households would engage in such short-long strategies in reality, and so rule it out..

[17]If the parameter of relative risk-aversion is denoted by γ and we set hedging demand equal to zero (i.e. $\rho_{sp} = 0$, this would then collapse exactly to the constant optimal share as derived by Merton (1971): $\omega' = \frac{r_s - r}{\gamma \cdot \sigma_s^2}$.

Discussion of assumptions. In constructing the households problem and the stochastic environment, we have made a particular set of assumptions. We discuss the implications of those here.

Nominal Risky Asset. We keep the asset structure of the economy simple for the sake of exposition. Even though the risky-asset is denominated in nominal units, by allowing for possibility of the risky asset returns and the inflation process to be correlated the environment remains flexible enough to allow the asset to provide a real return. This is the limiting case of $\rho_{sp} = 1$, where the risky-asset provides a perfect hedge against fluctuations in the nominal price level (i.e. is effectively denominated in units of the consumption good).

Nominal Bonds. Secondly, we allow for (limited) borrowing in the nominal bond. Again, we view this as reflective of the types of debt households most commonly have on their balance sheet (i.e. loans or mortgages which are denominated in units of currency). If bonds were inflation-indexed, their real payoff would be risk-free. This would make them a dominant savings vehicle for cautious households. Savers would seek to increase their holdings (and so reduce the risky-share of total assets), and borrowers would reduce their holdings, as unexpected inflation - when marginal utility is high - no longer erodes the value of the debt. In other words, the 'inflation put' that nominal debt represents disappears, raising the effective cost of borrowing using this instrument when inflation is high. Were the household able to trade real bonds, alongside nominal bonds, the real debt would dominate the nominal bonds if issued at the same rate. Were the nominal bonds to offer an excess return (i.e. an inflation risk premium), then the nominal debt would be priced in a similar manner to the risky-asset. This is what we would expect in general equilibrium. If the borrowing constraints were sufficiently lax, a sophisticated investor could hedge the inflation risk entirely if both assets are freely traded and correctly priced. Were that the case, the optimal distortion (h^{p*}) would be driven to zero for all households, as the agent's welfare would no longer be exposed to inflation. Indeed, once dW^p is spanned, the only source of risk is the orthogonal shock dW^z. Given that: 1) the purpose of writing the model is to investigate the implications of household's having persistent inflation forecast errors, and; 2) even the most sophisticated investor would struggle to hedge inflation costlessly (Campbell, Shiller, et al. (2009)); we wish to ensure that households *do* face nominal risk, and cannot hedge away their exposure. We therefore rule out such inflation hedging possibilities. In Section 3.1 we provide some empirical justification for this assumption.

Maturity Structure. Similarly, if the household were able to issue or purchase claims of different maturities, this would change the analysis too. First, we would have to introduce short-rate risk (of the Vasicek (1977) variety perhaps) and an extra state-variable. Poorer households, more worried about short-rate risk due to the non-homotheticities of the concern for model misspecification, would buy a greater proportion of long-term bonds to hedge this risk; this is the logic presented in Campbell and Viceira (2001). Whether the portfolio share changes would be dependent on the covariance structure of the short-rate and inflation, and the relative hedging quality of long-term bonds versus equities. The overall worst-case distortions h^{p*} and h^{z*} would likely decrease relative to the baseline model, as the agent would now be able to "push" the worst-case distortions into parts of the yield curve it doesn't hold. We keep the assets as maturing in one period for the sake of parsimony.

Short Sales. We also restrict households to not being able to short-sell the risky-asset. As we are interested in the effect of persistent household forecasting errors on portfolio choice, we view this as reflective of the investment problem of a household. This is therefore not a model of a sophisticated investor, who may have the ability to take short positions in particular assets. It is also why we restrict the portfolio share of risky assets to be 0 if the household is a net borrower: we want to rule out the possibility of the household engaging in short-long strategies across assets too. This does restrict further the ability of the household to hedge itself against inflation when $\rho_{sp} < 0$ somewhat. In Section 3.3.2 we highlight evidence which suggests $\rho_{sp} < 0$.

Multiple Risky Assets. We also restrict the household to only holding one type of stock. This is of course an abstraction from reality. We view the risky-asset as therefore representing a representative portfolio of stocks, rather than an individual stock. If we let the number of available risky assets increase, the analysis would remain the same under the condition that the returns did not complete the market for the inflation risk dW_t^p. However if that were the case, the household would be able to construct a portfolio that perfectly hedges the inflation risk in the economy. Again, introducing such a possibility would run counter to the purpose of the model.

Our assumptions on the asset structure of the economy ensure two things. First, that the model remain tractable, while maintaining a semblance of realism about the investment environment faced by households. Secondly, that households cannot hedge inflation. Were they able to do so, a model with nominal risk would be of no interest.

Inflation Process. In the model, $\log(P_t)$ follows (the continuous time equivalent) of a random walk. In reality, using an Ornstein-Uhlbeck process (i.e. an AR(1)) to model the inflation process would be a closer to fit to reality, as inflation is typically persistent but mean-reverting. Changing the stochastic process for the price level would be simple enough, but comes at the cost of introducing another state-variable. It also makes the numerical solution to the households problem less numerically stable when calibrated, as with persistent inflation the rate of return often turns negative for reasonable values of σ_p. Were we to modify the stochastic environment, our measure of distance between the robust agent's forecast and the expected-utility maximisers forecast would remain the same. However it would incentivise more complex savings and investment: with $\rho_{sp} < 0$, households would find it optimal to invest more in the risky asset when inflation is high: as the price-level mean reverts, returns will increase at the same time dis-inflation occurs. While sophisticated investors may restructure their portfolio like so during times of high inflation, such behaviour is unrealistic for the vast majority of households.

3.3.1 Properties of Optimal Policies

We first provide a set of results characterising the choices of the household in the presence of non-homothetic robustness. We characterise how the optimal distortions of the agents (and therefore their implied forecast errors) behave within this model.

First, define $\Delta_p^E(a, y) \equiv \mathbb{E}^{\text{robust}}[d\log(P_t)] - \mathbb{E}^{\text{EU}}[d\log(P_t)]$ as the difference in inflation expectations between the robust agent and their expected-utility maximising counterpart.

Similarly define $\Delta_r^E(a, y) \equiv \mathbb{E}^{\text{robust}}\left[\frac{dS_t}{S_t}\right] - \mathbb{E}^{\text{rational}}\left[\frac{dS_t}{S_t}\right]$ as the difference in the perceived drift of the risky-asset returns between a robust and expected utility maximising agent.

We introduce the following set of assumptions:

A1 **Preferences.** $u\colon \mathbb{R}_+ \longrightarrow \mathbb{R}$ is twice continuously differentiable, strictly increasing, strictly concave: $u \in \mathcal{C}^2,\ u'(\cdot) > 0,\ u''(\cdot) < 0$. Furthermore, let $u(\cdot)$ satisfy Inada conditions: $\lim_{c\downarrow 0} u'(c) = \infty$, $\lim_{c\to\infty} u'(c) = 0$.

A2 **Relative risk aversion bounds.** For every state (a, y), define $R(a, y) = -\dfrac{c^*(a, y) \cdot u''(c^*(a, y))}{u'(c^*(a, y))}$. Let relative risk aversion be finite and positive: i.e. there exist constants $0 < \underline{\gamma} \leq R(a, y) \leq \bar{\gamma} < \infty$.

A3 **Income and borrowing.** Labour income $y_t > 0$ in all states; assets satisfy $a_t \geq 0$; no short-selling so the portfolio share obeys $0 \leq \omega^*(a, y) \leq 1$.

A4 **Model parameters.** $\theta_p, \theta_z, \sigma_p, \sigma_s > 0$ and $\rho_{sp} \in [-1, 1]$. Throughout we impose the "low correlation" restriction $\rho_{sp} \leq \sigma_p/\sigma_s$.

A5 **Decay of marginal utility.**[18] There exists $\varepsilon > 0$ s.t. $a^{1+\varepsilon} u'(c^*(a, y)) \xrightarrow[a\to\infty]{} 0$.

A6 **Penalty dominance.** $\theta_z\,(1 - \rho_{sp}^2) \ \geq\ \theta_p\,\rho_{sp}\,\frac{\sigma_p}{\sigma_s}$.

[18] This is the only place where a slightly stronger tail assumption than Inada is used. It simply ensures that marginal utility declines slightly quicker than $\frac{1}{a}$.

Note that (A3) imposes the extra condition that $\underline{a} = 0$.

We present three key results: first, we show that households expect inflation to be higher than the rational benchmark; second, we show that the difference between the household's expectation and the rational benchmark declines in wealth; third, we show that for reasonable empirical calibrations (i.e. $\rho_{sp} \leq 0$), optimal portfolios are lower than in the expected utility case, but that this effect declines as wealth increases.

1. Households have positive forecast errors. The first order conditions with respect to the distortions are given by:

$$h^{p*} = -\theta_p \cdot \left(\sigma_s \rho_{sp} \cdot \omega^* - \sigma_p\right) \cdot a \cdot \partial_a V$$
$$h^{z*} = -\theta_z \cdot \sigma_s (1 - \rho_{sp}^2)^{\frac{1}{2}} \cdot \omega^* \cdot a \cdot \partial_a V$$

A positive inflation distortion h^{p*} shifts the perceived drift of inflation upwards; similarly, a negative h^{z*} lowers the perceived risky-asset returns. Lemma 3.3.1 details the conditions under which $h^{p*}(a, y) > 0$ and $h^{z*}(a, y) < 0$. If Assumptions (A1)-(A4) hold, then optimal distortion is one in which the mean inflation rate is perceived to be higher and the average risky return is perceived to be lower. This then implies that $\Delta_p^E(a, y) \geq 0$: that the robust agent perceives the process for inflation to be greater than the non-robust agent. Empirically, such behaviour is consistent with positive inflation forecast errors.

Lemma 3.3.1 (Signs of optimal distortions)**.** *Under Assumptions (A1)–(A4),*

$$h^{p*} \geq 0, \qquad \Delta_p^E \geq 0 \quad \forall a \geq 0,\ y > 0.$$

Assuming also (A6):

$$h^{z*} \leq 0, \quad \Delta_r^E \leq 0 \quad \forall a \geq 0,\ y > 0.$$

Both wedges equal zero only at $a = 0$ or $a = \infty$. Proof: See Appendix C.4.

If we were to relax the assumption that $\underline{a} = 0$, and instead let the consumer borrow in the nominal bond, the sign of h^{p*} would flip, and we would have $\Delta_p^E < 0$ for $a < 0$. That is to say, household borrowing in the nominal bond would display *negative* forecast errors. This would run in opposition to the empirical evidence (although Masolo and Monti (2024) find evidence to suggest that borrowers do indeed have lower inflation expectations than savers). This feature is an artefact of the simplicity of the model environment. All that has to be added to restore the positive forecast error for some values of $a < 0$ is for households to at least perceive inflation to be welfare-negative. This could be achieved by introducing nominal wages, for example.

2. Expectation distortions decline in wealth. How does pessimism change in the cross-section? As there is no known closed form solution for $V(a, y)$ when robustness is non-homothetic we cannot simply solve the HJB for an analytic solution. However we can characterise the distortions with respect to wealth. As wealth increases so too does the optimal (total) risky-asset holdings. At the same time the agents value the extra unit of wealth less as marginal utility declines. The question of which of these two effects wins out determines whether or not the distortion is increasing or decreasing with wealth. Proposition 3.3.1 characterises the conditions under which the inflation forecast error $\Delta_p^E(a, y)$ is initially increasing in a, and eventually decreasing in a. At very low values of wealth, the marginal utility effect wins out and the expectation wedge is increasing. But for high values of wealth, under Assumption (A5) the marginal utility of consumption declines sufficiently quickly so as to ensure the optimal distortion declines to zero.

Proposition 3.3.1 (Slope near the bounds)**.** *Assume two further refinements to (A2):*

- *A2(b)* ***Asymptotic constant RRA.*** $\lim_{a\to\infty} R(a, y) = \gamma_\infty \in (0, \infty)$
- *A2(c)* ***Relative Prudence Bounds.*** $u \in C^3$. $P(a, y) := -\frac{c^*(a,y)\cdot u'''(c^*(a,y))}{u''(c^*(a,y))}$ *where* $\sup_a P(a, y) < \infty$.

Then under Assumptions (A1)–(A5), there exist $0 < a_0(y) < a_1(y) < \infty$ *such that*

$$\frac{d\Delta_p^E}{da} > 0 \text{ on } (0, a_0), \quad \frac{d\Delta_p^E}{da} < 0 \text{ on } (a_1, \infty),$$

Further assume (A6):

$$\frac{d\Delta_r^E}{da} < 0 \text{ on } (0, a_0), \quad \frac{d\Delta_r^E}{da} > 0 \text{ on } (a_1, \infty).$$

If (A2(b)-(c)) hold alongside $\underline{\gamma} > 1$*, the large-a inequalities hold for all* $a \geq a_1(y)$*. Proof: See Appendix C.4.*

These conditions are worth discussing. Assumption (A4) requires the correlation between inflation and risky-asset returns are sufficiently low. The gradients of the worst-case distortions depend on how an extra unit of wealth alters the consumer's marginal utility, as well as the covariance structure embedded in the budget constraint. When $\rho_{sp} \leq \frac{\sigma_p}{\sigma_s}$ - the empirically relevant case - and the risky-asset is a poor hedge for inflation, then this is sufficient to ensure that the gradient doesn't change sign. Assumption (A5) requires that the marginal utility of consumption declines sufficiently quickly so as to ensure the term $\lim_{a\to\infty} \; a \cdot \partial_a V = 0$. While this is a stronger assumption than the typical Inada conditions, it is a property of many frequently used utility specifications: for example, CRRA utility with relative risk aversion > 1 and Epstein-Zin utility with an elasticity of inter-temporal substitution < 1 both satisfy this condition. The technical conditions (A2(b)) and (A2(c)) are both unrestrictive, and are also features of many popular utility functions employed in macroeconomics.

So far we have characterised the gradient of the inflation expectation wedge with respect to a, but we still are unsure of the asymptotic behaviour. We therefore characterise the asymptotic behaviour of the wedge in this environment in Lemma 3.3.2. Here, we illustrate that as wealth grows the optimal distortions converge to zero: in the limit, the consumer becomes an expected-utility maximiser.

Lemma 3.3.2 (Boundary limits)**.** *Under Assumptions (A1), (A3), (A5),*

$$\lim_{a\downarrow 0} \Delta_p^E(a,y) = 0, \qquad \lim_{a\uparrow\infty} \Delta_p^E(a,y) = 0$$

$$\lim_{a\downarrow 0} \Delta_r^E(a,y) = 0, \qquad \lim_{a\uparrow\infty} \Delta_r^E(a,y) = 0$$

Proof: See Appendix C.4.

Lemmas 3.3.1, 3.3.2 and Proposition 3.3.1 taken together allow us to characterise the full shape of the inflation expectation wedge in this model. Under Assumptions (A1)-(A5), the difference between the inflation expectation of a robust agent and an expected utility maximiser is hump-shaped[19]. Proposition 3.3.2 formalises this result.

Proposition 3.3.2 (Hump-shaped inflation-expectation wedge)**.** *Assume two further refinements to (A2):*

- *A2(b)* ***Asymptotic constant RRA.*** $\lim_{a\to\infty} R(a,y) = \gamma_\infty \in (0,\infty)$
- *A2(c)* ***Relative Prudence Bounds.*** $u \in C^3$. $P(a,y) := -\frac{c^*(a,y)\cdot u'''(c^*(a,y))}{u''(c^*(a,y))}$ *where* $\sup_a P(a,y) < \infty$.

Assume (A1)–(A6), alongside (A2(b)), (A2(c)) and $\underline{\gamma} > 1$*. Then, for every* $y > 0$*,*

- $\Delta_p^E(a,y)$ *is non-negative, single-peaked, and attains a unique interior maximum;*

Further assume (A6)

- $\Delta_r^E(a,y)$ *is non-positive, single-valleyed (inverse hump), and attains a unique interior minimum.*

[19]The evidence presented in Section 1.4.1 is suggestive of a hump-shape profile of inflation forecast errors across income.

Consequently, $\Delta_p^E(a,y)$ and $\Delta_r^E(a,y)$ admit unique extrema ($a^*_{\Delta p}(y)$ and $a^*_{\Delta r}(y)$) that satisfy the first-order conditions:

$$\frac{d\Delta_p^E(a^*_{\Delta p},y)}{da} = 0 \qquad \frac{d\Delta_r^E(a^*_{\Delta r},y)}{da} = 0$$

Proof: See Appendix C.4.

3. The distortion to portfolio allocation disappears for the richest households. We now turn our attention to the agent's portfolio allocation. Using the first order conditions, we can write the (interior) optimality condition on risky asset portfolio share as:

$$\omega' = \left[\frac{r_s - r}{\sigma_s^2} + \frac{\sigma_s\sigma_p\rho_{sp}}{\sigma_s^2} + \frac{\sigma_s\sigma_p\rho_{sp}}{\sigma_s^2}\cdot\frac{a\left(\theta_p(\partial_a V)^2 - \partial_{aa}V\right)}{\partial_a V}\right]\cdot\frac{-a\cdot\partial_a V}{\partial_{aa}V - (\theta_z(1-\rho_{sp}^2) + \theta_p\rho_{sp}^2)(\partial_a V)^2}$$

This expression consists of three terms. The first term is typical myopic demand as in Merton (1971) and the second represents the standard hedging demand. These mechanisms are features of the expected-utility maximisers problem. The final term captures the effect of ambiguity on hedging demand. If $\rho_{sp} < 0$ and the numerator is positive, i.e. $\theta_p(\partial_a V)^2 - \partial_{aa}V) > 0$, then this term attenuates the hedging demand of the agent. The hedging ability of the risky-asset is called into question when there is model uncertainty. This decomposition is as studied in Maenhout (2004).

To assess the impact of robustness on portfolio allocations, we define a portfolio "wedge" ($\Delta\omega$). This measures the difference between the portfolio share optimal for an expected utility maximiser relative to the optimal share for a robust agent. Making explicit the dependence of the portfolio share on the concern for robustness, denote the portfolio share of an agent as $\omega(a|\theta_p,\theta_z)$. Formally:

$$\Delta\omega(a) \equiv \omega(a|0,0) - \omega(a|\theta_p,\theta_z)$$

The portfolio wedge $\Delta\omega(a) = 0$ when the robust investor with wealth a assigns the same proportion of their financial wealth to the risky-asset as an expected utility maximiser would do.

Proposition 3.3.3 details two properties of this portfolio wedge. This portfolio wedge is positive: robust investors allocate less to risky assets in this environment than their expected utility maximising counterparts. In the limit, the wedge tends to 0. That is to say: robust agents invest less in risky assets, but as wealth increases their portfolio shares tend to the expected-utility allocation.

Proposition 3.3.3 (Robust Portfolio Wedge)**.** *Assume (A1)-(A5), and let $\underline{y}$ be the income in the lowest income* *state. Then:*

$$\lim_{a \to \infty} \Delta\omega^*(a) = 0$$

Furthermore, under a parameter restriction on θ_p, we have:

$$\Delta\omega^*(a) \;\geq\; 0 \quad \forall\, a \quad \textit{if} \quad \theta_p \;\leq\; 1 + \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s \sigma_p |\rho_{sp}| \cdot u'(\underline{y})}$$

Proof: see Appendix C.4.

The parameter restriction of Proposition 3.3.3 is a sufficient condition (but not necessary). It places restrictions on the concern for robustness with respect to the inflation process (θ_p), in terms of primitives of the model. Simply put, model-uncertainty over the inflation process cannot be so strong that it completely eradicates the benefits the excess return of the risky-asset provides, once one accounts for the worst-case inflation distortion faced by the poorest household. For typical calibrations, this condition is unrestrictive.

The model therefore captures the two primary features of the data we outlined in Section 3.2: households have positive forecast errors, but these decline with wealth. An implication of such behaviour is that the difference between the pessimistic agents of our model and expected-utility maximisers is that they under-invest in risky assets, but again this effect declines with income. With these three features of the model in mind, we now move onto assessing their importance. We therefore seek to quantify the magnitude of these effects and calibrate the model.

3.3.2 Calibration

We turn now to calibrating and solving the model. Our goal is to implement a calibration which captures the aggregate inflation forecasting error, while also matching the equity investment behaviour of US households. We parametrise instantaneous utility with a CRRA specification ($u(c) = \frac{c^{1-\gamma}}{1-\gamma}$) with relative risk aversion $\gamma = 2$. We set the rate of time preference as $\rho = 0.05$, which in consistent with $\beta = 0.95$ in a discrete-time model. We calibrate the return process of the risky asset as having an average (nominal) return $r_s = 8\%$ and a standard deviation of $\sigma_p = 16\%$. These are values commonly used in the literature (Maenhout (2004), Peijnenburg (2018)) and are close to the long-term mean and variation of the S&P500 Equity Index. We set the nominal asset return $r = 3.5\%$ to aim for an average real return of 1.5%. For the inflation process, we set the average inflation rate as $\bar{\pi} = 2\%$ and measure the standard deviation of the BLS CPI in annual year-on-year terms over 1990-2020 as 0.015. We therefore set inflation volatility to $\sigma_p = 1.5\%$. An important parameter for the model is the correlation between inflation and risky-asset returns. Estimates in the literature typically range from -0.15 to -0.25 for the US (see Fama and Schwert (1977), Cieslak and Pflueger (2023)). We pick the midpoint of this range, and set $\rho_{sp} = -0.2$.

For the income process, we use a 3-state Poisson process and calibrate it so as to generate the wealth distribution. We take the income process to be $[0.1,\ 1,\ 10]$, and use the same income transition matrix as Nord (2022). As discussed by the author, a high earning state with lower persistence is requires to generate the long-right tail of the US wealth distribution.

To calibrate the robustness preferences, we target two seperate moments of the data[20]. We target the population average forecast error for inflation as measured in the Michigan Survey of Consumers (MSC) over the period 1990-2020: this implies an average consumer inflation expectation of 2.96%[21]. Our second target is the average risky-asset portfolio share of US households, which allows the model to capture the aggregate dynamics of wealth accumulation. The Federal Flow of Funds records US households as owning \$56tr worth of equity, out of a \$169tr net-worth (33%), or \$122tr worth of financial assets (46%). The review paper of Gomes (2020) suggests that roughly 50% of US households participate in the stock-market, who allocate 59.6% of their financial wealth in stocks. We therefore target an average risky asset portfolio share of 30%[22].

Unlike the detection-error method that has become standard in the robust-control literature, our calibration pins down the penalty parameters θ_z and θ_p by matching observed forecast errors and portfolio behaviour, rather than an information-theoretic bound. Using the detection-error method, (see Anderson et al. (2003), Maenhout (2006)), one computes the probability that a likelihood-ratio test would let an econometrician distinguish the agent's worst-case model from the approximating model over a finite sample: Anderson et al. (2003) recommend targeting a 10–20% detection error over ten years. The chosen θ is therefore governed by a statistical criterion: how large a misspecification the agent can entertain without it being easily refuted by the data the econometrician expects to observe. Our moment-matching approach reverses the order of reasoning. We start from survey evidence on inflation

[20]Specifically, we find the θ_z and θ_p that solve:

$$\sum_y \int \omega^*(a, y|\theta_z, \theta_p) \cdot g(a, y|\theta_z, \theta_p)\, da = 0.3$$
$$\sum_y \int \sigma_p \cdot h^{p*}(a, y|\theta_z, \theta_p) \cdot g(a, y|\theta_z, \theta_p)\, da = 0.96$$

.

[21]Section 1.4 provides more detail.

[22]Limited equity market participation is a common result in the data, across a range of developed economies. In our baseline calibration, all households choose to participate. Including a fixed-cost of equity market participation, or a much higher degree of ambiguity over risky returns could also replicate this. Peijnenburg (2018) study a life-cycle portfolio allocation problem with Gilboa and Schmeidler (1989) preferences and learning which can match equity-market participation rates over the life-cycle.

Table 3.1: Parameter Values and Targets

	Parameter	Value	Target
Risk aversion	γ	2	Standard
Time preference	ρ	0.05	$\beta \approx 0.95$
Nominal risky return	r_s	0.08	Real return = 6%
Std. risky returns	σ_s	0.16	Literature
Nominal bond return	r	0.035	Real return = 1.5%
Mean inflation	$\bar{\pi}$	0.02	US inflation target
Std. inflation shocks	σ_p	0.015	BLS CPI (1990-2020)
Inflation/returns correlation	ρ_{sp}	-0.20	Cieslak and Pflueger (2023)
Income states	y	[0.1 , 1, 10]	Wealth distribution
Income transition	Π_y	$\begin{bmatrix} 0.975 & 0.025 & 0 \\ 0.057 & 0.931 & 0.012 \\ 0 & 0.15 & 0.85 \end{bmatrix}$	Nord (2022)
Robustness, inflation	θ_p	2.63	Aggregate forecast error = 0.96%
Robustness, returns	θ_z	0.163	Aggregate risky share = 30%

Calibrated parameters for the baseline model. The final two parameters are estimated using moment matching.

expectations and aggregate equity shares, taken from the MSC and the BEA Flow of Funds respectively. These are moments that directly reveal households' perceived distribution of inflation and returns. We then choose θ_z, θ_p so that the model reproduces those revealed distortions.

This has three advantages. First, it anchors robustness to economically meaningful moments of the data, rather than to an arbitrary choice of sample length and significance level. Second, it avoids the question of whether respondents to consumer surveys or agents within our model (who are not sophisticated econometricians, as per the evidence presented by persistent forecast errors) *can* perceive deviations from their reference model of a particular probability. Third, it allows us to side-step the question of exactly how to employ the detection-error method in a model with a continuum of agents, where the worst-case distortion is heterogeneous across agents. In future work with more sophisticated calibrations, exploiting this heterogeneity using moment-targeting calibrations opens up the possibility of using the cross-sectional variation in expectations itself as an identification device. This is something that the detection error method, being a representative-agent likelihood-based criterion, cannot exploit. Of course, our strategy implies that were the empirical forecast errors themselves to change meaningfully, θ_p, θ_z should be re-estimated. On the other hand, the detection error method would keeps θ fixed so long as the detection probability benchmark is deemed reasonable to the modeller (cf. Ilut and Schneider (2023)).

We are able to match the moments for inflation forecasting errors and aggregate risky-share allocation exactly. Table 3.2 illustrates the fit of the model with regards to the wealth distribution. We compare the model moments to an SCF+[23] derived measure of net-financial wealth. Financial wealth inequality is considerable in the US, with nearly 90% of wealth owned by the top 20%[24] Our calibration generates significant wealth disparity and is able to match the concentration of wealth in the top quintile. However as we have imposed a no-borrowing constraint we cannot match the net borrowing behaviour of the lowest quintile.

Table 3.2: Model vs. Data Wealth Distribution

Quintile	1	2	3	4	5
Data	-2.80%	0.0234%	2.18%	12.2%	88.5%
Model	1.04%	1.82%	2.54%	9.47%	85.13%

Model vs. data wealth distribution in the baseline calibration. The data used is 2016 net financial wealth (ffafin - pdebt), as measured in the SCF+ (see Kuhn et al. (2020) for more).

Computational Solution The model is solved using the implicit finite-difference scheme as in Achdou et al. (2022). We discretise the state space and use non-linear grids for financial wealth a and the idiosyncratic income process. More detail can be found in Appendix C.5.

The only departure we take from the (now standard) finite-difference scheme of Achdou et al. (2022) is in our calculation of the Kolmogorov-Forward Equation (KFE), which we use to determine the stationary distribution of households. We require this to match the model to the moments of the data. As discussed there, a major benefit of using continuous-time solutions to solve heterogeneous agent models is that the infinitesimal generator of the agent's HJB is the adjoint operator of the generator in the KFE.

However in our model, we require a slightly different operator in the KFE than simply the adjoint of the HJB's infinitesimal generator. As the agents perceive a different law of motion than is actually realised, their wealth does not evolve according to the laws of motion implied by the HJB generator. When solving the KFE for the stationary distribution $g(a, y)$, we therefore solve the following equation:

$$0 = \partial_a(s^*(a,y) \cdot g(a,y)) + \frac{1}{2}\partial_{aa}(\sigma_a^2(a,\omega^*(a,y))g(a,y)) - \sum_{y \neq y'} \lambda_{y' \to y}(g(a,y') - g(a,y)) \tag{3.2}$$

where $s^*(a,y) := y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2 + (r_s - r + \sigma_s\sigma_p\rho_{sp}) \cdot \omega^*(a,y)\right) \cdot a - c^*(a,y)$ is the optimal savings policy according to the HJB, and $c^*(a,y)$ and $\omega^*(a,y)$ are the policy functions that solve the HJB maximisation.

That is to say, we solve the KFE removing the effect of the optimal distortions h^{p*} and h^{z*}: the equation

[23] This is a longer version of the Survey of Consumer Finances, compiled by Kuhn et al. (2020). It is easier to work with than the SCF itself, as it contains fewer variables and has already been cleaned significantly by the authors.

[24] More stark in fact than total net-wealth, of which 78% was held by the top quintile in 2016.

for the distribution is calculated using the true drift of the wealth process, rather than the perceived drift. This allows the model to capture that the agent's perceptions of the economy diverge from reality.

3.4 Quantitative Results

Figure 3.2 presents the risky asset portfolio share policy function across income states and wealth $\omega(a, y_i|\theta_z, \theta_p)$. In the left panel, we compare these policy functions to an agent in the same environment but with no preference for robustness, $\omega(a, y_i|0, 0)$. This is shown for reference, as only the robust model has been calibrated to match a 30% aggregate equity share. For each income state y_i and each level of wealth a the optimal portfolio allocation to risky assets is lower for robust investors relative to the expected utility benchmark. This difference is most pronounced for agents with high marginal utility: those in the lowest income state (displayed in blue), or lower levels of wealth. In our baseline calibration, the differences across income and wealth states are stark: the minimum portfolio share for low income households is roughly 10%, whereas those in the top income group have a risky-share at the boundary of 100%. Low wealth, low income households therefore allocate more of their savings to assets with nominal exposure relative to the expected utility case. Ex-ante, both low and high wealth households face the same price-level risk. However, ex-post, low income and low wealth households expose themselves to greater inflation risk than they would do otherwise, without a preference for robustness.

Figure 3.2: Risky-Asset Portfolio Shares

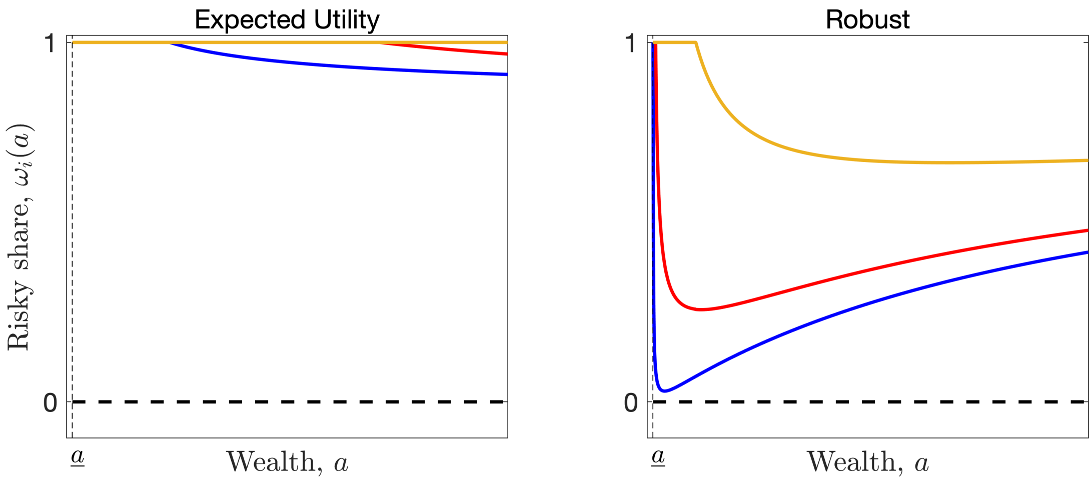

The left panel displays the risky-asset portfolio shares of an agent for whom $\theta_z,\ \theta_p = 0$; that is, the agent is an expected utility-maximiser. The second panel uses our calibrated values. The blue line denotes the policy function of agents in the lowest income state $\omega(a, y_1)$. The red and yellow lines denote $\omega(a, y_2)$ and $\omega(a, y_3)$ respectively.

Moreover, low (and medium) income households begin increasing their risky-asset share beyond a particular asset threshold. Studies such as Campbell (2006) and Fagereng et al. (2020) have found empirical evidence in support of this[25]. This behaviour is not captured in the standard portfolio model: the expected utility maximiser solving such a problem displays (weakly) declining portfolio shares.

Portfolio shares begin rising due to the effect of non-homothetic robustness. As touched upon earlier our calibration meets the criteria of Lemma 3.3.2 once away from the borrowing constraint. Therefore as wealth increases the optimal distortion declines, as marginal utility declines faster than the optimal portfolio share. Richer agents perceive the drift of the risky asset returns (and inflation) to now be higher, and so increase their risky-asset allocation.

Perceived Inflation. We next illustrate the perceived asset return and inflation drift across income states and wealth within the model. Substitute the distorted Brownian motions $\hat{W}_t^z$ and $\hat{W}_t^p$ for a household with wealth a_t in income state $y_{i,t}$ into the stochastic processes for dP_t and $\frac{dS_t}{S_t}$. Once we take expectations, we can write:

$$\begin{aligned} \mathbb{E}_t \left[d\log(P_t) \middle| \, a_t, y_{i,t} \right] &= \bar{\pi} + \sigma_p \cdot h^{p*}(a_t, y_{i,t}) \\ \mathbb{E}_t \left[\frac{dS_t}{S_t} \middle| \, a_t, y_{i,t} \right] &= r_s + \sigma_s \cdot \left(\rho_{sp} \cdot h^{p*}(a_t, y_{i,t}) + (1 - \rho_{sp}^2)^{\frac{1}{2}} h^{z*}(a_t, y_{i,t}) \right) \end{aligned}$$

Note that the perceived drift for the risky asset returns is a function of the optimal distortion for both the inflation process, and the uncorrelated component of risky-asset innovations. In Figure 3.3, we plot the perceived drift of the two stochastic processes between income groups and wealth levels.

The top panel displays the perceived drift of agents for the price-level process, and the bottom panel displays the risky returns process. The black dotted line represents the expectation of an expected utility maximising agent (i.e. the rational forecast under the true probability measure). The distance between these two lines, in an empirical setting, would therefore represent the measured forecast error of an agent with state vector $(a_t, y_{i,t})$. The blue, red and yellow lines denote households in the lowest, middle and highest income states, respectively. As illustrated, non-homothetic robustness generates significant heterogeneity across income and wealth. Using the solution to Equation (3.2), we calculate the average perceived drift between agents within a particular income group y_i. Those in the lowest income state perceive the drift of the price level to be 3.01%, whereas those in the middle and top income states perceive the drift to be 2.80% and 2.41% respectively.

Perceived Returns. Despite the low coefficient of robustness towards risky-asset returns estimated in the calibration exercise, the lower panel of Figure 3.3 displays large drift distortions for the risky returns

[25] The latter posits that one possible explanation is heterogeneous returns to wealth: in particular, wealthy households invest in higher return assets (even within an asset class).

Figure 3.3: Model Implied Expectations, Inflation and Risky Asset Returns

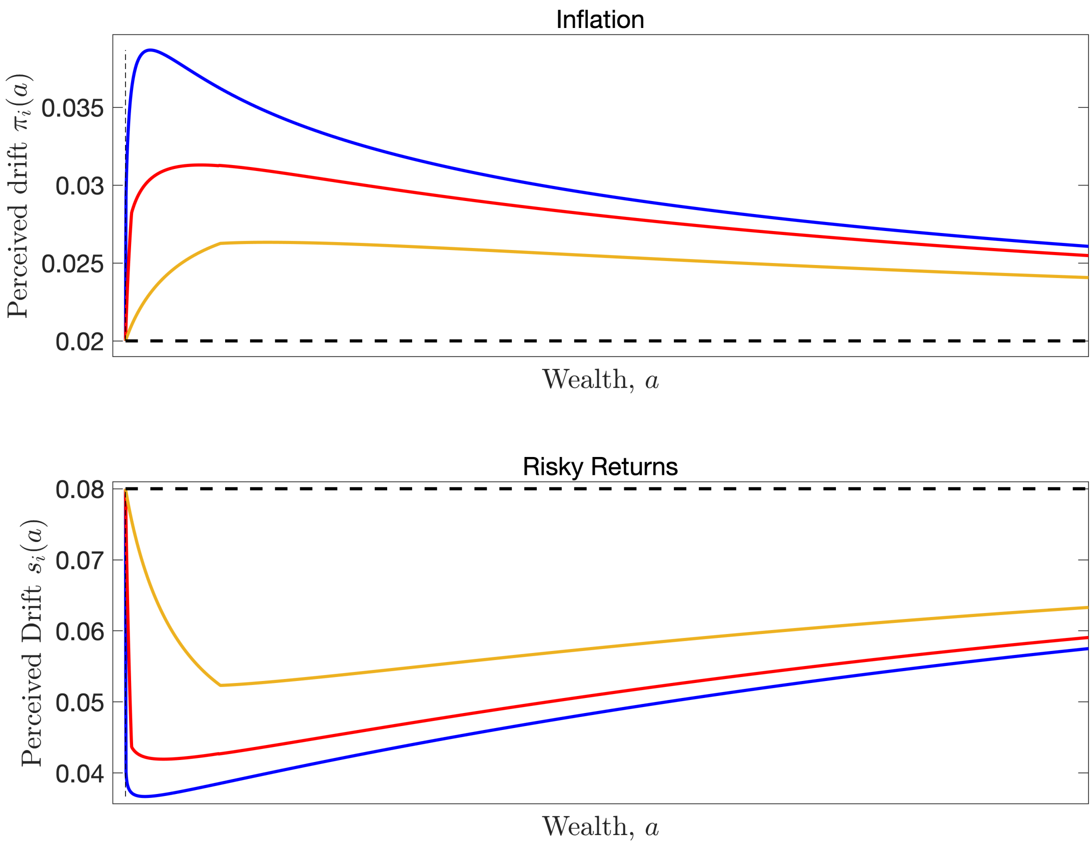

Here we plot model implied expectations by income state, across wealth. The top panel reports the agent's perceived drift of the price-level process: i.e. $\bar{\pi} + \sigma_p \cdot h^p(a, y_i)$. The bottom panel reports the agent's perceived drift of the risky-asset process, measured as $r_s + \sigma_s \cdot (\rho_{sp} \cdot h^p(a, y_i) + (1 - \rho_{sp}^2)^{\frac{1}{2}} \cdot h^z(a, y_i))$. The blue line denotes the policy function of agents in the lowest income state, y_1. The red and yellow lines denote y_2 and y_3 respectively.

process. This is because the perceived drift of the risky asset is a function of both the orthogonal returns distortion, and the distortion associated with the inflation process. The modest negative correlation between inflation and risky returns is sufficient to drive considerable divergence in the perceived drift of the risky asset: when the agent perceives the drift of the inflation process to be high, they also perceive the drift of the risky asset to be lower, irrespective of the distortion with respect to the uncorrelated component[26].

Model Fit. We next turn to assessing the fit of the distribution of inflation expectations within the model across the wealth distribution. In Table 3.4 we plot the average inflation forecast error, controlling for demographics and time-fixed effects[27] - as reported in the MSC. We contrast this with the model-

[26]The average perceived drift of stock returns in the model is 5.05%, compared to $r_s = 0.08$ in the baseline calibration. There is some evidence to suggest that consumers believe stock returns to be considerably lower than this. Von Gaudecker and Wogrolly (2022) find that in a sample of American Life Panel respondents, the average perceived stock market return is 0.5% per annum. However their estimate comes with the caveat that the expectation is indirectly elicited. More detailed surveys on household's perceptions of asset returns would be of use to the literature.

[27]We do not control for income here as we did in one specification of 3.2, as the model-implied moments are measured across income states.

implied inflation forecast errors across quintiles of the distribution of risky-asset holdings. Despite not targeting the shape of the forecast error distribution in our calibration exercise - we targeted only the average forecast error - we find that the shape of the distribution of forecast errors in the model matches the data well. The model overstates inflation expectations at the bottom of the distribution and somewhat understates it at the top of the distribution of stock holdings. However, given that the data is merely a proxy for net financial wealth and the basic model has no serious treatment of liabilities or borrowing, we view this as encouraging for robustness specifications.

Figure 3.4: Inflation Forecast Errors by Stock Holdings Quintile, Model vs. MSC Data

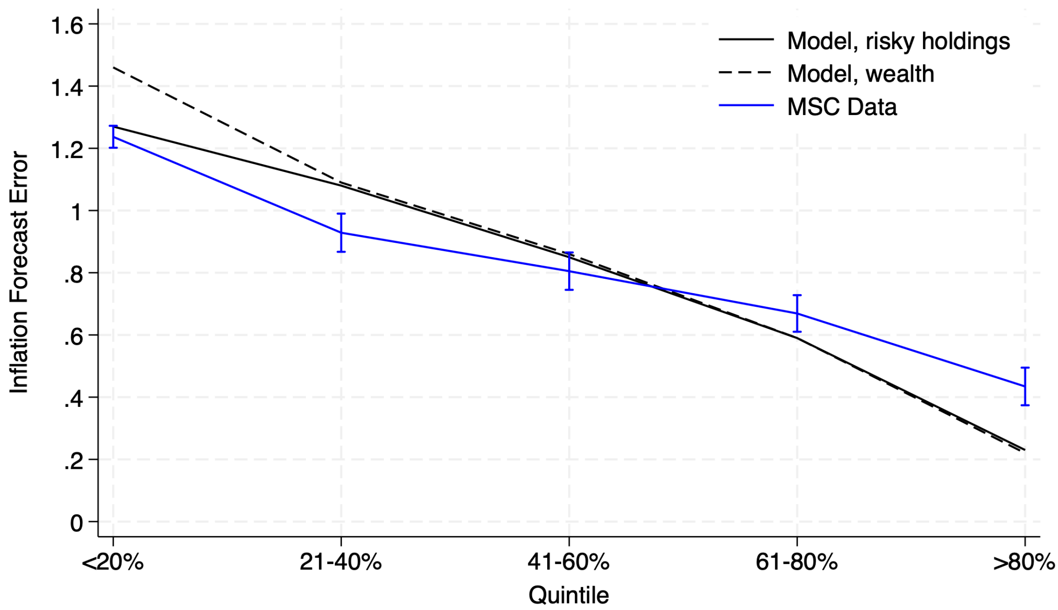

This graph displays two moments of the model. First, is the model implied forecast error for the inflation process - shown in solid black - across quintiles of (absolute) stock holdings. In dotted black, we show the model implied forecast error for the inflation process across quintiles of total wealth. This is compared to inflation forecast errors as measured in the Michigan Survey of Consumers, shown in blue. We compute the average inflation forecast error for respondents within a particular quintile of total stock holdings, after controlling for demographic covariates (education level, marital status, age, age squared, geographic region) and time fixed effects. This is the same data as shown in 3.1. In order to maintain consistency with the model moments, we report the estimates unconditional on income. The intervals denote the 95% confidence range. Standard errors are clustered at the individual level, and survey weights are employed.

Sensitivity to Inflation/Returns Covariance. The covariance structure of the inflation process and the risky-returns process amplifies the effects of robustness in our baseline calibration. When $\rho_{sp} < 0$, larger distortions to the inflation perception induce the agent to perceive lower returns to the risky-asset. This is true even if the agent doesn't seek robustness to the orthogonal component of the risky returns (dW_t^z) itself.

We show here the majority of the perceived drift for almost all agents[28] for the risky asset return is driven by the inflation distortion, rather that the distortion on the uncorrelated component. This is particularly true for those in the lowest income state y_1, for whom roughly 75% of the risky asset distortion is explained by robustness with respect to the inflation process.

In Figure 3.5, we illustrate how the perceived drift and the portfolio choice of the agent varies across different specifications of ρ_{sp}. The top row illustrates the perceived drift of the risky-asset returns and the bottom row illustrates portfolio shares. We vary ρ_{sp} across columns: the leftmost column illustrates the policy functions for a larger (absolute) correlation $\rho_{sp} = -0.3$, and the rightmost for a very small correlation $\rho_{sp} = -0.05$. We also show two calibrations for θ_z, for the sake of comparison. For the sake of legibility, we plot each set of results unconditional on y[29].

Figure 3.5: Proportion of Risky Asset Share Explained by Inflation Distortion

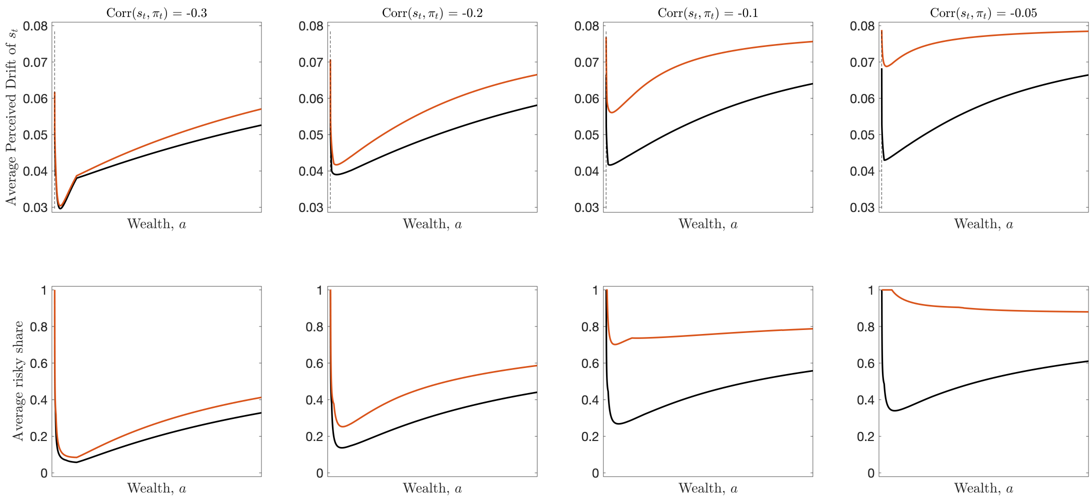

This graph displays the effect of changing ρ_{sp} on the perceived drift of the risky asset returns $\mathbb{E}_t\left[\frac{dS_t}{S_t}|a_t\right]$ (top row), and $\omega'(a)$, across wealth a. Both rows chart the average policy functions across states (i.e. not conditional on income state y). To calculate this, we use the ergodic distribution for income states. Across columns, we plot the results for $\rho_{sp} \in \{-0.3, -0.2, -0.1, -0.05\}$. The second column therefore aligns with the baseline calibration. In black, the standard calibration of $\theta_z = 0.16$ is used; the red lines compare this specification to one in which $\theta_z = 0$.

In black we show the baseline calibration for θ_z. In red, we set $\theta_z = 0$, isolating the effect on the policy functions of inflation model uncertainty only. The distance between these two lines is therefore the effect of seeking robustness to the orthogonal component of risky-returns.

The first thing to note is that for our baseline specification (shown in the second column), model uncertainty around the inflation process only is sufficient to generate large deviations from the expected

[28] Only those agents in the highest income-state but low wealth display perceived drifts of the risky asset returns that are (marginally) primarily explained by robustness towards the uncorrelated innovations.

[29] To do this, we take the average across income states, using the ergodic mean of the income process generated by y and Λ_y.

utility benchmark shown in Figure 3.2. Even for (unrealistically) low levels of correlation ($\rho_{sp} = -0.05$), uncertainty over the inflation process alone reduces the perceived risky return by approximately 1pp for low- wealth individuals. This is equivalent to 18% of the excess return provided by risky assets. For large covariance ($\rho_{sp} = -0.3$), model uncertainty over inflation alone is enough to eradicate the excess return provided by risky assets for the poorest groups.

Secondly, this exercise highlights that increasing the profile of the portfolio share is an artefact of the correlation between inflation. For low absolute covariance, the profile of the portfolio share across wealth becomes weakly declining when the consumer only seeks robustness towards inflation.

3.5 Welfare Implications of Robustness

Having outlined that the welfare implications of distorted portfolio allocations are substantial, we now focus on the welfare implications of seeking robustness in our baseline model.

The Hansen and Sargent (2001) specification we use raises two considerations when calculating welfare effects. First, under our specification of the concern for robustness, the value-function is non-homothetic. When $\theta_p > 0$ or $\theta_z > 0$ we therefore cannot easily calculate consumption-equivalent welfare measures as we would do in a standard CRRA setting (as in Lucas (1987), or later work such as Reis (2009)). Secondly, we must take a stand on the contribution of the "multiplier preferences" (i.e. terms related to h^z or h^p) in the value function. In Hansen and Sargent (2001), this utility specification is motivated as a modelling simplification of the "constraint" robust-control problem, where θ_i is to be interpreted as the implied Lagrangian multiplier on the specification error constraint[30]. In other discussions, such as that in Hansen and Sargent (2005), the problem is cast a two-player zero-sum game, in which an adversary chooses the distortion h_i such as to induce the worst-case stochastic process for the maximising agent. In either of these interpretations of the agent's HJB, the terms associated with robustness do not contribute to the "welfare" of the consumer: the terms associated with h_i are artefacts of the modelling environment only.

We therefore first assume that the "multiplier preferences" we employ are indeed nothing more than a modelling device, and do not contribute to the welfare of the agent directly. In order to calculate welfare differences in welfare between agents who seek robustness and those who do not, we assume that differences in welfare are generated only by their differing policy functions c^* and ω^* - i.e. their different consumption and risky-asset allocation profiles - and not by the terms associated with h_i.

[30] Using the multiplier specification rather than a constraint specification is both more simple to solve and easier to cast in a time-consistent form

Setting $\theta_z = \theta_p = 0$ has the added benefit of recovering a homothetic value function. We can therefore use the standard consumption-equivalent welfare metric for CRRA utility as employed in Lucas (1987). As instantaneous utility from consumption is CRRA, we calculate the difference between agents who have instantaneous consumption and risky-asset allocation pair $x = (c, \omega)$ rather than x' as[31]:

$$\lambda(a, y) = 1 - \left(\frac{V(a, y \mid x, \theta_p = 0, \theta_z = 0)}{V(a, y \mid x', \theta_p = 0, \theta_z = 0)} \right)^{\frac{1}{1-\gamma}}$$

Our welfare metric compares the instantaneous consumption an expected utility maximising agent with policy function x' would be prepared to sacrifice each period to instead employ the allocation x.

Back of the envelope. In our baseline calibration, reasonably modest pessimism over the drift of the inflation process has outsized effects for welfare. To see why this is the case, we first sketch a back-of-the-envelope illustration of the welfare cost for an individual who misses out on the excess returns provided by optimal investment in the risky-asset in an environment with no price-level uncertainty. The consumption-equivalent implications of setting the risky-asset share ω equal to ω_1 rather than ω_0 is approximately[32]:

$$\lambda_{\text{BOTE}} \approx 1 - \left[\frac{\rho - (1-\gamma) \cdot g(\omega_0)}{\rho - (1-\gamma) \cdot g(\omega_1)} \right]^{\frac{1}{1-\gamma}}$$

Where $g(\omega)$ is the certainty equivalent growth rate of consumption if the investor allocates ω percent of their portfolio to the risky-asset each period. Taking ω_0 to be the standard Merton (1971) market share($\omega_0 = \frac{r_s - r}{\gamma \sigma_s^2} = \frac{0.08-0.035}{2 \cdot (0.016)^2} = 88\%$) and ω_1 to be the population average risky share in our calibrated model ($\omega_1 = 30\%$), we find for our calibration that $\lambda = 8.2\%$. In other words, the representative investor allocating 30% of their wealth to the risky asset would be willing to relinquish 8.2% of their per-period consumption to switch to $\omega_0 = 88\%$. While this figure is substantial, it should not be surprising: investing ω_1 instead of ω_0 results in an effective per-period lower drift of $(\omega_0 - \omega_1)(r_s - r) \approx 2.6\%$ *each period*. Over the course of an infinitely lived lifetime these missed returns accumulate substantially[33].

[31]To see where this comes from, consider the λ that solves $u((1-\lambda) \cdot c(x)) = u(c(x'))$. When utility is CRRA, and the value function is $V(\cdot) = \mathbb{E} \int e^{-\rho t} u(c(\cdot)) dt$ we can write $u((1-\lambda)c) = (1-\lambda)^{(1-\gamma)} u(c)$ and $V(x) = (1-\lambda)^{(1-\gamma)} V(x)$. Use this in the first equation, and then solve for λ.

[32]Consider a CRRA investor who allocates $\omega \in [0,1]$ to a log-normal stock with drift r_s and volatility σ_s^2, and therefore allocates $1-\omega$ to a risk-free asset with per period return r. The certainty equivalent drift of c_t is given by $g(\omega) = r + \omega \cdot (r_s - r) - \frac{1}{2} \gamma \omega^2 \sigma_s^2$. This investor therefore has a consumption path given by $c_t = c_0 e^{g \cdot t}$. Their value function (lifetime welfare) is given by $V(g) = \frac{c_0^{1-\gamma}}{1-\gamma} \cdot \frac{1}{\rho - (1-\gamma) g}$. Therefore the consumption-equivalent welfare difference between an investor allocating ω_0 versus ω_1 to the risky-asset is given by $V(g(\omega_0)) = (1-\lambda)^{1-\gamma} V(g(\omega_1))$. Solving this for λ yields the back-of-the-envelope measure used in the main body of the text.

[33]This is reminiscent of Reis (2009), who finds that the effect of persistent shocks (there income) increases the welfare costs of uncertainty due to the effects on savings and investment behaviour.

3.5.1 Welfare in the Baseline Model

To calculate the impact of robustness, we calculate the optimal policy functions under the robust specification. Formally, these are $x^{\text{Rob}}(a, y) = (c^*(a, y), \omega^*(a, y))$, where:

$$(c^*(a, y), \omega^*(a, y)) = \arg\max_{c,\omega} V(a, y, c, \omega \mid \theta_p, \theta_z)$$

We then calculate the value function $V^{\text{Rob}}(a, y) = V(a, y \mid x^{\text{Rob}}, \theta_p = 0, \theta_z = 0)$. This is the welfare of an expected utility maximising agent, but one who takes the consumption c and ω to be fixed exogenous controls at x^{Rob}. We compare the welfare of this hypothetical agent to the welfare obtained by the standard expected utility maximising agent in our environment $V(a, y \mid \theta_p = 0, \theta_z = 0)$.

Figure 3.6 plots the value of consumption equivalent welfare λ across the state variables (a, y). For the sake of comparison, we also show the consumption equivalent welfare losses in a model in which the agent only seeks robustness towards the inflation process. The measure is reduced, but due to the correlation between risky returns and inflation only by approximately 25%.

Figure 3.6: Consumption-Equivalent Welfare, Robust vs. Non-Robust Agent

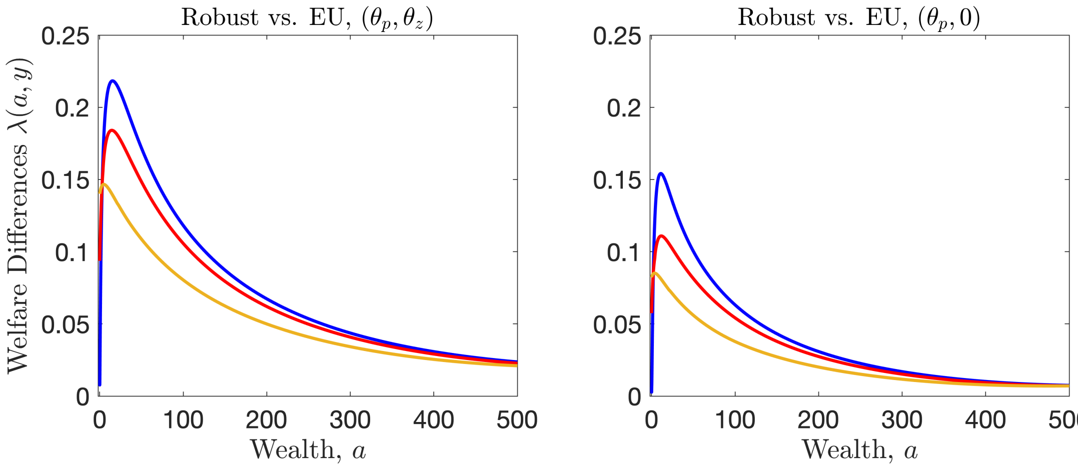

This graph plots consumption-equivalent welfare losses $\lambda(a, y)$ in the baseline model in the left hand panel. In the right hand panel it compares this to the consumption equivalent losses in a model in which the agent only seeks robustness towards the inflation process only. We use the same calibration for this second model, and only set $\theta_z = 0$ The blue line denotes the policy function of agents in the lowest income state, y_1. The red and yellow lines denote y_2 and y_3 respectively.

We find substantial variation across states. Those in the lowest income group, with low values of wealth would be willing to forgo around 20% of their consumption per-period to exist in a model in which they are certain of the data-generating process governing both the price-level and risky-asset returns. On the other hand, those in high wealth states would be willing to forgo roughly 11% of their per period consumption.

Caveats on the Magnitude of λ. Of course these numbers are much higher than similar exercises in the literature[34], and comes with considerable caveats.

First, is that this is the welfare of a dynastic household, rather than an individual with a finite life. Much of the welfare loss comes from the effect of lower risky-asset portfolio shares. Over long horizons, the utility benefit of such compounded can be very large[35].

Second, is our choice of utility function. We have employed a CRRA specification common in the macro literature, but the finance literature frequently uses higher values of γ for portfolio-choice models, often in the range of 4-6. This would serve to decrease our welfare calculation substantially. Moving to an Epstein-Zin style utility specification[36], while complicating the computation of the model, would allow a future researcher to separate the effect of inter-temporal substitution from risk aversion and refine this number further.

Third, we have also restricted the household from hedging inflation risk, by restricting the household to only long positions in the nominal bond. If the household is able to protect itself from nominal risk even partially, this does not only protect the returns on their wealth, but also serves to reduce the magnitude of their worst-case distortions. Finally, this is a partial equilibrium model, and therefore unlikely to wholly representative of general equilibrium outcomes. For all these reasons, our consumption welfare metric should be taken to be an upper-bound.

Average welfare. After obtaining $\lambda(a, y)$, we aggregate the consumption equivalent welfare measure using the ergodic distribution of agents obtained in the expected-utility maximising environment. We denote this as $g(a, y \mid \theta_p = 0, \theta_z = 0)$. We use this to calculate the population average welfare metric $\bar{\lambda}$, calculated as:

$$\bar{\lambda} = \sum_y \int_a \lambda(a, y) g(a, y) da$$

We find $\bar{\lambda} = 9.2\%$. This is comparable in magnitude to our back of the envelope calculation. The average agent employing the robust policy functions in this economy would be willing to forgo 9.2% of their consumption each period to employ the policy functions of the rational expected utility maximiser. To use the language of Rawls (2005), an agent behind the *veil of ignorance* would be willing to forgo 9.2% consumption each period to be born as an agent with rational expectations, rather than one with

[34] Reis (2009) find an upper bound of 5% consumption equivalent welfare losses caused by under-investment.

[35] In Viceira (2018), the welfare effects over long-horizons versus short horizons can be orders of magnitude larger.

[36] The most common continuous time version of these preferences is Stochastic Differential Utility, Duffie (1990).

expectations distorted by an adversarial agent (but expectations that are those reported in survey data).

Here we are using a social-welfare criterion with equal Pareto weights. We use the equal-weighting social welfare criterion without making a claim to its superiority or value, and employ it only for its immediacy. Other weightings would change this result. Were we to use weights that place greater weight on those with larger marginal utility (i.e. low values of (a, y)), the consumption equivalent welfare metric would be larger, as those with lower marginal welfare face greater distortions.

Again, we do not interpret this to be the *true* welfare cost of inflation forecast errors for households, but only seek to highlight: i) the welfare effect could be large, and; ii) that the interaction between forecasting errors and the investment channel is a possible driver of such welfare costs.

Welfare sensitivity to Robustness Finally, we explore the effects of the magnitude of inflation forecast errors on aggregate consumption equivalent welfare loss. Here we aim to illustrate the effect of the magnitude of forecast errors on our welfare metric. We therefore wish to emphasise the relative magnitudes of our consumption equivalent metric.

We perform the following exercise. We calibrate the robustness parameters to target aggregate inflation forecasts of varying sizes. We perform two calibrations. First, we hold θ_z fixed at 0, isolating the effect on welfare of seeking robustness towards inflation only. In a second set of calibrations, we target both the aggregate forecast error, as well as target an aggregate portfolio share of 30%, as in the baseline calibration. For both sets of calibrations we present the aggregate consumption equivalent welfare measure ($\bar{\lambda}$), as well as the standard deviation of the consumption equivalent welfare losses. These are calculated like so:

$$\sigma(\lambda) = \left[\sum_y \int_a \left(\lambda(a, y) - \bar{\lambda} \right)^2 \cdot g(a, y) \; da \right]^{\frac{1}{2}}$$

As before we aggregate the welfare losses using the $g(a, y)$ that solves the Kolmogorov Forward Equation under the expected utility problem. Therefore only the $\lambda(a, y)$s change as we change parameters θ_p and θ_z.

In Table 3.3 we present these results. The first set of rows presents calibrations where we adjust θ_p to target aggregate inflation expectations of 1% (the rational expectation) up to 4%. Note that for an aggregate expectation *below* the actual mean of inflation, we require a negative value of θ_p. Turning to the aggregate welfare losses: these are roughly linear in the average inflation forecast. Starting from the baseline of $\bar{\mathbb{E}}[\pi] = 0$, welfare losses increase from 7.2% for aggregate forecast error of 1 percentage

Table 3.3: Comparison of Aggregate Welfare Differences by Aggregate Forecast Error

	$\bar{\mathbb{E}}[\pi]$	0.01	0.02	0.025	0.03	0.035	0.04
Targeting Inflation	θ_p	-1.44	0	1.53	3.25	4.74	6.16
	θ_z	0	0	0	0	0	0
	$\bar{\lambda}$	0.0262	0	0.0229	0.0717	0.112	0.145
	$\sigma(\lambda)$	0.0397	0	0.0226	0.0566	0.0771	0.0886
Targeting both Inflation & Agg. Risky-share	θ_p	-2.03	0	1.202	2.79	4.88	7.51
	θ_z	0.619	0.416	0.299	0.148	-0.0340	-0.267
	$\bar{\lambda}$	0.0784	0.0819	0.0858	0.0929	0.109	0.128
	$\sigma(\lambda)$	0.0596	0.0619	0.0638	0.0673	0.0763	0.856

This table compares aggregate welfare differences under different policy targets and forecast error scenarios. The first set of rows holds θ_z fixed at 0, and varies θ_p to target various aggregate inflation expectations. The second set of rows allows both θ_p and θ_z to vary, targeting an aggregate portfolio share of 30%, as in the baseline model, but varying the targeted average inflation forecast $\bar{\mathbb{E}}[\pi]$ across calibrations. The average consumption-equivalent welfare loss ($\bar{\lambda}$) and the standard deviation of consumption equivalent welfare losses $\sigma(\lambda)$ are displayed. These are calculated as in the main text: all use the ergodic wealth distribution $g(a, y)$ of the expected utility maximising problem (i.e. where $\theta_p = \theta_z = 0$. Therefore only $\lambda(a, y)$ varies across calibrations.

point, to 14.5% when the aggregate forecast error is 2 percentage points. The effects on welfare are asymmetric: an average forecast error of 1% generates an average consumption equivalent welfare loss of 7.2%; whereas an average forecast error of -1% generates an average welfare loss of only 2.6%. A positive forecast error causes welfare losses as households pick sub-optimal portfolio shares, leading them to under-accumulate assets. A negative forecast error, on the other-hand, leads agents to over accumulate risky assets. The welfare loss is therefore generated by consumers exposing themselves to too much risk. The effect of under-accumulating wealth therefore has much greater welfare implications over the (infinite) course of the agent's life.

Secondly, the dispersion of welfare losses increases as the aggregate forecast error increases: i.e. welfare losses become more unequally distributed as the aggregate forecast error increases. This dispersion is driven by higher wealth individuals primarily. Higher perceived inflation affects the decisions of high wealth individuals the most, whose balance sheet are most exposed. While the welfare losses at the bottom of the distribution also increase, those at the top of the wealth distribution increase by more.

The second set of rows presents calibrations in which we vary the target aggregate inflation forecast, while keeping the aggregate portfolio share of risky holdings fixed at 30%, as in the baseline model. Here, the utility losses even when the average forecast error is only 0.5% is large (8.6%). This reflects the fact that an aggregate portfolio share of 30% is much lower than the optimal allocation implied by expected utility maximisation. When the aggregate expectation is lower than the rational expectation, this mitigates the effect of sub-optimal allocations somewhat. For large aggregate errors, the welfare cost is of roughly the same magnitude as in the calibration in which agent's only seek robustness to the inflation process. This reflects that when the aggregate forecast error is high and $\rho_{sp} = -0.2$, risky-asset allocation is already driven towards very low values. Indeed to maintain an aggregate risky-share of 30% for such a calibration, a negative value of θ_z is required to compensate for the effect generated by inflation model uncertainty.

3.6 Portfolio Adjustment Costs

We now adjust our baseline model to incorporate the possibility of fixed portfolio adjustment costs. We consider this extension for two reasons. First, it is a realistic feature of the real world that has outsized effects: Vissing-Jørgensen (2002) find that very small adjustment costs help explain half of all non-participation. Second, and most crucially, fixed costs and concerns regarding model misspecification disproportionately affect lower-income households. Low-income households that expect higher inflation (and lower returns) will assess the benefit of paying the fixed cost to be lower *ceteris paribus*, even when the cost of a suboptimal portfolio has the most detrimental effects for these agents.

Households have access to a liquid asset (B_t) and an illiquid asset (S_t). As before, the returns process for both assets is nominal. The returns processes and the process for the price-level for these two assets is as in Section 3.3. Agents continue to seek robustness towards the process governing innovations to the (log) price level, as well as innovations in the uncorrelated component of risky-asset demand. The key difference is that now deposits and withdrawals into the risky-asset account can only be made after payments of a transaction cost κ. The risky-asset provides higher returns, but is illiquid. We assume that instantaneous consumption can only be purchased using the stock of the liquid, but low return, asset. We can therefore write the household's problem as:

$$V(a,b,y) = \max_{\{c_t\},\tau} \min_{q} \int_0^\tau e^{-\rho t} \left[\int u(c_t)\, dq_t \right] dt + \rho \int_0^\tau e^{-\rho t} \left(\log\left(\frac{dq_t}{dq_t^0}\right) dq_t \right) dt + \\ + e^{-\rho\tau} \int V_k^*(a_\tau + b_\tau, y_\tau)\, dq_t$$

$$\text{where:} \quad V^*(a+b,y) \equiv \max_{a',b'} V(a',b',y) \text{ s.t. } a'+b' = a+b-\kappa$$

$$ds_t = (r_s - \bar{\pi} - \frac{1}{2}\sigma_p^2 + \sigma_s\sigma_p\rho_{sp})s_t dt + (\sigma_s\rho_{sp} - \sigma_p)s_t dW_t^p + \sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}} s_t dW_t^z$$

$$db_t = \left[\left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) b_t + y_t - c_t\right] dt - \sigma_p b_t dW_t^p$$

The inclusion of a fixed-adjustment cost means the household must choose at what "time" to adjust their portfolio. This therefore introduces a stopping time problem to the households decision. Problems of this form can be written in recursive form using a "HJB quasi-variational inequality" (HJBQVI)[37] as follows:

$$
\begin{aligned}
0 = \min\{\rho \cdot V(s,b,y) - \max_{c}\min_{h^p,h^z} \mathcal{H}(s,b,y)\ ,\quad V(s,b,y) - V^*(s+b,y)\} \\
\mathcal{H}(s,b,y) \equiv u(c) + \frac{1}{2}\left(\frac{1}{\theta_z}(h^z)^2 + \frac{1}{\theta_p}(h^p)^2\right) + \\
+ \left[(r_s - \bar{\pi} - \frac{1}{2}\sigma_p^2 + \sigma_s\sigma_p\rho_{sp})\cdot s + (\sigma_s\rho_{sp} - \sigma_p)\cdot s\cdot h^p + \sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}}\cdot s\cdot h^z\right]\cdot \partial_s V(s,b,y) + \\
+ \left[\left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right)\cdot b + y - c - \sigma_p b\cdot h^p\right]\cdot \partial_b V(s,b,y) + \\
+ \frac{1}{2}(\sigma_s^2 + \sigma_p^2 - 2\sigma_s\sigma_p\rho_{sp})\cdot s^2\cdot \partial_{ss}V(s,b,y) + \\
+ (-\sigma_p\cdot(\rho_{sp}\sigma_s - \sigma_p)\cdot s\cdot b)\cdot \partial_{sb}V(s,b,y) + \\
+ \frac{1}{2}\sigma_p^2\cdot b^2\cdot \partial_{bb}V(s,b,y) + \\
+ \sum_{y\neq y'}\lambda_{y\to y'}(V(s,b,y') - V(s,b,y))
\end{aligned}
\tag{3.3}
$$

Note that as we two endogenous state variables, the cross-correlation between s and b is non-zero if $\rho_{sp} \neq 0$. We account for this by including the term $(-\sigma_p\cdot(\rho_{sp}\sigma_s - \sigma_p)\cdot s\cdot b)\cdot \partial_{sb}V(s,b,y)$ in the HJB. In the limiting case of $\kappa \to 0$, we recover the baseline model of Section 3.3.

We then discretize the problem and cast it as a Linear Complementarity Problem (LCP), and then solve this using a finite-difference upwind scheme. For more detail on this, see Appendix C.6.

Calibration For this extension, we use the calibration as used in the baseline model of section 3.3, for common parameters. We therefore do not recalibrate this model to hit the same moments of the data. Instead, the objectives are to: i) examine how fixed adjustment costs alter the model in comparison to the baseline, and; ii) demonstrate that concerns regarding misspecification exacerbate the welfare losses associated with adjustment costs.

We set the cost of adjustment κ to be equal to be 1% of the ergodic mean of the income process ($\kappa = 0.0601$). In 2015, this would be equivalent to approximately 35$ per month[38]. Vissing-Jørgensen (2002) find very low adjustment costs ($50, or 0.2% of average disposable income per capita in 2002) can explain roughly half of non-participation in the US. On the other-hand, Khorunzhina (2013) estimates there to be evidence of large adjustment costs (4-6% of annual income) using Panel Survey of Income

[37]See Stokey (2009) for a canonical reference for the use of equations of this form in economics. An earlier paper, similar in spirit ours is Reis (2006).

[38]Disposable income per capita was roughly $42,000 in the U.S. in 2015 (BEA Data, taken from FRED).

Dynamics data. We therefore pick a conservative value within this range.

3.6.1 Perceived Inflation

We first show how perceived inflation across agents varies across asset states. Figure 3.7 displays the perceived drift of the inflation process for each of three income states in the model. Within the fixed-

Figure 3.7: Perceived Inflation Across Income States, Fixed Adjustment Costs

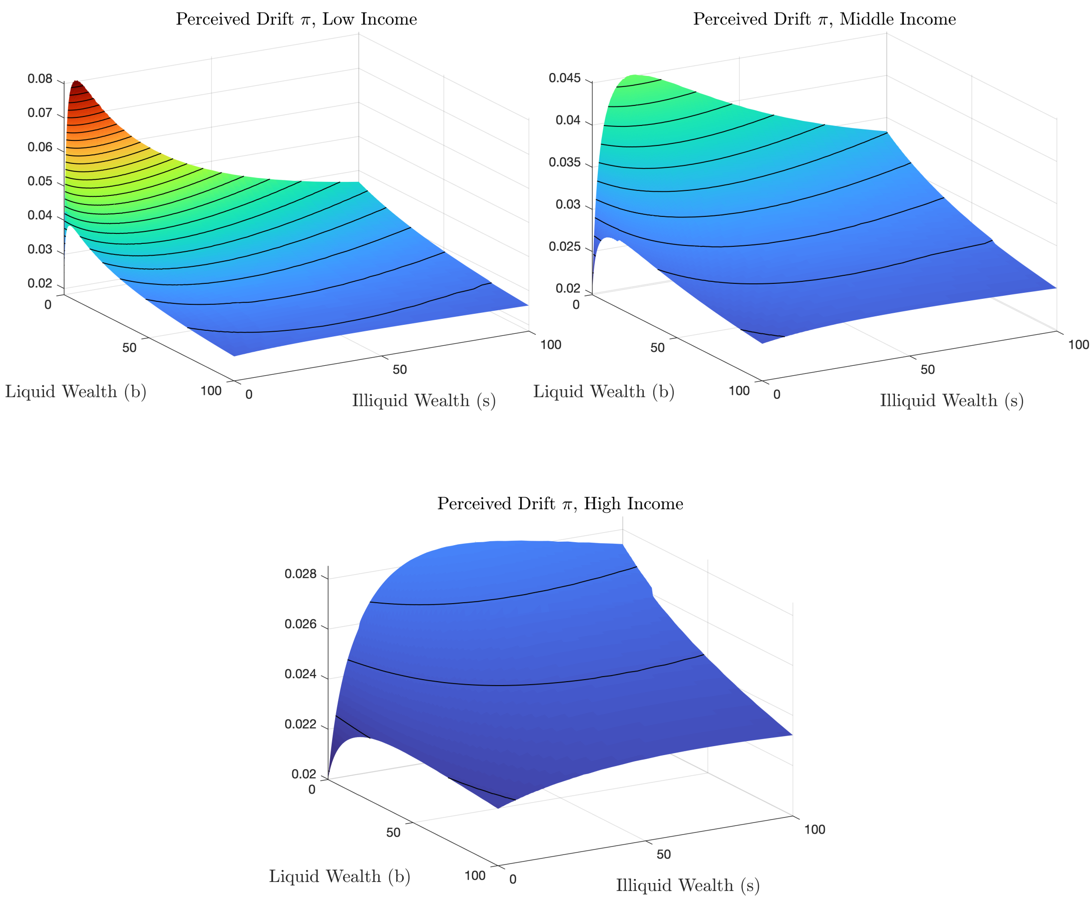

This graph displays the model implied perceived inflation across the three income states. The perceived drift of the inflation process is calculated as $\bar{\pi} + \sigma_p \cdot h^{p*}(s, b, y)$. The colouration denotes the z-axis values: the perceived drift of the inflation process. The black lines are iso-lines. The z-axis starts at $\bar{\pi} = 2\%$ but for income state are show with varying upper limits for the sake of legibility. To the same end, we have rotated the axes such that the liquid wealth (b) axis runs from the background to the foreground.

adjustment costs model perceived inflation shows a similar pattern across wealth and income as in the model with no adjustment costs. The perceived mean of the inflation process is highest for those agents with low wealth (s, b) and low income y. Holding either liquid wealth (b) or illiquid wealth fixed (s), the perceived drift of the inflation process displays a pronounced hump shape for all three income states. As in the no adjustment cost model, the peak of this hump as a function of (b, s) increases as the income state increases.

Note that the iso-lines - which here denote combinations of the state-vector with the same perceived drift of inflation - are not semi-circular, but rather stretched for lower values of b. Between asset classes, the difference between the expectation of the robust-agent and the expected utility maximiser is therefore larger for those with low liquid asset holdings (b) relative to those with low illiquid asset holdings (s). The intuition for this is simple: as consumption can only be paid for from the liquid asset account, the marginal value to the agent of a unit of b relative to a unit of s is higher. Therefore at low levels of liquidity (high risky-asset portfolio shares), the agent perceives the mean of inflation to be higher *ceteris paribus*: a given inflationary shock will have a larger effect on consumption than at high values of b, and increases the likelihood the agent will need to re-optimise their portfolio and therefore pay the adjustment cost.

Relative to the no-adjustment cost model the peak of the perceived inflation process within an income group is substantially higher. For the lowest income group in the baseline model, the peak expectation wedge is $\Delta_p^E \approx 1.75\%$. In a model with the same calibration and modest transaction costs (1% of the ergodic mean of annual income) the peak expectation wedge for the same income state is $\Delta_p^E \approx 8\%$. Within the model, transaction costs can have outsized effects on inflation expectations. For low income and wealth states paying the transaction cost is particularly harmful for the agent due to the curvature of utility. Therefore the worst-case drift of inflation is very high. In particular, for those agents with low liquidity *and* low wealth (i.e. high ω but low $s + b$) this effect is most pronounced.

3.6.2 Robustness and Adjustment

How does the concern for robustness interact with the adjustment cost? In Figure 3.8 we display the action/inaction regions of agents across states. The middle row displays our baseline calibration. The upper and lower rows show low and high θ_p specifications respectively, for the sake of comparison. White regions are areas in which the agent does not choose to pay the transaction cost and adjust their portfolio. Red regions denote those in which agents adjust from illiquid to liquid assets; in green regions, vice-versa.

There are three effects of robustness to highlight here. First, is that greater concern for robustness - and all else equal, greater forecast errors - moves the inaction region towards portfolios with a lower proportion of risky assets and higher proportion of liquid assets. Graphically, the white region moves towards the south-west. This reflects the behaviour of agents in the baseline model. For our baseline calibration (the middle row), only those agents with very high levels of liquid asset holdings find it optimal to increase the risky-share of their portfolio (and even then only mildly).

Secondly, the boundary flattens. Graphically, the boundary line between the white and the coloured

Figure 3.8: Adjustment Regions For Each Income State, Varying θ_p

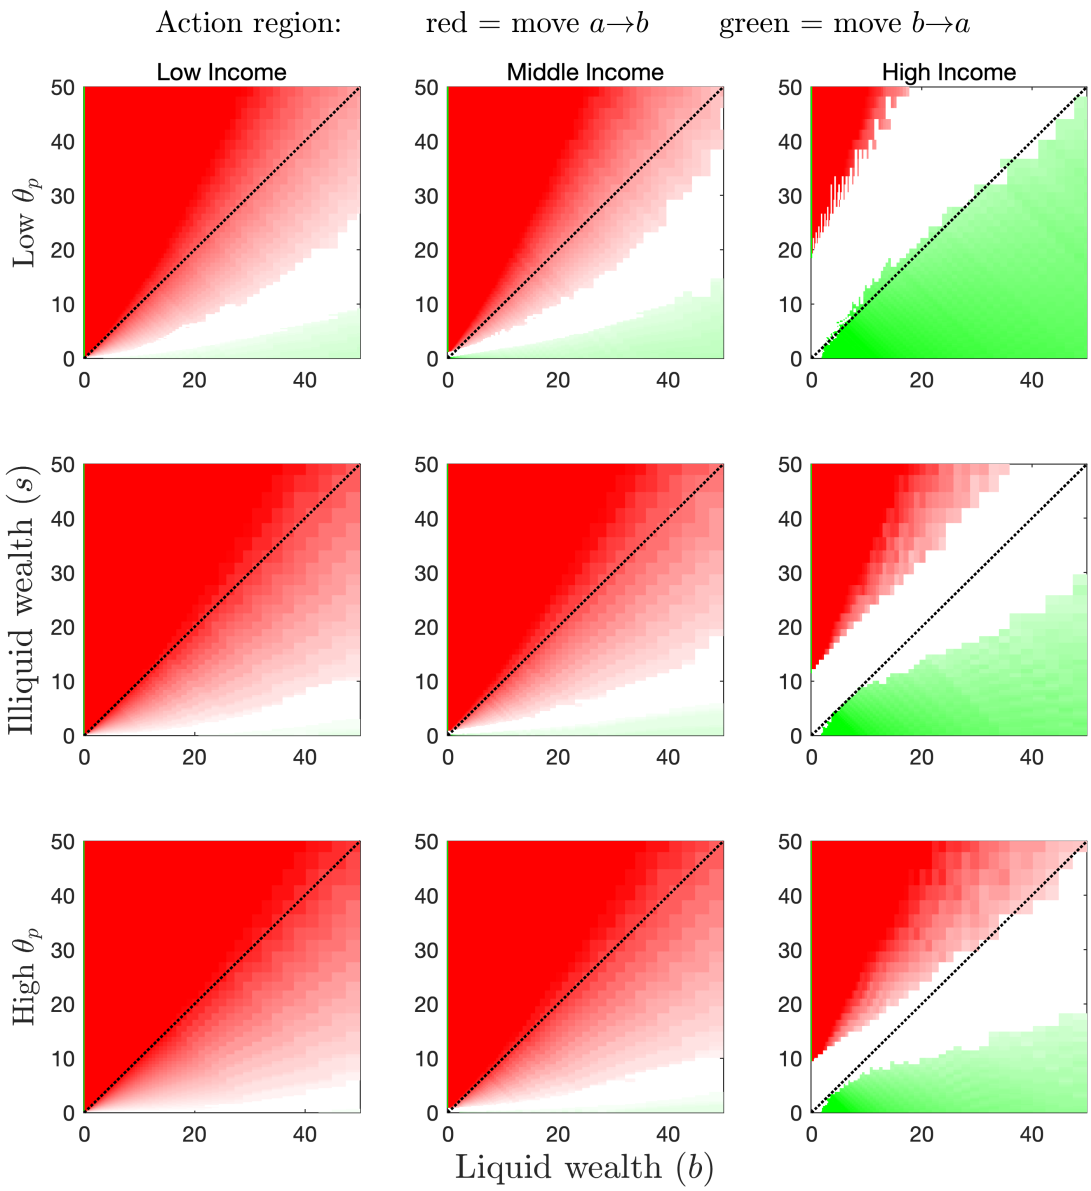

The top row displays this for a low concern for robustness specification ($\theta_p = 0.1 \times \theta_p^{\text{calib}}$). The middle row displays the inaction region for the baseline calibration ($\theta_p^{\text{calib}} = 2.6332$). The bottom row for a high concern for robustness specification ($\theta_p = 1.5 \times \theta_p^{\text{calib}}$). This graph displays the action/inaction region across pairs (b, s) for each income state y. The region in which agents find no-action optimal (i.e. to not adjust their portfolio) is shown in white. Action regions are coloured. Red areas denote the region in which agents opt to pay the transaction cost and drawdown assets from their liquid account. The green areas denote regions in which agents choose to pay the fee and purchase in the risky-asset. The intensity of the colouration varies with the proportion of assets they choose to move: the highest intensity regions are ones in which the agent chooses to move their entire portfolio. The dotted black line denotes the line $s = b$, where the portfolio share of risky assets is 50%. This is shown as a visual aide.

region has a shallower gradient as concern for robustness increases. To understand why the slope of the inaction boundary decreases, consider the relative effect of a change in θ_p on the perceived drift of each asset. For s, increasing θ_p changes the perceived drift in proportion to $(\sigma_s \rho_{sp} - \sigma_p) < 0$; whereas for b it is in proportion to $-\sigma_p < 0$. In our calibration, as $\rho_{sp} < 0$, a given change in θ_p has a larger effect on the drift of s than b. Therefore the marginal benefit of holding an extra unit of wealth in the risky asset falls quicker than holding it in b. However at the boundary we know that the smooth-pasting condition must hold (i.e. that $V_s(s, b, y) = V_b(s, b, y) = V^*_{s+b}(s + b - \kappa, y)$). Therefore the equality $V_s = V_b$ is met at lower values of s for large b, when θ_p increases. Intuitively, robust households become less and less willing to let their illiquid asset share increase as their liquid buffer grows when concern for robustness increases, as the worst-case inflation process hurts s disproportionately.

Finally, Figure 3.8 also shows how the inaction region narrows. Greater concern for robustness makes the worst-case drift of inflation greater, and therefore more costly to the agent. The agent therefore finds being further away from the optimal portfolio mix is now more costly. The marginal gain from adjusting therefore increases, and as the adjustment cost is fixed at κ the option value of waiting (not-adjusting) decreases. Households are willing to tolerate smaller deviations from their optimal portfolio mix before paying the fixed cost κ as concern for robustness increases. This serves to increase the *frequency* of trades[39]. The corollary of this is concern for robustness - larger forecasting errors - increases the frequency with which households pay the fixed cost to adjust[40].

3.6.3 Welfare with Adjustment Costs

As a final exercise, we consider the welfare cost in the model with portfolio adjustment costs. Specifically we are interested in how the welfare of the agent in the model with adjustment costs, relative to the frictionless model, changes as the concern for model misspecification increases. We construct consumption equivalent welfare measures in a manner similar to Section 3.5. We compare the value functions across four different specifications: $V^{\text{EU}}(\kappa)$, $V^{\text{rob}}(\kappa)$, $V^{\text{EU}}(\kappa')$ and $V^{\text{rob}}(\kappa')$, where $\kappa' > \kappa$. We calculate V^{EU} as equation (3.3) solved under $\theta_p = \theta_z = 0$. In order to maintain the homotheticity of the value functions, and remove the effect of worst-case distortion on the utility function[41], we remove the robustness terms from V^{rob}. We calculate this like so:

$$V^{\text{rob}}(\kappa) = V(s, b, y|\kappa, \theta_p, \theta_z) - \frac{1}{\rho} \sum_y \int_s \int_b \frac{1}{2} \left(\frac{1}{\theta_z} (h^{z*}(s, b, y))^2 + \frac{1}{\theta_p} (h^{p*}(s, b, y))^2 \right) \cdot g'(s, b, y) \, ds \, db$$

Where h^{p*} and h^{z*} are the worst-case distortion policy functions associated with $V(s, b, y|\kappa, \theta_p, \theta_z)$ and $g'(s, b, y)$ is the solution to the KFE equation evaluated *with* the worst-case drift distortions: this represents the agent's perceived expectation, rather than the *actual* equation of motion. Simply put:

[39] That agent's with more pessimistic expectations trade more frequently (conditional on market participation) is something which could possibly be empirically verified. The data-set is not off-the-shelf however, and so is left to future researchers.

[40] This would also be a feature of the homothetic-robustness specification. There, an increase in the concern for robustness is akin to an increase risk-aversion (when utility is CRRA). As an increase in risk-aversion has a similar effect on compressing the inaction region, so too would an increase in concern for robustness.

[41] as per the discussion in Section 3.5.

as the value function is additive in the worst-case drifts, by removing the expected 'utility' flow that stems from them we restore homotheticity to the value function in the non-robust case. Were the worst-case distortion contribution not-additive, this would not be possible. This is also consistent with the consumption equivalent welfare metric derived in Section 3.5.

We can still employ the homothetic consumption equivalent welfare metric of Section 3.5 as the Bellman operator of (3.3) remains linear. Therefore even under the HJBQVI specification, the construction of the consumption equivalent welfare measure remains unchanged. The objective function is the present value of consumption utility. Even if consumption is permanently rescaled by $(1-\lambda)$, this is conditional on the state vector $c(s,b,y)$. Therefore how the state variables move between trading periods is irrelevant, as multiplying the path of consumption by some factor also multiplies the path of consumption that already reflects the jumps. In fact, this is the same logic that allows us to remain unconcerned in the face of the Poisson process over income: the reasoning is the same. Were the instantaneous utility function to also directly be a value of the state variables - say, liquid asset holdings b - we could not construct a consumption equivalent welfare measure in this manner. We therefore compute:

$$\bar{\lambda}^{\mathrm{i}} = \sum_{y} \int_{s} \int_{b} \lambda^{i}(s,b,y) \cdot g^{\mathrm{EU}}(s,b,y|\kappa) \, ds \, db \;, \quad \text{where} \quad \lambda^{i}(s,b,y) = 1 - \left(\frac{V^{i}(\kappa')}{V^{i}(\kappa)} \right)^{\frac{1}{1-\gamma}} \quad \text{for } \; i \in \{\mathrm{EU},\ \mathrm{rob}\}$$

In words, we compare an agent in the model of low fixed portfolio adjustment costs to a model in which fixed portfolio adjustment costs are higher. λ represents the instantaneous consumption the κ-model agent would require in compensation each period, in order to live in the κ' model. We then aggregate this measure, to compute the average consumption-equivalent welfare metric. We employ the ergodic *realised* distribution of the expected utility low adjustment-cost model κ for both, as we did in Section 3.5, for the sake of fair comparison.

We take $\kappa = 1\%$ as before, and compare it with a value at the upper range of the estimates of transaction costs (Khorunzhina (2013)). We therefore set $\kappa' = 5\%$ of the ergodic mean of the income process. We use the calibrated robustness parameters as before. We find $\bar{\lambda}^{\mathrm{EU}} = 0.66\%$, and $\bar{\lambda}^{\mathrm{rob}} = 1.26\%$. In the expected-utility model, moving from a low transaction cost to a high transaction cost is worth 0.66% of consumption per-period. However, in the robust model - i.e. the model in which agents have pessimistic forecasts - the cost of the same change in the transaction cost rises to 1.26%. The intuition of this stems from the discussion in the previous section. As the robust agent has pessimistic views over the stochastic process governing inflation and risky-returns, the risk these processes present are more costly to the agent [42]. Being away from their desired portfolio share is therefore more costly, and so they are willing to pay the adjustment cost after smaller shocks - i.e. more frequently. They therefore pay the adjustment cost more often, and which increases the presence of the friction to the agent.

Concern for model misspecification generates pessimistic expectations - in turn, this amplifies the welfare cost of portfolio cost frictions. This has an interesting policy implication: it tells us that policies which reduce transaction frictions are *more valuable* in economies where agents distrust the data-generating process more. Alternatively, it highlights that removing transaction frictions in a model in which households have pessimistic expectations has greater effect than one in which their expectations are rational.

3.7 Conclusion

Our analysis makes three central contributions. First, we document a new empirical fact—households that hold larger equity positions systematically make smaller, though still positive, inflation-forecast errors. Second, we show that a portfolio-choice model with non-homothetic robust preferences can replicate this pattern: poorer (high-marginal-utility) agents endogenously distort the mean of the inflation process upward, perceive a lower equity premium, and therefore allocate far less to risky assets than an expected-utility agent would. Third, we quantify the stakes. Calibrated to U.S.survey data, the model implies that misperceptions only of the order of one percentage point can dramatically reduce optimal risky asset shares for the poorest agents by roughly 90%, and generate large consumption-equivalent welfare losses. When we embed a modest fixed portfolio-adjustment cost, the "inaction region" shrinks as concern for robustness increases. This leads the welfare losses presented by moving from a low fixed-cost to a high fixed-cost regime to roughly double: transaction frictions matter far more in economies where agents distrust the data-generating process.

On the empirical side, these results highlight that heterogeneity in *beliefs*, rather than in preferences or participation costs alone, can account for much of the observed dispersion in risky-asset holdings. On the theoretical side, they demonstrate that abandoning the usual homothetic robustness shortcut is essential to accurately capture the cross-sectional profile of expectations. Only the fully non-homothetic formulation delivers the declining forecast error profile the data demand. Moreover, by calibrating robustness parameters directly to survey moments instead of relying on detection-error probabilities, we offer a transparent, data-driven discipline for future quantitative work with ambiguity or model uncertainty.

[42] This is as in Maenhout (2004): In the homothetic robustness case, ambiguity is homomorphic to greater risk-aversion.

Several promising extensions remain. A natural next step is to endogenise learning about the true data-generating process, as in Hansen and Sargent (2011). Second, introducing additional assets (e.g. inflation-indexed bonds, housing, or long-duration nominal debt) would let us study further how incomplete markets shape the demand for inflation protection and amplify cross-sectional welfare effects. Of course, most pertinently, embedding our household block inside a general-equilibrium model would reveal how the *distribution* of belief-driven portfolio distortions feed back into asset prices, the equity premium, and optimal monetary or financial-stability policy. Incorporating more representative expectations into heterogeneous-agent general equilibrium models is currently a fertile area of research[43]. Finally, richer micro-evidence - especially direct surveys of expected real returns — would help separate ambiguity over inflation from ambiguity over asset fundamentals, sharpening both identification and policy design.

[43] See Moll (2024) for a recent call-to-arms.

C Supplementary Material for Chapter 3

C.1 Inflation Forecast Errors and Stock Holdings

In this section, we provide supporting evidence to Section 3.2. Table C.1 illustrates the summary statistics for our data. This subset contains only those observations for which we have both inflation forecast errors, and reported values of stock holdings. The data has been trimmed at 5-95% level, by date. The units are deflated by the BLS CPI, and so are denominated in 2015$.

Table C.1: Summary Statistics for MSC Stock-Inflation FE Sample, by Quintile

	Stock Holdings Quintile						
	Non-participant	<20%	21-40%	41-60%	61-80%	>80%	Total
Stock Holdings (2015$)							
Mean	0	7738	37263	96931	237429	569052	94608
Median	0	6136	15122	36990	97540	244289	179246
St. Dev	0	6565	35520	91707	219597	517930	10100
N	46304	13414	14715	15698	16881	11130	118142
Inflation Forecast Error							
Mean	1.35	.976	.888	.768	.583	.411	.997
Median	3.91	3.76	3.71	3.63	3.49	3.32	3.75
St. Dev	.857	.508	.481	.386	.269	.23	.559
N	46304	13414	14715	15698	16881	11130	118142

This table provides summary statistics for the variables of interest in Section 3.2. We show the reported value of stock holdings (in 2015$) and the inflation forecast error, broken down across stock holdings quintile. We only report summary statistics for the sub-set of the data for which we have stock and inflation forecast error observations (post trimming). Summary statistics are calculated using survey weights.

In Table C.2, we show the regression supporting 3.1. Column 1) only includes time fixed effects, Column 2) includes demographic controls[44], and Column 3) additionally controls for income. Our income control groups households into 11 bins.

The declining profile of inflation forecast errors across respondents grouped by stock-holdings holds true across all three specifications. The inclusion of a control for income approximately halves the effect of being in the top quintile compared to non-participants; however, even when conditioned on income, inflation forecast errors exhibit a marked and monotonic negative gradient concerning stock holdings.

[44] For a full list of the demographic controls we use in the MSC, as well as their precise definitions, please see A.1.

Table C.2: Regression Supporting Figure 3.1

	(1)	(2)	(3)
1st quantile	-0.387***	-0.237***	-0.166***
	(-10.79)	(-6.42)	(-4.41)
2nd quantile	-0.488***	-0.376***	-0.221***
	(-13.80)	(-10.30)	(-5.83)
3rd quantile	-0.599***	-0.499***	-0.300***
	(-17.64)	(-14.11)	(-7.93)
4th quantile	-0.769***	-0.679***	-0.424***
	(-23.77)	(-19.58)	(-11.19)
5th quantile	-0.993***	-0.894***	-0.585***
	(-27.34)	(-22.82)	(-13.46)
Obs	118,142	116,898	112,187
Adj. R^2	0.238	0.242	0.246
Mean dep var	.971	.969	.956
First obs	1990m1	1990m1	1990m1
Demo. Controls		✓	✓
Income Control			✓

This table details the regression output supporting Figure 3.1, where the inflation forecast error is regressed on stock-holdings group in the MSC. The coefficients display marginal effects of belonging to a given stock holdings quantile, measured within a survey, relative to the group of stock market non-participants. Column (1) shows the basic regression, where only time-fixed effects are included, corresponding to the black line. Column (2) also includes the vector of demographic controls, corresponding to the blue line. Column (3) also includes (real) income group fixed effects, corresponding to the red line. In brackets, t-statistics are displayed. These are calculated using standard errors clustered at the individual level. One, two, or three stars denote significance at the 10%, 5% and 1% level respectively.

C.2 Hansen-Sargent Style Robustness

We first present a short recap of a Hansen and Sargent (2001) style agent, who fears model misspecification and therefore seeks robustness. At time 0 let an agent's problem in state x with control c be of the form:

$$V(x_0) = \max_{\{c_t\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho_{sp} t} u(c_t, x_t) dt \right]$$
$$\text{s.t. } dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dW_t$$

Here, x is stochastic, and follows $dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dW_t$ with a given initial condition x_0. Here, dW_t is a standard Brownian motion.

The consumer seeks Hansen and Sargent (2001) style robustness across the stochastic environment, and worry that their understanding of the world is misspecified. They therefore evaluate other competing models of the world by evaluating the statistical similarity of their reference model to alternatives. Models which are far away from the reference model are penalised more than those which are observationally similar to the agent's reference. The distance between the reference and alternative models is captured by the relative entropy (also known as Kullback-Leibler divergence). In essence, this is the expected log-likelihood ratio of alternative models, relative to the reference model. Let q^0 be the probability measure defined by the Brownian motion dW_t, and q be the probability measure of an alternative model which is absolutely continuous with respect to q_0[45]. Then, the distance between this model and alternatives is given by:

$$\mathcal{R}(q) = \rho_{sp} \int_0^\infty e^{-\rho_{sp} t} \left(\log \left(\frac{dq_t}{dq_t^0} \right) dq_t \right) dt$$

The insight of Hansen and Sargent (2001) is that we can model a decision-maker in such an environment as playing a zero-sum game with an adversary who systematically attempts to reduce the welfare of the agent by distorting the probability distribution of the underlying stochastic state. This problem is then written as:

$$V(x_0) = \max_{\{c_t\}_{t\geq 0}} \min_q \int_0^\infty e^{-\rho_{sp} t} \left[\int u(c_t, x_t) dq_t \right] dt + \frac{1}{\theta} \mathcal{R}(q)$$
$$\text{s.t. } dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)dW_t$$

[45] If q is not absolutely continuous with respect to q_0, this implies that there is a zero-probability event which the two models disagree on. These are easy to detect statistically.

Where θ captures preferences for robustness. The case of $\theta \to 0$ nests the expected utility-maximising consumer. As θ grows, so does the fear of misspecification. Letting h be a progressively measurable process, an application of Girsanov's Theorem let the relative entropy be written as:

$$\mathcal{R}(q) = \frac{1}{2}\mathbb{E}_0\left[\int_0^\infty e^{-\rho_{sp}t}|h_t|^2\right]$$

A further application of the theorem allows us to view the alternative model q as induced by a perturbation of the reference model. That is, $\hat{W}_t = W_t + \int_0^t h_s ds$. Given this, we can write the consumer's problem as:

$$V(x_0) = \max_{\{c_t\}_{t\geq 0}} \min_{\{h_t\}_{t\geq 0}} \mathbb{E}_0\left[\int_0^\infty e^{-\rho_{sp}t}\left(u(c_t, x_t) + \frac{1}{2\theta}|h_t|^2\right)\right]$$
$$\text{s.t. } dx_t = \mu(c_t, x_t)dt + \sigma(c_t, x_t)\left(h_t dt + d\hat{W}_t\right)$$

This problem can also be written recursively. The Hamilton-Jacobi-Bellman equation associated with this problem is:

$$V(x) = \max_c \min_h u(c, x) - \frac{1}{2\theta}h^2 + \left(\mu(c, x) + \sigma(c, x)\cdot h\right)\cdot \partial_x V(x) + \frac{1}{2}\sigma(c, x)^2 \partial_{xx} V(x)$$

Where ∂_k represents the derivative operator with respect to variable k. Therefore the decision maker acts as if the drift of the state x were $\left(\mu(c, x) + \sigma(c, x)\cdot h\right)$ instead of simply $\mu(c, x)$. Doing this means the agent incurs a utility cost of $\frac{1}{2\theta}h^2$. This is the value function as written in Section 3.2.

C.3 Derivation of the Real-Wealth Process

Throughout the paper, time is continuous and uncertainty is generated by two standard Brownian motions W^p, W^s on a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$ with instantaneous correlation $\mathbb{E}[dW_t^p\, dW_t^s] = \rho_{sp}\, dt$ where $-1 < \rho_{sp} < 1$.

The nominal dynamics of the price-level, risky-asset and safe asset values are given by:

$$d\log(P_t) = \bar{\pi}dt + \sigma_p dW_t^p$$
$$dS_t = r_s S_t\, dt + \sigma_s S_t\, dW_t^s$$
$$dB_t = r\, B_t\, dt$$

Where we have used Itô's Lemma for the process on the price level:

$$d\log(X_t) = \mu(X_t)dt + \sigma(X_t)dW_t \implies dX_t = \left(\mu(X_t) + \frac{1}{2}\sigma(X_t)^2\right) \cdot X_t dt + \sigma(X_t)X_t dW_t$$

We can therefore rewrite the process for $\log(P_t)dt$ as $dP_t = (\bar{\pi} + \frac{1}{2}\sigma_p^2)P_t\ dt + \sigma_p P_t\ dW_t^p$. The term $\frac{1}{2}\sigma_p^2$ is sometimes referred to as an Itô drift distortion.

Then, let $A_t = S_t + B_t$ be total nominal wealth and Y_t exogenous nominal income. With nominal consumption $P_t c_t$, the budget constraint is: $dA_t = (Y_t + rA_t + (r_s - r)S_t - P_t c_t)\,dt + \sigma_s S_t\, dW_t^s$. Next, define the real value of variable X_t as $x_t \equiv \frac{X_t}{P_t}$. For any ratio $x_t = X_t/P_t$, Itô's quotient rule gives:

$$dx_t = \frac{1}{P_t}dX_t - \frac{X_t}{P_t^2}dP_t + \frac{1}{P_t^2}d[X,P]_t.$$

From the diffusion parts of the stochastic processes, write the quadratic covariation of S_t and P_t as:

$$\begin{aligned} d[S,P]_t &= (\sigma_s S_t)(\sigma_p P_t)\,\rho_{sp}\,dt \\ &= \sigma_s\sigma_p\rho_{sp} \cdot\ S_t P_t\, dt \end{aligned}$$

Applying Itô's quotient rule to $s_t \equiv \frac{S_t}{P_t}$ and using our expression for the quadratic covariation, we can write:

$$ds_t = \left[r_s - \bar{\pi} - \frac{1}{2}\sigma_p^2 + \sigma_s\sigma_p\rho_{sp}\right] \cdot s_t\ dt + s_t\ (\sigma_s\, dW_t^s - \sigma_p\, dW_t^p).$$

Next, define real variables $a_t \equiv A_t/P_t,\ y_t \equiv Y_t/P_t$. As B_t contains no diffusion, $d[A,P]_t = d[S,P]_t$. Using the quotient rule on a_t, and substituting the above gives:

$$da_t = \left[y_t + (r - \bar{\pi} - \frac{1}{2}\sigma_p^2)\,a_t + \left(r_s - r + \sigma_s\sigma_p\rho_{sp}\right)s_t - c_t\right]dt + (\sigma_s s_t\, dW_t^s - \sigma_p a_t\, dW_t^p)$$

As Hansen and Sargent (2001) robustness requires independent stochastic processes, we rewrite the stochastic differential equation for a_t in terms of two independent Brownian motions. Introducing a Brownian motion W^z independent of W^p, we can state: $dW_t^s = \rho_{sp}\, dW_t^p + \sqrt{1-\rho_{sp}^2}\, dW_t^z$. Finally, substituting this into our equation for da_t yields the final form used in the text:

$$da_t = [y_t + (r - \bar{\pi} - \frac{1}{2}\sigma_p^2)\,a_t + (r_s - r + \sigma_s\sigma_p\rho_{sp})s_t - c_t]dt + (\sigma_s\rho_{sp}\, s_t - \sigma_p a_t)\, dW_t^p + \sigma_s\sqrt{1-\rho_{sp}^2}\ s_t\, dW_t^z$$

Finally, defining $\omega_t \equiv \frac{s_t}{a_t}$ as the portfolio share of risky-assets, we can write:

$$da_t = \left[y_t + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2 + \left(r_s - r + \sigma_s \sigma_p \rho_{sp} \right) \cdot \omega_t \right) \cdot a_t - c_t \right] dt +$$
$$+ (\sigma_s \rho_{sp} \cdot \omega_t - \sigma_p) \cdot a_t \; dW_t^p + \sigma_s \sqrt{1 - \rho_{sp}^2} \cdot \omega_t \cdot a_t \; dW_t^z$$

C.4 Basic Model Proofs

This section details proofs supporting Section 3.3. Below is the set of Assumptions we will refer to in the proofs:

A1 **Preferences.** $u \in \mathcal{C}^2$, $u' > 0$, $u'' < 0$, Inada at 0 and ∞.

A2 **Relative risk aversion bounds.** $R(a,y) := -\dfrac{c^*(a,y)\,u''(c^*(a,y))}{u'(c^*(a,y))} \in [\underline{\gamma}, \bar{\gamma}]$ with $0 < \underline{\gamma} \leq \bar{\gamma} < \infty$.

A3 **Budget and constraints.** Wealth $a_t \geq 0$, labour income $y_t > 0$, portfolio share $\omega^* \in [0,1]$ (no short-selling).

A4 **Model parameters.** $\theta_p, \theta_z, \sigma_p, \sigma_s > 0$, correlation $\rho_{sp} \in [-1,1]$ with $\rho_{sp} \leq \sigma_p/\sigma_s$ ("low correlation").

A5 **Decay of marginal utility.** There exists $\varepsilon > 0$ s.t. $a^{1+\varepsilon} u'(c^*(a,y)) \to 0$ as $a \to \infty$.

A6 **Penalty dominance.** $\theta_z\,(1-\rho_{sp}^2) \;\geq\; \theta_p\,\rho_{sp}\,\frac{\sigma_p}{\sigma_s}$.

Definition of Expectation Wedges. We are interested in how the difference between the robust expectation and the rational expectation vary as wealth varies. Define the expectation difference as:

$$\Delta_p^E(a,y) \;\equiv\; \mathbb{E}^{\text{robust}}[d\log(P_t)] - \mathbb{E}^{\text{rational}}[d\log(P_t)] = \sigma_p\;h^{p*}(a,y)$$
$$\Delta_r^E(a,y) \;\equiv\; \mathbb{E}^{\text{robust}}\left[\frac{dS_t}{S_t}\right] - \mathbb{E}^{\text{rational}}\left[\frac{dS_t}{S_t}\right] = \sigma_s\cdot\left(\rho_{sp}\cdot h^{p*}(a,y) + (1-\rho_{sp}^2)^{\frac{1}{2}}\cdot h^{z*}(a,y)\right)$$

First-order conditions (for reference). The (interior) first-order conditions of the Hamilton-Jacobi-Bellman (3.1) equation. These are:

$$u'(c^*) = \partial_a V \tag{3.4}$$
$$\omega' = \frac{-\partial_a V}{\sigma_s^2\cdot a\cdot\partial_{aa}V}\cdot\left[\left(r_s - r + \sigma_s\sigma_p\rho_{sp}\right) + \sigma_s\rho_{sp}\cdot h^{p*} + \sigma_s\sqrt{1-\rho_{sp}^2}\cdot h^{z*}\right] + \frac{\sigma_s\sigma_p\rho_{sp}}{\sigma_s^2} \tag{3.5}$$
$$h^{p*} = -\theta_p\cdot\left(\sigma_s\rho_{sp}\cdot\omega^* - \sigma_p\right)\cdot a\cdot\partial_a V \tag{3.6}$$
$$h^{z*} = -\theta_z\cdot\sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}}\cdot\omega^*\cdot a\cdot\partial_a V \tag{3.7}$$

Substituting in the FOCs for the optimal distortions, we can write the interior solution for the optimal portfolio share as

$$\omega' = \frac{-\partial_a V}{a \cdot \left(\partial_{aa} V - (\theta_z(1-\rho_{sp}^2) + \theta_p \rho_{sp}^2) \cdot (\partial_a V)^2\right)} \cdot \left[\frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} - \frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot \frac{a \cdot (\partial_{aa} V - \theta_p (\partial_a V)^2)}{\partial_a V}\right] \tag{3.8}$$

Where $\theta_p,\ \theta_z,\ \sigma_p,\ \sigma_s > 0$ and $\rho_{sp} \in [-1,\ 1]$. Note that $\omega^* = \min\{\max\{\omega', 0\}, 1\}$. Our proofs employ that ω^* is bounded by assumption[46].

Shorthand. Set $s^* = a\omega^*$. Throughout, $O(\cdot)$ and $o(\cdot)$ are standard Landau symbols.

Proof of Lemma 3.3.1 (sign of Distortions)

Lemma. *Under Assumptions (A1)–(A4),*

$$h^{p*} \geq 0, \qquad \Delta_p^E \geq 0 \quad \forall a \geq 0,\ y > 0.$$

Assuming also (A6):

$$h^{z*} \leq 0, \quad \Delta_r^E \leq 0 \quad \forall a \geq 0,\ y > 0.$$

Both wedges equal zero only at $a = 0$ or $a = \infty$.

Proof. *Assumptions used: (A2), (A3), (A4).*

1. $u' > 0 \Rightarrow \partial_a V > 0$ by (3.4).

2. With $\omega^* \in [0, 1]$ and $\rho_{sp} \leq \sigma_p/\sigma_s$ (A4), the bracket in (3.6) is non-positive, hence $h^{p*} \geq 0$; sign of h^{z*} follows analogously.

3. $\Delta_p^E = \sigma_p h^{p*} \geq 0$ since $\sigma_p > 0$.

4. By (3.6) and (3.7): $\Delta_r^E = -\sigma_s\, a\, u'(c^*) \left[\rho_{sp}\, \theta_p(\sigma_s \rho_{sp} \omega^* - \sigma_p) + \theta_z\, \sigma_s (1 - \rho_{sp}^2)\, \omega^*\right]$. By (A1), the first multiplicand is negative. By (A6) and (A4), the bracket term is non-negative. Therefore $\Delta_r^E \leq 0$. For any $a > 0$, $\Delta_r^E > 0$.

□

[46] In fact, it is a required condition for several proofs.

Proof of Lemma 3.3.2 (boundary limits)

Lemma. *Under Assumptions (A1), (A3), (A5),*

$$\lim_{a\downarrow 0} \Delta_p^E(a,y) = 0, \qquad \lim_{a\uparrow\infty} \Delta_p^E(a,y) = 0$$

$$\lim_{a\downarrow 0} \Delta_r^E(a,y) = 0, \qquad \lim_{a\uparrow\infty} \Delta_r^E(a,y) = 0$$

Proof. *Assumptions used: (A1), (A3), (A5).*

(i) $a \downarrow 0$. From the consumer's budget constraint, we can see that as $a \downarrow 0$ we have $\frac{da_t}{dt}|_{a\downarrow 0} = y_t - c_t$. Therefore as $y_t > 0$ in all states by (A3) and $\lim_{c\to 0} u'(c) \to \infty$ by (A1) we have $c^*(a,y) > 0 \;\; \forall\, a \geq 0\,,\; y > 0$.

We therefore know that $\lim_{a\to 0} u'(c^*(a,y)) < \infty,\; \forall\, y > 0$. From (3.6) this implies that $\lim_{a\to 0} \partial_a V(a,y) < \infty$. As this limit is finite, and the term $\lim_{a\to 0} \left(\sigma_s \rho_{sp} \cdot \omega^* - \sigma_p\right) \cdot a = 0$, we can say that $\lim_{a\downarrow 0} h^{p*}(a,y) = 0$ and therefore $\lim_{a\downarrow 0} \Delta_p^E(a,y) = 0$.

(ii) $a \uparrow \infty$. (A5) gives $a\, u'(c^*) \to 0$; multiply (3.6) by a to obtain $h^{p*} \to 0$.

(iii) Because $h^{p*}, h^{z*} \to 0$ by the same argument, the returns wedge Δ_r^E also converges to 0 at both boundaries. □

Proof of Proposition 3.3.1 (slope near bounds)

Proposition. *Assume two further refinements to (A2):*

- *A2(b)* ***Asymptotic constant RRA.*** $\lim_{a\to\infty} R(a,y) = \gamma_\infty \in (0,\infty)$
- *A2(c)* ***Relative Prudence Bounds.*** $u \in \mathcal{C}^3$. $P(a,y) := -\frac{c^*(a,y)\cdot u'''(c^*(a,y))}{u''(c^*(a,y))}$ *where* $\sup_a P(a,y) < \infty$.

Then under Assumptions (A1)–(A5), there exist $0 < a_0(y) < a_1(y) < \infty$ *such that*

$$\frac{d\Delta_p^E}{da} > 0 \text{ on } (0,a_0), \qquad \frac{d\Delta_p^E}{da} < 0 \text{ on } (a_1,\infty),$$

Further assume (A6):

$$\frac{d\Delta_r^E}{da} < 0 \text{ on } (0,a_0), \qquad \frac{d\Delta_r^E}{da} > 0 \text{ on } (a_1,\infty).$$

If (A2(b)-(c)) hold alongside $\underline{\gamma} > 1$, *the large-*a *inequalities hold for all* $a \geq a_1(y)$.

Proof. *Assumptions used: (A1)-(A6).*

Proof for Inflation Wedge: We first prove the limits for $\Delta_p^E(a,y)$. From (3.4), note that $\partial_a V = u'(c^*) > 0$ and $\partial_{aa} V = u''(c^*) \cdot \frac{dc^*}{da} < 0$. Write:

$$\frac{d\Delta_p^E}{da} = -\sigma_p\theta_p \left[\underbrace{\left(\sigma_s\rho_{sp}\frac{ds^*}{da} - \sigma_p\right)\cdot u'(c^*)}_{\equiv A_1(a,y)} + \underbrace{\left(\sigma_s\rho_{sp}\cdot s^* - \sigma_p \cdot a\right)\cdot u''(c^*)\frac{dc^*}{da}}_{\equiv A_2(a,y)} \right]$$

Small-a. $a = \varepsilon \downarrow 0$. By (A3), as $y > 0$ then $u'(c^*) \geq u'(c(0,\underline{y})) > 0 \implies u'(c^*) = O(1)$. As $\omega \in [0,1] \implies s^* \in [0,a]$. Therefore $s^*(a,y) = a\omega^*(a,y) = O(a) = O(\varepsilon)$ and $\frac{ds(a,y)}{da} = \omega(a,y) + a\frac{d\omega(a,y)}{da} \leq 1 = O(1)$. This implies that at $a = \varepsilon$ we have:

$$\begin{aligned} A_1(a,y) &= \left(\sigma_s\rho_{sp}\frac{ds^*(a,y)}{da} - \sigma_p\right)\cdot u'(c^*(a,y)) \\ &= (\sigma_s\rho_{sp}\cdot O(1) - \sigma_p)\cdot O(1) \\ &= -\sigma_p \cdot O(1) \end{aligned}$$

Where the final line follows from (A4).

By (A2) write $|u''(c^*)| = \frac{R(a,y)u'(c^*)}{c^*} \leq \bar{\gamma}\frac{u'(c^*)}{c^*}$. As established earlier, $u'(c^*) = O(1)$ and $c^* > 0$. Therefore, $u''(c^*) = -O(1)$.

By (A3), $\frac{dc^*(a,y)}{da} \in [0,1]$.[47] Therefore:

$$\begin{aligned} A_2(a,y) &= (\sigma_s\rho_{sp}\cdot s^*(a,y) - \sigma_p\cdot a)\cdot u''(c^*(a,y))\cdot\frac{dc^*(a,y)}{da} \\ &= (\sigma_s\rho_{sp}\cdot O(\varepsilon) - \sigma_p\cdot O(\varepsilon))\cdot(-O(1))\cdot O(1) \\ &= O(\varepsilon) \end{aligned}$$

Therefore the term $A_1(a,y)$ dominates $A_2(a,y)$ and:

$$\begin{aligned} \frac{d\Delta_p^E}{da} &= -\sigma_p\theta_p\left[-\sigma_p O(1) - O(1)\right] \\ \frac{d\Delta_p^E}{da} &= \sigma_p^2\theta_p O(1) > 0 \end{aligned}$$

[47] We know $\frac{dc^*}{da} \leq 1$ as agents cannot issue new debt within the instant dt, and so instantaneous consumption dc^* cannot exceed da. I.e. $dc^* = da - ds^*$. As $s^* = a\omega^* \in [0,a] \implies \frac{ds^*}{da} \in [0,1]$. So $\frac{dc^*}{da} \in [0,1]$.

And so for $a \downarrow 0$ (i.e. $a = \varepsilon > 0$) we know $\frac{d\Delta_p^E}{da} > 0$.

Large-a. $= \varepsilon \to \infty$. By (A3) $\omega \in [0,1] \implies s^* \leq a$ which implies $\lim_{a\to\infty} \frac{ds^*(a,y)}{da} \leq 1 \implies \frac{ds^*(a,y)}{da} = O(1)$. Therefore write:

$$
\begin{aligned}
A_1(a,y) &= \left(\sigma_s \rho_{sp} \frac{ds^*(a,y)}{da} - \sigma_p\right) \cdot u'(c^*(a,y)) \\
&= (\sigma_s \rho_{sp} \cdot O(1) - \sigma_p) \cdot u'(c^*(a,y))
\end{aligned}
$$

By (A2) we know $\left(R(a,y) := -\frac{c^*(a,y)u''(c^*(a,y))}{u'(c^*(a,y))} \geq \underline{\gamma} \implies R(a,y) = \underline{\gamma} + o(1)\right)$. Therefore:

$$
\begin{aligned}
u''(c^*(a,y)) &= -R(a,y) \cdot \frac{u'(c^*(a,y))}{c^*(a,y)} \\
\implies u''(c^*(a,y)) \cdot a &= -R(a,y) \cdot u'(c^*(a,y)) \cdot \frac{}{c^*(a,y)} \\
&= -[\underline{\gamma} + o(1)] \cdot \left(u'(c^*(a,y) \cdot \frac{a}{c^*(a,y)}\right)
\end{aligned}
$$

Under two additional (non-restrictive) refinements to (A2), we can characterise the asymptotic character of the final term. If:

- A2(b) **Asymptotic constant RRA.** $\lim_{a\to\infty} R(a,y) = \gamma_\infty \in (0,\infty)$
- A2(c) **Relative Prudence Bounds.** $u \in \mathcal{C}^3$. $P(a,y) := -\frac{c^*(a,y)\cdot u'''(c^*(a,y))}{u''(c^*(a,y))}$. Impose $\sup_a P(a,y) < \infty$.

Then under (A1)-(A3) and A2(b),(c), consumption is asymptotically linear (see Ma and Toda (2022) for a discussion and proof). This implies $\lim_{a\to\infty} \frac{c^*(a,y)}{a} = \lim_{a\to\infty} \frac{dc^*(a,y)}{da} = m_\infty \in (0,1)$[48]. Therefore $\lim_{a\to\infty} \frac{a}{c^*} \cdot \frac{dc^*}{da} = 1$. Using these in $A_2(a,y)$ yields:

$$
\begin{aligned}
A_2(a,y) &= (\sigma_s \rho_{sp} \cdot s^*(a,y) - \sigma_p \cdot a) \cdot u''(c^*(a,y)) \cdot \frac{dc^*(a,y)}{da} \\
&= (\sigma_s \rho_{sp} \cdot \omega^*(a,y) - \sigma_p) \cdot a \cdot u''(c^*(a,y)) \cdot \frac{dc^*(aa,y)}{da} \\
&= (\sigma_s \rho_{sp} \cdot O(1) - \sigma_p) \cdot \left(-\left[\underline{\gamma} + o(1)\right] \cdot u'(c^*(a,y))\right) \cdot \frac{a}{c^*(a,y)} \cdot \frac{dc^*(a,y)}{da} \\
&= (\sigma_s \rho_{sp} \cdot O(1) - \sigma_p) \cdot \left(-\underline{\gamma} \cdot u'(c^*(a,y))\right)
\end{aligned}
$$

Where the last line follows from (A4). We now compare the magnitudes of A_1 and A_2. Taking their

[48] The graphical intuition for this is simple. If the consumption policy function asymptotes to a line, then in the limit the policy function and the ray from the origin will have the same gradient (i.e. MPC = APC).

ratio:

$$\frac{A_1(a,y)}{A_2(a,y)} \leq \frac{(\sigma_s \rho_{sp} \cdot O(1) - \sigma_p) \cdot u'(c''^*(a,y))}{-(\sigma_s \rho_{sp} \cdot O(1) - \sigma_p) \cdot \underline{\gamma} \cdot u'(c^*(a,y))}$$
$$= -\frac{1}{\underline{\gamma}}$$

Therefore $\underline{\gamma} > 1$ is a sufficient condition to ensure that $A_2(a,y)$ eventually dominates $A_1(a,y)$ for large a. To see this, write:

$$\begin{aligned}\frac{d\Delta_p^E}{da} &= -\sigma_p \theta_p [A_1(a,y) + A_2(a,y)] \\ &\leq -\sigma_p^2 \theta_p \left[\underline{\gamma} - 1\right] \cdot u'(c^*(a,y)) \\ &= -\sigma_p^2 \theta_p \left[\underline{\gamma} - 1\right] \cdot o\left(\frac{1}{a}\right)\end{aligned}$$

And the term in the square bracket is positive. And so for large enough a:

$$\frac{d\Delta_p^E}{da} = -\sigma_p^2 \theta_p \cdot o\left(\frac{1}{a}\right) < 0$$

Proof for Returns Wedge: Write the derivative of the returns wedge as:

$$\frac{d\Delta_r^E}{da} = -\sigma_s \Big[\underbrace{u'(c^*)\,\Psi}_{\equiv B_1} + \underbrace{a\,u''(c^*)\,\frac{dc^*}{da}\,\Psi}_{\equiv B_2} + \underbrace{a\,u'(c^*)\,\frac{d\Psi}{da}}_{\equiv B_3}\Big] \tag{3.9}$$

with $\Psi = \sigma_s \frac{d\omega}{da}\left[\rho_{sp}^2 + (1-\rho_{sp}^2)\theta_z\right]$ and $\frac{d\Psi}{da} = \sigma_s \omega_a^* (\rho_{sp}^2 \theta_p + (1-\rho_{sp}^2)\theta_z) > 0$. Note (A1), (A3) imply $u'(c^*) > 0,\ u''(c^*) < 0$ and $0 \leq \omega^* \leq 1,\ 0 \leq \omega_a^* \leq 1/a$, where $\omega_a^* \equiv \frac{d\omega^*(a,y)}{da}$. These, with (A6) imply $\Psi \geq 0$ for every a, y.

(i) Behaviour near $a = 0$. By (A1) and (A3),

$$u'(c^*) = O(1), \quad u''(c^*) = -O(1), \quad \Psi = O(1), \quad \omega_a^* = O(1), \quad \frac{dc^*}{da} = O(1).$$

Hence

$$B_1 = -O(1), \qquad B_2 = a\,O(1) = O(a), \qquad B_3 = -a\,O(1) = O(a).$$

For sufficiently small a the order-one negative contribution B_1 dominates the $O(a)$ terms B_2 and B_3, so

$$\frac{d\Delta_r^E}{da} < 0 \quad \text{for } 0 < a < a_0(y)$$

(ii) Behaviour as $a \longrightarrow \infty$. Use (A2) write $B_2 = -R(a, y) \cdot u'(c^*) \left(\frac{a}{c^*}\right) \cdot \frac{dc^*}{da} \cdot \Psi$. As in the previous subsection, assume A2(b)-(c). Therefore for large a:

$$\begin{aligned} |B_2| &= R(a, y) \cdot u'(c^*) \cdot \Psi \\ &\geq \underline{\gamma} u'(c^*) \cdot \Psi \end{aligned}$$

Also note that as $\omega_a \leq \frac{1}{a}$, we have $a \cdot \frac{d\Psi}{da} = O(a\omega_a^*) = O(1)$. Therefore:

$$\begin{aligned} B_3 &= a \cdot u'(c^*) \cdot \frac{d\Psi}{da} \\ &= O(1) \cdot O(a^{-(1+\varepsilon)}) \\ &= o(1) \end{aligned}$$

by (A5). Therefore we can write:

$$\begin{aligned} \frac{d\Delta_r^E}{da} &= -\sigma_s \left[B_1 + B_2 + B_3\right] \\ &> -\sigma_s \left[u'(c^*)\Psi - \underline{\gamma} u'(c^*)\Psi + o(1)\right] \end{aligned}$$

Where the last line follows from $B_2 < 0$ by (A1). As B_2 dominates the other terms in magnitude, the term in the square brackets is postive. Therefore there exists $a_1(y)$ such that $B_1 + B_2 + B_3 > 0$ for all $a > a_1(y)$.This yields:

$$\frac{d\Delta_r^E}{da}(a, y) > 0 \quad \text{for all } a > a_1(y).$$

□

Proof of Proposition 3.3.2 (hump shape)

Proposition. *Assume two further refinements to (A2):*

- *A2(b)* ***Asymptotic constant RRA.*** $\lim_{a\to\infty} R(a, y) = \gamma_\infty \in (0, \infty)$
- *A2(c)* ***Relative Prudence Bounds.*** $u \in C^3$. $P(a, y) := -\frac{c^*(a,y)\cdot u'''(c^*(a,y))}{u''(c^*(a,y))}$ *where* $\sup_a P(a, y) < \infty$.

Assume (A1)–(A6), alongside (A2(b)), (A2(c)) and $\underline{\gamma} > 1$. Then, for every $y > 0$,

- *$\Delta_p^E(a, y)$ is non-negative, single-peaked, and attains a unique interior maximum;*

Further assume (A6)

- *$\Delta_r^E(a, y)$ is non-positive, single-valleyed (inverse hump), and attains a unique interior minimum.*

Consequently, $\Delta_p^E(a, y)$ and $\Delta_r^E(a, y)$ admit unique extrema ($a^*_{\Delta p}(y)$ and $a^*_{\Delta r}(y)$) that satisfy the first-order conditions:

$$\frac{d\Delta_p^E(a^*_{\Delta p}, y)}{da} = 0 \qquad\qquad \frac{d\Delta_r^E(a^*_{\Delta r}, y)}{da} = 0$$

Proof. Combine Lemmas 3.3.1, 3.3.2 and Proposition 3.3.1. Continuity of Δ_p^E (Weierstrass) yields an interior maximum; the single sign change of its derivative guarantees uniqueness. The same argument applies to Δ_r^E. □

Proof of Proposition 3.3.3

Proposition. *Assume (A1)-(A5), and let $\underline{y}$ be the income in the lowest income state. Then:*

$$\lim_{a\to\infty} \Delta\omega^*(a) = 0$$

Furthermore, under a parameter restriction on θ_p, we have:

$$\Delta\omega^*(a) \;\geq\; 0 \quad \forall\, a \quad \textit{if} \quad \theta_p \;\leq\; 1 + \frac{r_s - r + \sigma_s\sigma_p\rho_{sp}}{\sigma_s\sigma_p|\rho_{sp}| \cdot u'(\underline{y})}$$

Proof. *Assumptions used: (A1)-(A5)*

Positive portfolio wedge We first prove that $\Delta\omega(a) \geq 0\ \forall a$. From (3.8), write:

$$\begin{aligned}
\omega^*(\theta_p, \theta_z) &= -\frac{V_a}{a}\frac{K_1(\theta_p)}{K_2(\theta_p, \theta_z)} \\
\text{where: } K_1(\theta_p) &:= \frac{r_s - r + \sigma_s\sigma_p\rho_{sp}}{\sigma_s^2} - \frac{\sigma_s\sigma_p\rho_{sp}}{\sigma_s^2}\cdot a\left(\frac{V_{aa}}{V_a} - \theta_p\cdot V_a^2\right) \\
K_2(\theta_p, \theta_z) &:= V_{aa} - \lambda(\theta_p, \theta_z)\cdot V_a^2 \;< 0\ \forall\ \theta_p, \theta_z, a, y \geq 0 \\
\lambda(\theta_p, \theta_z) &:= \theta_p\rho_{sp}^2 + \theta_z(1 - \rho_{sp}^2) > 0\ \forall\ \theta_p, \theta_z > 0
\end{aligned} \tag{3.10}$$

We first take the sign the (partial) derivative of ω w.r.t. θ_z

$$\frac{\partial \omega^*}{\partial \theta_z} = -\frac{V_a}{a} \cdot \frac{-[-(1-\rho_{sp}^2) \cdot V_a^2] \cdot K_1(\theta_p)}{K_2(\theta_p, \theta_z)^2} \quad \leq 0 \quad \text{if } K_1(\theta_p) \geq 0$$

Next, consider $\frac{\partial \omega^*}{\partial \theta_p}$:

$$\text{sign}\left(\frac{\partial \omega^*}{\partial \theta_p}\right) = \text{sign}\left(-\frac{V_a}{a}\right) \cdot \text{sign}\left(\left(\frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot a \cdot V_a^2\right) \cdot K_2 - K_1 \cdot (-\rho_{sp}^2 \cdot V_a^2)\right)$$

If $\rho_{sp} \leq 0$, then $\left(\frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot a \cdot V_a^2\right) \leq 0$. Therefore the first term inside the final bracket is positive. Similarly, if $K_1(\theta_p) \geq 0$ then $-K_1 \cdot (-\rho_{sp}^2 \cdot V_a^2) \geq 0$. Hence the final term is positive. As the first term is negative, we have:

$$\frac{\partial \omega^*}{\partial \theta_p} \leq 0 \quad \text{if } K_1(\theta_p) \geq 0$$

We now demonstrate the conditions under which $K_1(\theta_p) \geq 0$. As $K_1(\theta_p)$ is linear in θ_p, we can write $K_1(\theta_p) = K_1(0) + \theta_p \cdot \frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot a \cdot V_a^2$. Note that $K_1(0) = \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} - \frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot a \cdot \frac{V_{aa}}{V_a}$. Bounding this using the above:

$$K_1(0) \geq \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} - \frac{\sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} \cdot (-\bar{\gamma}) = \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s^2} + \frac{\sigma_s \sigma_p |\rho_{sp}|}{\sigma_s^2} \cdot \bar{\gamma}$$

Next we bound the denominator. By (A2), we have:

$$\begin{aligned} C := -a \cdot \frac{V_{aa}}{V_a} &= -a \cdot \frac{u''(c^*)}{u'(c^*)} \cdot \frac{dc^*}{da} \\ &= \frac{R(a,y)}{c^*/a} \cdot \frac{dc^*}{da} \\ &= R(a,y) \cdot \frac{\frac{dc^*}{da}}{\frac{c^*}{a}} \\ &\leq \bar{\gamma} \end{aligned}$$

where the final line follows from $\frac{dc^*}{da} \leq \frac{c^*}{a}$ (i.e. MPC $\leq$ APC).

$$\begin{aligned} a \cdot V_a^2 &= -a \cdot \left(\frac{V_{aa}}{V_a}\right) \cdot \left(\frac{V_a^2 \cdot V_a}{-V_{aa}}\right) \\ &= C \cdot \left(\frac{V_a^2}{-(V_{aa}/V_a)}\right) \end{aligned}$$

By (A2) and (A4) $-V_{aa}/V_a = -\frac{u''(c^*)}{u'(c^*)} \cdot \frac{dc^*}{da} \geq 0$. Furthermore, $\left(\frac{V_a^2}{-(V_{aa}/V_a)}\right) \geq 0$. Together, these two conditions imply:

$$\left(\frac{V_a^2}{-(V_{aa}/V_a)}\right) \leq |V_a|$$
$$\implies aV_a^2 \leq C \cdot |V_a|$$
$$\implies aV_a^2 \leq \bar{\gamma} \cdot |V_a|$$

Finally, note that by (A1) and (A3), as utility is strictly drecreasing in concavity and there is no borrowing, $|V_a| \leq u'(\underline{y})$ where $\underline{y} = \min y$. Using these bounds for the numerator and denominator into $K_1(\theta_p)$, we can write:

$$\begin{aligned} K_1(\theta_p) &\geq \frac{r_s - r + \sigma_s\sigma_p\rho_{sp}}{\sigma_s^2} + \frac{\sigma_s\sigma_p|\rho_{sp}|}{\sigma_s^2} \cdot \bar{\gamma} \; - \theta_p \cdot \bar{\gamma} \cdot \frac{\sigma_s\sigma_p|\rho_{sp}|}{\sigma_s^2}|V_a| \\ \implies \theta_p &\leq 1 + \frac{r_s - r + \sigma_s\sigma_p\rho_{sp}}{\sigma_s\sigma_p|\rho_{sp}|u'(\underline{y})} \end{aligned}$$

This condition is sufficient to ensure $K_1(\theta_p) \geq 0$ and so both $\frac{\partial \omega^*}{\partial \theta_p}, \frac{\partial \omega^*}{\partial \theta_z} \leq 0$. Therefore integrating from $(0,0)$ to (θ_p, θ_z) along any path ensures $\omega^*(\theta_p, \theta_z) \leq \omega^*(0,0)$. Therefore $\Delta\omega(a) \geq 0 \; \forall a$. □

Zero Portfolio Wedge for large a Under (A1), (A3) and (A5), $h^{p*}, \; h^{z*} \to 0$ as $a \to \infty$, as in the proof for Lemma 3.3.2. Therefore $\omega^*(\theta_p, \theta_z) \to \omega^*(0,0)$. □

C.5 Computation of Baseline Model

For the sake of ease of reading, we outline again the HJB associated with the investors problem. However we modify the problem slightly so that the investor picks the risky-asset position outright, rather than a portfolio share. This is merely a substitution $s = \omega \cdot a$, and does not materially change the problem. However it aligns our description of the computational process with Achdou et al. (2022). We have also found such a transformation to be more numerical stable. The investor HJB is therefore:

$$
\begin{aligned}
V(a,y) = \max_{c,k} \min_{h^z,h^p} & \frac{c^{1-\gamma}}{1-\gamma} + \frac{1}{2}\left(\frac{1}{\theta_z}(h^z)^2 + \frac{1}{\theta_p}(h^p)^2\right) + \\
& + \left[s(a,y,c,k) + (\sigma_s \rho_{sp} k - \sigma_p a)\cdot h^p + \sigma_s (1-\rho_{sp}^2)^{\frac{1}{2}} k \cdot h^z\right] \cdot \partial_a V(a,y) + \\
& + \frac{1}{2}\left(\sigma_s^2 k^2 - 2\sigma_s \sigma_p \rho_{sp} \cdot ka + \sigma_p^2 a^2\right) \partial_{aa} V(a,y) + \sum_{y \neq y'} \lambda_{y \to y'} (V(a,y') - V(a,y))
\end{aligned}
$$

where $s(a,y,c,k) \equiv y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) a + (r_s - r + \sigma_s \sigma_p \rho_{sp}) \cdot k - c$ represents the instantaneous savings of an agent in state (a, y) with policy functions (c, k).

Boundary Conditions Our problem includes a state constraint, such that $a \geq a_{\min}$. As is standard in HJB problems of this form, the first order condition $u'(c(a_{\min}, y)) = \partial_a V(a_{\min}, y)$ still holds at this borrowing constraint. To ensure that the constraint is respected, we therefore require that the drift of a at $a_{\min}$ is weakly positive. In a problem with no concern for robustness, the state constraint would generate a boundary condition:

$$
\partial_a V(a_{\min}, y) \geq \left(y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) a_{\min}\right)^{-\gamma}
$$

However, in the robustness problem, the drift of a is no longer simply savings $s(a, y)$ but rather $s(a,y,c,k) + \mu(a,k,h^p,h^z)$, where $\mu(a,k,h^p,h^z) = (\sigma_s \rho_{sp} k - \sigma_p a) \cdot h^p + \sigma_s (1-\rho_{sp}^2)^{\frac{1}{2}} k \cdot h^z$. The first order conditions on h^p and h^z do not yield closed-form solutions for low values of $a_{\min}$. In practice, we therefore impose the following boundary condition:

$$
\partial_a V(a_{\min}, y) \geq \left(y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) a_{\min} + \mu\left(a_{\min}, y\right)\right)^{-\gamma}
$$

$$
\text{where:} \quad \mu(a_{\min}, y) = (\sigma_s \rho_{sp} k(a_{\min}, y) - \sigma_p a_{\min}) \cdot h^p(a_{\min}, y) + \sigma_s (1-\rho_{sp}^2)^{\frac{1}{2}} k(a_{\min}, y) \cdot h^z(a_{\min}, y)
$$

with $k(a_{\min}, y)$, $h^p(a_{\min}, y)$, $h^{z*}(a_{\min}, y)$ given by the first order conditions evaluated at state vector $(a_{\min}, y)$. As we impose $a_{\min} = 0$ in the main body of the text, the extra contribution generated by robustness concerns to the drift of a at $a_{\min} = 0$ is simply zero and the boundary constraint is simply:

$$\partial_a V(0, y) \geq (y)^{-\gamma}$$

However, for other values of $a_{\min}$, we have found that numerically we can impose this condition by first assuming an initial value for the optimal drift distortion at the bottom of the a grid, and then updating this value with each iteration.

As the HJB must be solved in practice over a bounded interval, we must also impose a boundary condition at $a_{\max}$. As we use CRRA utility, and the impact of robustness over time declines under the conditions outlined in the paper (and met by our calibration), we can employing the boundary condition used in Achdou et al. (2022). As discussed there, for large a we know that $V(a, y) = \alpha_0(y) + \alpha_1(y)a^{1-\gamma}$ for some unknown $\alpha_0(y)$ and $\alpha_1(y)$. We impose the bondary condition:

$$\partial_{aa} V(a_{\max}, y) = -\gamma \frac{\partial_a V(a_{\max}, y)}{a_{\max}}.$$

We know that for large a we obtain the expected-utility (i.e. Merton problem) solution for k. From the first order condition on k, we can therefore write:

$$k(a_{\max}, y) = \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\gamma \sigma_s^2} a_{\max}$$

We can therefore bound the second-order term as the following:

$$\frac{1}{2}\left(\sigma_s^2 k(a_{\max}, y)^2 - 2\sigma_s \sigma_p \rho_{sp} \cdot k(a_{\max}, y) \cdot a_{\max} + \sigma_p^2 a_{\max}^2\right) \cdot \partial_{aa} V(a_{\max}, y) = \xi \cdot \partial_a V(a_{\max}, y)$$

where $\xi \equiv -a_{\max}\left(\frac{(r_s - r + \sigma_s \sigma_p \rho_{sp})^2}{2\gamma\sigma_s^2} - \sigma_p \rho_{sp} \cdot \frac{r_s - r + \sigma_s \sigma_p \rho_{sp}}{\sigma_s} + \frac{\sigma_p^2 \gamma}{2}\right)$. We impose this boundary condition when using a finite difference solution method.

Therefore our boundary conditions are:

$$\partial_a V(0, y) \geq (y)^{-\gamma}$$
$$\partial_a V(a_{\max}, y) \leq \left(y + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) a_{\max} + \frac{(r_s - r + \sigma_s \sigma_p \rho_{sp})^2}{\gamma \sigma_s^2} a_{\max}\right)^{-\gamma}$$

Upwind Scheme We follow Appendix 8 of Achdou et al. (2022) and employ an implicit upwind scheme and discretize the problem on a non-uniform grid. This allows us to add more grid points at low values of a: i.e. where the value function has more curvature. Denoting by $\Delta a_{i,+} = a_{i+1} - a_i$and $\Delta a_{i,-} = a_i - a_{i-1}$ the forward and backwards distance between two grid-points, we can write our forward and backwards discrete approximations to the first partial derivatives as:

$$v_j'(a_i) \equiv \partial_a V(a_i, y_j) \approx \frac{v_{i+1,j} - v_{i,j}}{a_{i+1} - a_i} = \frac{v_{i+1,j} - v_{i,j}}{\Delta a_{i,+}} \equiv v'_{i,j,F}$$
$$v_j'(a_i) \equiv \partial_a V(a_i, y_j) \approx \frac{v_{i,j} - v_{i-1,j}}{a_i - a_{i-1}} = \frac{v_{i,j} - v_{i-1,j}}{\Delta a_{i,-}} \equiv v'_{i,j,B}$$

Similarly, the approximation of a second-derivative is given by:

$$\partial_{aa} V(a_i, y_j) \equiv v_j''(a_i) \approx \frac{\Delta a_{i,-} v_{i+1,j} - (\Delta a_{i,-} + \Delta a_{i,+}) v_{i,j} + \Delta a_{i,+} v_{i-1,j}}{\frac{1}{2}(\Delta a_{i,+} + \Delta a_{i,-}) \Delta a_{i,-} \Delta a_{i,+}}$$

We define $s^n_{i,j,F} = y_j + \left(r - \bar{\pi} - \frac{1}{2}\sigma_p^2\right) a_i + (r_s - r + \sigma_s \sigma_p \rho_{sp}) \cdot k^n_{i,j,F} - c^n_{i,j,F}$, where subscript F denotes a solution calculated using a forward difference approximations. We similarly define $s^n_{i,j,B}$. We also define $\mu^n_{i,j,F} = (\sigma_s \rho_{sp} \cdot k^n_{i,j,F} - \sigma_p a_i) \cdot h^{p\,n}_{i,j,F} + \sigma_s (1 - \rho_{sp}^2)^{\frac{1}{2}} \cdot k^n_{i,j,F} \cdot h^{z\,n}_{i,j,F}$. The discretisation of our HJB is therefore:

$$\frac{v^{n+1}_{i,j} - v^n_{i,j}}{\Delta} + \rho_{sp} v^{n+1}_{i,j} = u\left(c^n_{i,j}\right) + v^{n+1}_{i-1,j} x_{i,j} + v^{n+1}_{i,j} y_{i,j} + v^{n+1}_{i+1,j} z_{i,j} + \sum_{l \neq j} v^{n+1}_{i,l} \lambda_{j \to l}$$

$$\text{where: } x_{i,j} = -\frac{\left(s^n_{i,j,B} + \mu^n_{i,j,B}\right)^-}{\Delta a_{i,-}} + \frac{1}{2} \frac{\left(\sigma_s^2 k_{i,j}^2 - 2\sigma_s \sigma_p \rho_{sp} \cdot k_{i,j} a_i + \sigma_p^2 a_i^2\right)}{\frac{1}{2}(\Delta a_{i,+} + \Delta a_{i,-}) \Delta a_{i,-} \Delta a_{i,+}}$$

$$y_{i,j} = -\frac{\left(s^n_{i,j,F} + \mu^n_{i,j,F}\right)^+}{\Delta a_{i,+}} + \frac{\left(s^n_{i,j,B} + \mu^n_{i,j,B}\right)^-}{\Delta a_{i,-}} - \frac{\left(\sigma_s^2 k_{i,j}^2 - 2\sigma_s \sigma_p \rho_{sp} \cdot k_{i,j} a_i + \sigma_p^2 a_i^2\right)}{\frac{1}{2}(\Delta a_{i,+} + \Delta a_{i,-}) \Delta a_{i,-} \Delta a_{i,+}} - \sum_{l \neq j} \lambda_{j \to l}$$

$$z_{i,j} = \frac{\left(s^n_{i,j,F} + \mu^n_{i,j,F}\right)^+}{\Delta a_{i,+}} + \frac{1}{2} \frac{\left(\sigma_s^2 k_{i,j}^2 - 2\sigma_s \sigma_p \rho_{sp} \cdot k_{i,j} a_i + \sigma_p^2 a_i^2\right)}{\frac{1}{2}(\Delta a_{i,+} + \Delta a_{i,-}) \Delta a_{i,-} \Delta a_{i,+}}$$

where the notation $(\cdot)^+ = \max\{0, \cdot\}$ and $(\cdot)^- = \min\{\cdot, 0\}$ is employed. For the second-order term, we use the solution to the FOC on k using the upwinded first derivative approximation, here notated as $k_{i,j}$. The upwind scheme requires:

$$v'_{i,j} = v'_{i,j,F} \mathbf{1}_{\{s_{i,j,F} + \mu_{i,j,F} > 0\}} + v'_{i,j,B} \mathbf{1}_{\{s_{i,j,B} + \mu_{i,j,B} < 0\}} + \bar{v}'_{i,j} \mathbf{1}_{\{s_{i,j,F} + \mu_{i,j,F} \leq 0 \leq s_{i,j,B} + \mu_{i,j,B}\}}$$

For a grid over a of length I, at the upper boundary $a_{\max}$ we use:

$$x_{I,j} = -\frac{\left(s^n_{I,j,B} + \mu^n_{i,j,B}\right)^-}{\Delta a_{I,-}} - \frac{\xi}{\frac{1}{2}\left(\Delta a_{I,+} + \Delta a_{I,-}\right)\Delta a_{I,-}\Delta a_{I,+}},$$

$$y_{I,j} = -\frac{\left(s^n_{I,j,F} + \mu^n_{i,j,F}\right)^+}{\Delta a_{I,+}} + \frac{\left(s^n_{I,j,B} + \mu^n_{i,j,B}\right)^-}{\Delta a_{I,-}} + \frac{\xi}{\frac{1}{2}\left(\Delta a_{I,+} + \Delta a_{I,-}\right)\Delta a_{I,-}\Delta a_{I,+}} - \sum_{l \neq j} \lambda_{j \to l},$$

$$z_{I,j} = \frac{\left(s^n_{I,j,F} + \mu^n_{i,j,F}\right)^+}{\Delta a_{I,+}}$$

Where ξ is as described above. These equations form a system of $3 \times I$ linear equations. In matrix notation, we have:

$$\frac{1}{\Delta}(v^{n+1} - v^n) + \rho_{sp} v^{n+1} = u^n + \mathbf{A}^n v^{n+1}$$

Where A is sparse intensity matrix consisting of $x_{i,j}$, $y_{i,j}$, and $z_{i,j}$. This linear system is then solved as in Achdou et al. (2022). For the Kolmogorov Forward Equation, we use artificial reflecting barriers at $a_{\min}$ and $a_{\max}$, instead of the state constraint boundaries used above. This avoids the issue of the stochastic drift of a pushing agents near the boundaries outside of the grid.

C.6 Computation of the Adjustment Cost Model

In this section we detail the numerical method used to solve the model with fixed portfolio adjustment costs. The solution involves three main steps which depart from the standard finite difference approach of Achdou et al. (2022): (1) constructing the discrete state transition generator A for the two-asset state space, (2) formulating the Hamilton-Jacobi-Bellman quasi-variational inequality (HJBQVI) as a Linear Complementarity Problem (LCP) to solve for the value function, and (3) solving the Kolmogorov Forward Equation (KFE) to obtain the stationary distribution of agents across states in the presence of portfolio adjustments. We draw heavily on the resources of Prof. Benjamin Moll's website for parts (2) and (3). Below, we describe how the inclusion of an endogenous adjustment decision (with a fixed transaction cost) modifies these standard methods.

Discretising generator A (inaction region). Let $\{s_i\}_{i=1}^{I}$ be the illiquid-asset grid and $\{b_j\}_{j=1}^{J}$ the liquid-asset grid, with non-uniform steps $\Delta s_{i,\pm}$ and $\Delta b_{j,\pm}$. Let k denote income states. For every income state z_k we store a value $V_{ij,k}$. The HJB operator without adjustment is:

$$\mathcal{L}V \;=\; (s_s(s)+\mu_s(s))\,\partial_s V \;+\; (s_b(b,y)+\mu_b(b))\,\partial_b V + \frac{1}{2}\varsigma_1(s)^2\partial_{ss}V + \frac{1}{2}\varsigma_2(s,b)^2\partial_{bb}V + \varsigma_3(s,b)\partial^2_{sb}V + \Lambda(y)V$$

where:

$$\begin{aligned}
s_s(s) &= (r_s - \bar{\pi} - \frac{1}{2}\sigma_p^2 + \sigma_s\sigma_p\rho_{sp})\cdot s & s_b(b,y) &= (r - \bar{\pi} - \frac{1}{2}\sigma_p^2)\cdot b + y - c^* \\
\mu_s(s) &= (\sigma_s\rho_{sp})\cdot s\cdot h^{p*} + \sigma_s(1-\rho_{sp}^2)^{\frac{1}{2}}\cdot s\cdot h^{z*} & \mu_b(b) &= -\sigma_p\cdot b\cdot h^{p*} \\
\varsigma_1(s)^2 &= (\sigma_s^2 + \sigma_p^2 - 2\sigma_s\sigma_p\rho_{sp})\cdot s^2 & \varsigma_2(s,b)^2 &= \sigma_p^2\cdot b^2 \\
\varsigma_3(s,b) &= -\sigma_p(\rho_{sp}\sigma_s - \sigma_p)\cdot s\cdot b
\end{aligned} \tag{3.11}$$

And $\Lambda(y)$ is the generator associated with the Poisson process on income, and starred variables denote the policy functions which solve the agent's FOCs.

This is discretised at each interior node (i,j) into *eight* non-zero coefficients:

$$x_s^{ijk}, y_s^{ijk}, z_s^{ijk},\; x_b^{ijk}, y_b^{ijk}, z_b^{ijk},\; w_{NE}^{ijk}, w_{NW}^{ijk}, w_{SE}^{ijk}, w_{SW}^{ijk}$$

There are three for the s-direction, three for the b-direction, and four cross-weights. The first 6 are analogous to any upwind scheme solved with two endogenous state variables. All first derivatives are up winded; $s_{ijk,F/B}$, $\mu_{ijk,F/B}$ and their b-analogues are computed using forward/backward differences.

The diagonal diffusion terms enter as usual. The final four coefficients account for the cross-partial term in the HJB operator. The mixed term uses the symmetric four-point stencil:

$$\partial^2_{sb} V \approx \frac{(V_{i+1,j+1,k} - V_{i+1,j-1,k} - V_{i-1,j+1,k} + V_{i-1,j-1,k})}{(4\Delta s_{i,+}\Delta b_{j,+})}$$

Cross-weights are therefore:

$$w^{ijk}_{NE} = \frac{\varsigma_3(s_i, b_j)}{4\Delta s_{i,+}\Delta b_{j,+}}, \qquad w^{ijk}_{NW} = -w^{ijk}_{NE}, \qquad w^{ijk}_{SE} = -w^{ijk}_{NE}, \qquad w^{ijk}_{SW} = +w^{ijk}_{NE}.$$

This maintains that every row sums to zero.

With these coefficients the discretized HJB at node (i, j) can be written as:

$$\begin{aligned}\frac{V^{n+1}_{ijk} - V^{n}_{ijk}}{\Delta} + \rho V^{n+1}_{ijk} = u\left(c^n_{ijk}\right) &+ x^{ijk}_s V^{n+1}_{i-1,j,k} + z^{ijk}_s V^{n+1}_{i+1,j,k} + x^{ijk}_b V^{n+1}_{i,j-1,k} + z^{ijk}_b V^{n+1}_{i,j+1,k} \\ &+ w^{ijk}_{NE} V^{n+1}_{i+1,j+1,k} + w^{ijk}_{NW} V^{n+1}_{i-1,j+1,k} + w^{ijk}_{SE} V^{n+1}_{i+1,j-1,k} + w^{ijk}_{SW} V^{n+1}_{i-1,j-1,k} \\ &- \sum_{l\neq k} \lambda_{k\to l} V^{n+1}_{ijl}.\end{aligned}$$

Stacking these equations (over all (i, j, k)) gives the vector form

$$\frac{1}{\Delta}(v^{n+1} - v^n) + \rho v^{n+1} = u^n + A^n v^{n+1},$$

where v stacks the entries $V_{ij,k}$ and A^n is the sparse generator populated by the x–, y–, z– and w–coefficients above.

HJBQVI as an LCP. Let[49] $V^\star$ be the *post–adjustment value*, found after solving for the optimal reallocation vector by interpolating over a fine of cash on hand $x = s + b - \kappa$. At every iteration we seek v^{n+1} such that:

$$\min\{\rho v^{n+1} - u^n - A^n v^{n+1},\ v^{n+1} - V^*\} = 0$$

We can write this more efficiently in the following way:

$$(v - V^*)^\top(\rho v - u^n - A^n v) = 0,\ v \geq V^*,\ \rho v - u^n - A^n v \geq 0$$

[49]We refer the reader to the code snippets and associated documents on Prof. Benjamin Moll's website for more on solving LCP problems of this form, along with the associated KFEs.

Define the *excess-value* vector $z := v - V^*$, set $B := \rho I - A^n$ and $q := -u^n + BV^*$. The three conditions become the standard LCP:

$$z^\top (Bz + q) = 0, \qquad z \geq 0, \qquad Bz + q \geq 0$$

Because B is an M-matrix (diagonally dominant with non-positive off-diagonals), the LCP has a unique solution that can be obtained with any standard LCP solver. Convergence delivers the value function and—via the sign of z— the endogenous inaction/adjustment regions.

Stationary Kolmogorov–Forward equation We have to adjust our solution method (for both the perceived and the realised) KFE to account for discrete adjustments. Let $\mathcal{I}$ be the set of inaction nodes.

Construct the intervention matrix M, the elements of which are given by:

$$M_{\ell,j,k} = \begin{cases} 1, & \text{if } \ell \in \mathcal{I} \text{ and } \ell = j \\ 1, & \text{if } \ell \notin \mathcal{I} \text{ and } s^*(\ell) = j \\ 0, & \text{otherwise} \end{cases}$$

where $s^*(\ell)$ is the target after adjustment. The modified generator is $A_M = A\,M$, which routes any flow that lands in an adjustment node straight to its target. The ergodic density g satisfies:

$$g = M^\top g, \qquad (A_M)^\top g = 0, \qquad \sum g = 1.$$

Which can be solved by time-stepping and normalised.

To solve for the realised distribution (rather than the perceived distribution), simply omit $\mu_s(s)$ and $\mu_b(b)$ from the construction of A. For more information on this, see Kaplan, Maxted and Moll.

Bibliography

Abramson, B., J. Boerma, and A. Tsyvinski (2024). *Macroeconomics of Mental Health*. Technical report w32354. Cambridge, MA: National Bureau of Economic Research.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022). "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach". *The Review of Economic Studies* 89.1.

Aït-Sahalia, Y., F. Matthys, E. Osambela, and R. Sircar (2021). "When Uncertainty and Volatility Are Disconnected: Implications for Asset Pricing and Portfolio Performance".

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003). "A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection". *Journal of the European Economic Association* 1.1.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart (2022). "Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples". *The Review of Economic Studies* 89.6.

Ang, A., G. Bekaert, and M. Wei (2006). "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?"

Angeletos, G.-M., Z. Huo, and K. Sastry (2020). "Imperfect Macroeconomic Expectations: Evidence and Theory". *SSRN Electronic Journal*.

Angelico, C. and F. D. Giacomo (2024). "Heterogeneous inflation expectations and forecast errors across the income distribution: the role of memory".

Armantier, O., W. Bruine de Bruin, S. Potter, G. Topa, W. Van Der Klaauw, and B. Zafar (2013). "Measuring Inflation Expectations". *Annual Review of Economics* 5.1.

Armantier, O., W. Bruine de Bruin, G. Topa, W. Van Der Klaauw, and B. Zafar (2015). "Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?" *International Economic Review* 56.2.

Auclert, A. (2019). "Monetary Policy and the Redistribution Channel". *American Economic Review* 109.6.

Bachmann, R., T. O. Berg, and E. R. Sims (2015). "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence". *American Economic Journal: Economic Policy* 7.1.

Bahaj, S., R. Czech, S. Ding, and R. A. Reis (2023). "The Market for Inflation Risk". *SSRN Electronic Journal*.

Bańkowska, K., Ana Maria Borlescu, Evangelos Charalambakis, António Dias Da Silva, Davide Di Laurea, Maarten Dossche, Dimitris Georgarakos, Juha Honkkila, Neale Kennedy, Geoff Kenny, Aleksandra Kolndrekaj, Justus Meyer, Desislava Rusinova, Federica Teppa, and Veli-Matti Törmälehto (2021). "ECB Consumer Expectations Survey: an overview and first evaluation".

Baqaee, D. R., A. Burstein, and Y. Koike-Mori (2024). "A Sufficient Statistics Approach to Measuring Forward-Looking Welfare".

Baqaee, D. R. (2020). "Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles". *Journal of Monetary Economics* 114.

Bhandari, A., J. Borovička, and P. Ho (2024). "Survey Data and Subjective Beliefs in Business Cycle Models". *Review of Economic Studies*.

Bianchi, F., C. Ilut, and H. Saijo (2024). "Smooth Diagnostic Expectations". *SSRN Electronic Journal*.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). "Overreaction in Macroeconomic Expectations". *American Economic Review* 110.9.

Bordalo, P., N. Gennaioli, and A. Shleifer (2022). "Overreaction and Diagnostic Expectations in Macroeconomics". *Journal of Economic Perspectives* 36.3.

Brennan, M. J. and Y. Xia (2002). "Dynamic Asset Allocation under Inflation". *The Journal of Finance* 57.3.

Brenner, M. and Y. Izhakian (2018). "Asset pricing and ambiguity: Empirical evidence". *Journal of Financial Economics* 130.3.

Broer, T., A. Kohlhas, K. Mitman, and K. Schlafmann (2022). "Expectation and Wealth Heterogeneity in the Macroeconomy".

Bruine de Bruin, W., W. van der Klaauw, G. Topa, J. S. Downs, B. Fischhoff, and O. Armantier (2012). "The effect of question wording on consumers' reported inflation expectations". *Journal of Economic Psychology* 33.4.

Bruine de Bruin, W., W. Vanderklaauw, J. S. Downs, B. Fischhoff, G. Topa, and O. Armantier (2010). "Expectations of Inflation: The Role of Demographic Variables, Expectation Formation, and Financial Literacy". *Journal of Consumer Affairs* 44.2.

Bullard, J. (2013). *CPI Vs. PCE Inflation: Choosing a Standard Measure | St. Louis Fed*.

Cagetti, M., L. P. Hansen, T. Sargent, and N. Williams (2002). "Robustness and Pricing with Uncertain Growth". *Review of Financial Studies* 15.2.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2009). "Measuring the Financial Sophistication of Households". *American Economic Review* 99.2.

Campbell, J. Y. and L. M. Viceira (2001). "Who Should Buy Long-Term Bonds?" *American Economic Review* 91.1.

Campbell, J. Y. (2006). "Household Finance". *The Journal of Finance* 61.4.

Campbell, J. Y., R. J. Shiller, and L. M. Viceira (2009). "Understanding Inflation-Indexed Bond Markets". *SSRN Electronic Journal*.

Candia, B., O. Coibion, and Y. Gorodnichenko (2020). "Communication and the Beliefs of Economic Agents". *SSRN Electronic Journal*.

Carroll, C. D. (2003). "Macroeconomic Expectations of Households and Professional Forecasters". *The Quarterly Journal of Economics* 118.1.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024). "On Binscatter". *American Economic Review* 114.5.

Cavallo, A., F. Lippi, and K. Miyahara (2023). "Inflation and misallocation in New Keynesian models". *ECB Forum on Central Banking 26-28 June 2023, Sintra, Portugal: Macroeconomic Stabilisation in a Volatile Inflation Environment*.

CEA (2023). *Crosswalk Talk: What's the difference between the PCE and the CPI? | CEA*.

Chen, Z. and L. Epstein (2002). "Ambiguity, Risk, and Asset Returns in Continuous Time". *Econometrica* 70.4.

Cieslak, A. and C. Pflueger (2023). "Inflation and Asset Returns".

Cocco, J. F., F. Gomes, and P. Lopes (2019). "Evidence on Expectations of Household Finances". *SSRN Electronic Journal*.

Coibion, O., D. Georgarakos, Y. Gorodnichenko, and M. Van Rooij (2023). "How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial". *American Economic Journal: Macroeconomics* 15.3.

Coibion, O. and Y. Gorodnichenko (2015). "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts". *The American Economic Review* 105.8.

Coibion, O., Y. Gorodnichenko, and R. Kamdar (2018). "The Formation of Expectations, Inflation, and the Phillips Curve". *Journal of Economic Literature* 56.4.

Cooley, T. and G. Hansen (1998). "The Inflation Tax in a Real Business Cycle Model". *Real Business Cycles*. Edited by K. Salyer, J. Hartley, and K. Hoover.

Crump, R. K., S. Eusepi, A. Tambalotti, and G. Topa (2022). "Subjective intertemporal substitution". *Journal of Monetary Economics* 126.

D'Agostino, R. B. and A. Belanger (1990). "A Suggestion for Using Powerful and Informative Tests of Normality". *The American Statistician* 44.4.

D'Acunto, F., E. Charalambakis, D. Georgarakos, G. Kenny, J. Meyer, and M. Weber (2024). *Household Inflation Expectations: An Overview of Recent Insights for Monetary Policy*. Technical report w32488. Cambridge, MA: National Bureau of Economic Research.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2023). "IQ, Expectations, and Choice". *The Review of Economic Studies* 90.5.

D'Acunto, F., U. Malmendier, and M. Weber (2021). "Gender roles produce divergent economic expectations". *Proceedings of the National Academy of Sciences* 118.21.

D'Acunto, F. and M. Weber (2024). *Why Survey-Based Subjective Expectations are Meaningful and Important*. Technical report w32199. Cambridge, MA: National Bureau of Economic Research.

Del Canto, F., J. Grigsby, E. Qian, and C. Walsh (2023). *Are Inflationary Shocks Regressive? A Feasible Set Approach*. Technical report w31124. Cambridge, MA: National Bureau of Economic Research.

Dimmock, S. G., R. Kouwenberg, O. Mitchell, and K. Peijnenburg (2013). *Ambiguity Aversion and Household Portfolio Choice: Empirical Evidence*. Technical report w18743. Cambridge, MA: National Bureau of Economic Research.

Dimmock, S. G., R. Kouwenberg, O. S. Mitchell, and K. Peijnenburg (2015). "Estimating ambiguity preferences and perceptions in multiple prior models: Evidence from the field". *Journal of Risk and Uncertainty* 51.3.

Doepke, M. and M. Schneider (2006). "Inflation and the Redistribution of Nominal Wealth". *Journal of Political Economy* 114.6.

Drechsler, I. (2013). "Uncertainty, Time-Varying Fear, and Asset Prices". *The Journal of Finance* 68.5.

Duffie, D. (1990). "Chapter 3 Money in general equilibrium theory". *Handbook of Monetary Economics*. Volume 1. Elsevier.

Ellsberg, D. (1961). "Risk, Ambiguity, and the Savage Axioms". *The Quarterly Journal of Economics* 75.4.

Engelberg, J., C. F. Manski, and J. Williams (2009). "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters". *Journal of Business & Economic Statistics* 27.1.

Epstein, L. and M. Schneider (2008). "Ambiguity, Information Quality, and Asset Pricing". *The Journal of Finance* 63.1.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). "Heterogeneity and Persistence in Returns to Wealth". *Econometrica* 88.1.

Fama, E. F. and G. Schwert (1977). "Asset returns and inflation". *Journal of Financial Economics* 5.2.

Ferreira, C., J. M. Leiva Murillo, G. Nuno, Á. Ortiz, T. Rodrigo, and S. Vazquez (2024). "The Heterogeneous Impact of Inflation on Households Balance Sheets." *SSRN Electronic Journal*.

Fisher, I. (1933). "The Debt-Deflation Theory of Great Depressions". *Econometrica* 1.4.

Fofana, S., P. Patzelt, and R. Reis (2024). "Household Disagreement about Expected Inflation".

Gabaix, X. (2019). "Behavioral inattention". *Handbook of Behavioral Economics: Applications and Foundations 1*. Volume 2. Elsevier.

Garlappi, L., R. Uppal, and T. Wang (2007). "Portfolio Selection with Parameter and Model Uncertainty: A Multi-Prior Approach". *Review of Financial Studies* 20.1.

Gilboa, I. and D. Schmeidler (1989). "Maxmin expected utility with non-unique prior". *Journal of Mathematical Economics* 18.2.

Gomes, F. (2020). "Portfolio Choice Over the Life Cycle: A Survey". *Annual Review of Financial Economics* 12.1.

Gorodnichenko, Y. and D. Sergeyev (2021). *Zero Lower Bound on Inflation Expectations*. Technical report w29496. Cambridge, MA: National Bureau of Economic Research.

Haidari, Y. and G. Nolan (2022). "Sentiment, Uncertainty and Households' Inflation Expectations".

Hansen, L. P. and T. J. Sargent (2001). "Robust Control and Model Uncertainty". *American Economic Review* 91.2.

— (2005). "Robust Estimation and Control Under Commitment". *Journal of Economic Theory* 124.2.

— (2010a). "Fragile beliefs and the price of uncertainty". *Quantitative Economics* 1.1.

— (2010b). "Wanting Robustness in Macroeconomics". *Handbook of Monetary Economics*. Volume 3. Elsevier.

— (2011). "Robustness and ambiguity in continuous time". *Journal of Economic Theory* 146.3.

Hartley, J. and J. Mejia (2024). "Inflation Inequality Across Time and Space in the United States". *SSRN Electronic Journal*.

Huo, Z., M. Pedroni, and G. Pei (2023). "Bias and Sensitivity Under Ambiguity". *SSRN Electronic Journal*.

Ilut, C. and M. Schneider (2014). "Ambiguous Business Cycles". *American Economic Review* 104.8.

— (2023). "Ambiguity". *Handbook of Economic Expectations*. Elsevier.

Jaravel, X. (2019). "The unequal gains from product innovations: Evidence from the U.S. retail sector". *The Quarterly Journal of Economics* 134.2.

— (2021). "Inflation Inequality: Measurement, Causes, and Policy Implications". *Annual Review of Economics* 13.1.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). "Household Expenditure and the Income Tax Rebates of 2001". *American Economic Review* 96.5.

Ju, N. and J. Miao (2012). "Ambiguity, Learning, and Asset Returns". *Econometrica* 80.2.

Kamdar, R. and W. Ray (2024). "Attention-Driven Sentiment and the Business Cycle".

Kaplan, G. and S. Schulhofer-Wohl (2017). "Inflation at the household level". *Journal of Monetary Economics* 91.

Khorunzhina, N. (2013). "Structural estimation of stock market participation costs". *Journal of Economic Dynamics and Control* 37.12.

Kim, G. and C. Binder (2023). "Learning-through-Survey in Inflation Expectations". *American Economic Journal: Macroeconomics* 15.2.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005). "A Smooth Model of Decision Making under Ambiguity". *Econometrica* 73.6.

— (2009). "Recursive smooth ambiguity preferences". *Journal of Economic Theory* 144.3.

Kohlhas, A. N. and A. Walther (2021). "Asymmetric Attention". *American Economic Review* 111.9.

Kuhn, M., M. Schularick, and U. I. Steins (2020). "Income and Wealth Inequality in America, 1949–2016". *Journal of Political Economy* 128.9.

Li, K. (2019). "Portfolio selection with inflation-linked bonds and indexation lags". *Journal of Economic Dynamics and Control* 107.

Lucas, R. E. (1987). *Models of Business Cycles*. Paperback, August 1988. Yrjö Jahnsson lectures. Oxford: Blackwell.

Lucas Jr., R. E. (2000). "Inflation and Welfare". *Econometrica* 68.2.

Lusardi, A. and O. S. Mitchell (2014). "The Economic Importance of Financial Literacy: Theory and Evidence". *Journal of Economic Literature* 52.1.

Ma, Q. and A. A. Toda (2022). "Asymptotic Linearity of Consumption Functions and Computational Efficiency". *Journal of Mathematical Economics* 98.

Macaulay, A. (2022). "Heterogeneous Information, Subjective Model Beliefs, and the Time-Varying Transmission of Shocks". *SSRN Electronic Journal*.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences". *Econometrica* 74.6.

Machado, J. A. and J. Santos Silva (2019). "Quantiles via moments". *Journal of Econometrics* 213.1.

Maćkowiak, B., F. Matějka, and M. Wiederhold (2021). *Rational inattention: a review.* LU: Publications Office.

Maenhout, P. J. (2004). "Robust Portfolio Rules and Asset Pricing". *The Review of Financial Studies* 17.4.

— (2006). "Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium". *Journal of Economic Theory* 128.1.

Malmendier, U. and S. Nagel (2016). "Learning from Inflation Experiences". *The Quarterly Journal of Economics* 131.1.

Mankiw, N. G. and R. Reis (2002). "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve". *The Quarterly Journal of Economics* 117.4.

Mankiw, N. G., R. Reis, and J. Wolfers (2003). "Disagreement about Inflation Expectations". *NBER Macroeconomics Annual* 18.

Masolo, R. M. and F. Monti (2021). "Ambiguity, Monetary Policy and Trend Inflation". *Journal of the European Economic Association* 19.2.

— (2024). "Inflation Expectations of Savers and Borrowers".

Mejia, J. (2022). "Inflation's Compounding Impact on the Poor".

Merton, R. C. (1971). "Optimum consumption and portfolio rules in a continuous-time model". *Journal of Economic Theory* 3.4.

Moll, B. (2024). "The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics".

Montag, H. (2024). "On the Welfare Costs of Perceptions Biases". *Journal of Money, Credit and Banking*.

Nord, L. (2022). "Who Cares about Inflation? Endogenous Expectation Formation of Heterogeneous Households". *SSRN Electronic Journal*.

Pathak, P. A. (2002). "Notes on Robust Portfolio Choice".

Peijnenburg, K. (2018). "Life-Cycle Asset Allocation with Ambiguity Aversion and Learning". *The Journal of Financial and Quantitative Analysis* 53.5.

Pfäuti, O. (2023). *The Inflation Attention Threshold and Inflation Surges*.

Rawls, J. (2005). *A Theory of Justice: Original Edition*. Harvard University Press.

Reis, R. (2006). "Inattentive consumers". *Journal of Monetary Economics* 53.8.

— (2009). "The Time-Series Properties of Aggregate Consumption: Implications for the Costs of Fluctuations". *Journal of the European Economic Association* 7.4.

— (2021a). "Comment". *NBER Macroeconomics Annual* 35.

— (2021b). "The People versus the Markets: A Parsimonious Model of Inflation Expectations".

Roth, C., M. Wiederholt, and J. Wohlfart (2023). "The Effects of Monetary Policy: Theory with Measured Expectations". *SSRN Electronic Journal*.

Sahm, C. R., M. D. Shapiro, and J. Slemrod (2010). "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications". *Tax Policy and the Economy* 24.1.

Sargent, T. J. (1993). *Bounded rationality in macroeconomics: the Arne Ryde memorial lectures*. Oxford: Clarendon press.

Shapiro, A. H., M. Sudhof, and D. J. Wilson (2022). "Measuring news sentiment". *Journal of Econometrics* 228.2.

Stantcheva, S. (2024). *Why Do We Dislike Inflation?* Technical report w32300. Cambridge, MA: National Bureau of Economic Research.

Stokey, N. L. (2009). *The economics of inaction: stochastic control models with fixed costs*. Princeton: Princeton University Press.

Uppal, R. and T. Wang (2003). "Model Misspecification and Underdiversification". *The Journal of Finance* 58.6.

Van Der Klaauw, W., W. Bruine de Bruin, G. Topa, S. Potter, and M. F. Bryan (2008). "Rethinking the Measurement of Household Inflation Expectations: Preliminary Findings". *SSRN Electronic Journal*.

Vasicek, O. (1977). "An equilibrium characterization of the term structure". *Journal of Financial Economics* 5.2.

Viceira, L. M. (2018). "Global Portfolio Diversification for Long-Horizon Investors".

Vissing-Jørgensen, A. (2002). "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution". *Journal of Political Economy* 110.4.

Von Gaudecker, H.-M. and A. Wogrolly (2022). "Heterogeneity in households' stock market beliefs". *Journal of Econometrics* 231.1.

Weber, M., F. D'Acunto, Y. Gorodnichenko, and O. Coibion (2021). "Reality Check: How People Form Inflation Expectations and Why You Should Care".

— (2022). "The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications". *Journal of Economic Perspectives* 36.3.

Wolff, E. (2023). *Is There Really an Inflation Tax? Not For the Middle Class and the Ultra-Wealthy*. Technical report w31775. Cambridge, MA: National Bureau of Economic Research.

Yang, P. (2020). "Robust Time-Consistent Portfolio Selection for an Investor under CEV Model with Inflation Influence". *Mathematical Problems in Engineering* 2020.

Zhang, F., W. Fei, M. Shen, and K. Jiang (2019). "A study on optimal consumption and portfolio with labour income under inflation". *Systems Science & Control Engineering* 7.3.