Heterogeneous information, subjective model beliefs, and the time-varying transmission of shocks*

Alistair Macaulay†

December 10, 2021

Abstract

This paper explores the interaction of information and subjective models in the formation of expectations. In a general class of macroeconomic models, a novel decomposition highlights that systematic relationships between information and subjective models across agents distort the aggregate transmission of shocks. Using unique features of the Bank of England Inflation Attitudes Survey, I document evidence of such a systematic correlation between household information about inflation, and subjective models of how inflation affects the real economy. A model in which households face costs of acquiring information about inflation, and Knightian uncertainty about the way inflation affects real incomes, can explain the empirical variation in information and subjective models in the cross-section and over time. The model generates time-varying shock transmission, a selection effect that weakens the role of information frictions in aggregate dynamics, and a novel channel through which transitory spikes in inflation may become 'baked in' to the expectations of the households with the most pessimistic models of inflation, leading to large changes in the effects of subsequent shocks.

JEL codes: D83, D84, E31, E71

*I thank George-Marios Angeletos, Edouard Challe, Martin Ellison, Yuriy Gorodnichenko, Alexandre Kohlhas, Jennifer La'O, Michael McMahon, Ricardo Reis, Wenting Song, Andrew Usher, Laura Veldkamp, and seminar participants at the University of Oxford for valuable feedback.

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

1 Introduction

The Full Information Rational Expectations (FIRE) assumption underlying many macroeconomic models is composed of two parts: the first (FI) concerns what agents observe, while the second (RE) concerns the models agents use to turn their information into expectations of the future. Indeed, whether it is a sophisticated forecaster running inflation data through a structural VAR, or a naive household using a rule-of-thumb to map their current experiences to future expectations, forming an expectation relies on the combination of information and a model.

However, when studying plausible departures from the extreme sophistication of FIRE, recent literature has tended to consider these two assumptions separately. The noisy information and rational inattention literatures acknowledge that agents may not fully observe all variables in their environment in real time, but those papers typically assume agents know the true model governing the evolution of those variables.[1] In contrast, the literatures on learning, model uncertainty, imperfect common knowledge, and others explore misperceptions about equilibrium laws of motion, but assume that agents have full information about current realisations.[2] Both strands of literature have amassed empirical evidence in support of relaxing their component of the FIRE assumptions.[3]

In this paper I consider information and subjective models jointly. Specifically, I ask two questions: how do agent information and subjective models interact? And how does that interaction affect macroeconomic dynamics? In answering, the paper makes three main contributions:

1. It introduces a novel decomposition of the aggregate response to shocks in a general log-linear model with arbitrary information and subjective models. Shock transmission is shown to depend critically on the covariance of information and subjective models across the population.
2. It documents key facts about information and subjective models around inflation in UK household survey data, highlighting their relationship and behaviour over time.
3. It builds a model in which the interaction of rational inattention and model uncertainty accounts for the empirical findings. The interaction has implications for the role of information frictions in aggregate dynamics, the interpretation of information-treatment experiments, and the evolution of expectations and consumption following a transitory shock to inflation.

[1] e.g. Sims (2003), Reis (2006), Angeletos and La'O (2010), Maćkowiak and Wiederholt (2009, 2015), Angeletos and Sastry (2021), among many others.

[2] e.g. Bullard and Mitra (2002), Eusepi and Preston (2011), Ilut and Schneider (2014), Angeletos and Lian (2018), Farhi and Werning (2019), among many others.

[3] For information, see e.g. Link et al. (2021). For subjective models, see e.g. Andre et al. (2019).

In studying the joint behaviour of information and subjective models, this paper also provides a theoretical framework to explore the role of narratives in macroeconomics. Despite their use by policymakers to explain aggregate behaviour (e.g. Haldane, 2020), and the evidence collected by Shiller (2017) and others, they have not so far been incorporated into macroeconomic theory. A broad definition of narratives, following Gibbons and Prusak (2020), is that they comprise a state of the world and a resulting action. This is extremely similar to the information and subjective model pair studied in this paper, as an agent's response to information on a state of the world is determined by their subjective model.[4] The general decomposition in section 2 can therefore be interpreted as identifying the channels through which narratives can drive the transmission of aggregate shocks, and drawing attention in particular to the heterogeneity in narratives across households.

The role of narrative heterogeneity in the aggregate response to shocks shows the importance of any feedback that exists between information and models. Agents with different subjective models will react differently to the same piece of information, and so the aggregate response differs according to the distribution of information among those with heterogeneous subjective models.[5] Any feedback in which an agent's information affects their subjective model, or their subjective model affects the information they receive, will lead to systematic relationships between the two components of expectations, which will in turn alter the aggregate transmission of shocks.

In the second part of the paper I explore information and subjective models empirically. Since expectations are determined by the combination of information and models, survey or asset-price based data on expectations is insufficient. I therefore use unique features of the Bank of England's Inflation Attitudes Survey, which allow for the separation of information and subjective models about inflation. Specifically, respondents are asked about the information sources they used to arrive at their expectations, and how a hypothetical rise in inflation would affect the strength of the UK economy. The first of these therefore concerns information without involving the conclusions drawn from it. The second concerns the respondent's subjectve model of how inflation relates to the rest of the economy, without asking about what they believe inflation is or will be.

Using this data, I document three facts about household subjective models and information about inflation. First, the proportion of households who believe inflation is associated with a weaker economy is strongly positively correlated with recent realised

[4]Eliaz and Spiegler (2020) offer a formal model of narratives which is nested within this broad definition. In particular, their use of causal chains of variables represented by a DAG can be seen as an assumption about the kinds of subjective models that qualify as part of narratives.

[5]This is related to the earnings heterogeneity channel explored in Auclert (2019), in which heterogeneous marginal propensities to consume imply households react differently to income shocks, and so the transmission of aggregate shocks depends on how income changes are distributed among households.

inflation. Second, households who believe inflation makes no difference to the strength of the economy are less likely to use information about inflation when forming their expectations, while households who believe it is positively or negatively associated with the strength of the economy make use of similar information sources. Third, compared with households with other subjective models, households with negative subjective models of the effects of inflation expect substantially higher inflation, and perceive that inflation has been higher in the last year. Altogether, these facts show that in the case of inflation, information and subjective models are not independent of each other, and their joint distribution is not constant over time. The decomposition in the first part of the paper indicates that this interaction has implications for the aggregate transmission of shocks.

To explore those implications, in the final part of the paper I develop a model that is consistent with the empirical facts. The two key ingredients are that households face Knightian uncertainty over the effect of inflation on real incomes, and also face costs of acquiring information about inflation.

In the face of Knightian uncertainty, ambiguity averse households make choices based on the worst case model within some plausible set (Hansen and Sargent, 2008). For a household expecting high inflation, the worst case is that inflation erodes real wages, and makes the economy weaker. However, if a household expects very low inflation, then the worst case is that inflation does not hurt real incomes, and instead makes the economy stronger.[6] Information about inflation therefore affects subjective models. When realised inflation rises, households processing some information increase their expectations, and so the worst case belief becomes that inflation harms real incomes for a greater proportion of households. This therefore accounts for the time-variation in subjective models.

With costly information of the form in Sims (2003), household information is noisy, giving rise to heterogeneous inflation expectations and so heterogeneous worst-case models. In addition, subjective models influence optimal information choices, as the expected benefit of information depends on the extent to which it affects household choices. If a household has a subjective model that implies inflation is irrelevant for their choices, there is no benefit to information, so they will not pay the costs to process any. Subjective models therefore affect information choices, and this explains the cross-sectional relationship of information and subjective models.

Finally, Knightian uncertainty and noisy information can also explain the expectation and perception differences between households with different subjective models. Households receiving signals that inflation is high are more likely to believe that the worst case has inflation harming the economy. This relationship is strengthened if households also

[6]This is what a rational expectations model would predict at the zero lower bound on nominal interest rates, hence the use of forward guidance by many central banks after the Great Recession (Del Negro et al., 2015).

face Knightian uncertainty about the relationship between their signals and underlying inflation. If they believe inflation weakens the economy, the worst case is that the signals they received were in fact an under-estimate of true inflation, and they revise their perceptions and expectations up accordingly. This therefore generates a kind of confirmation bias, in which households expecting high inflation are more likely to interpret a given piece of information as indicating high future inflation.

This interaction between information and subjective models has several implications for aggregate behaviour. First, the dependence of information choice on heterogeneous subjective models implies a selection effect: the households who pay most attention to inflation are the ones who react the most to information about it. As in menu cost models of price setting (Golosov and Lucas, 2007), this reduces the imprtance of information frictions at the aggregate level, relative to the level of frictions implied by micro-evidence on inattention. This same selection effect also means that information treatment experiments aimed at estimating the causal effects of expectations (e.g. Coibion et al., 2019a) are driven largely by the households who respond the least to that information. Intuitively, households who respond weakly to new information choose to pay less attention to inflation, and so have more uncertain beliefs going into such an experiment. They therefore update their beliefs by more when shown the information treatment, compared with a more attentive household, and so they drive the estimates. However, when a shock hits the economy it is the attentive households who get the most information about it, and react most strongly.[7]

Second, the interaction of information and subjective models implies the transmission of shocks to aggregate consumption is history-dependent. A transitory spike in inflation causes more households to switch to believing inflation harms the economy. This means that subsequent inflation will be met with a more negative consumption response, particularly among those who are attentive to inflation, as their expectations will have been more strongly affected by the initial shock. Worse, if the confirmation bias mechanism operates, those attentive households will also begin to believe that inflation data is under-estimating true inflation, further increasing expectations in future periods among those most likely to observe subsequent shocks, and most likely to strongly reduce consumption in the face of further inflation. This therefore offers a novel channel through which expectations of high inflation may become 'baked in', as policymakers and journalists have worried about during 2021.[8] In particular, the model predicts inflation expectations are most likely to become baked in among those with the most negative reactions to inflation.

[7]Of course, in some settings the response of inattentive households is precisely the object of interest in these experiments. This is discussed further in section 5.1.

[8]See for example Michael Schumacher (Wells Fargo), quoted in Domm (2021): "[Jerome] Powell has sounded concerned about [inflation] expectations getting baked in."

Related literature. This paper principally contributes to the broad literatures on information frictions and subjective models in macroeconomics. A large number of papers have studied the role of limited information in business cycles (see Hubert and Ricco, 2018, for a review), either assuming that information is exogenous (e.g. Angeletos and La'O, 2010), or in the rational inattention literature that information is endogenous (Maćkowiak et al., 2020). These models typically assume that agents know the true equilibrium model of the economy: if they could observe the realisations of the exogenous shocks, they would be able to map perfectly from those into all endogenous variables. Similarly, literatures on learning (Eusepi and Preston, 2018), model uncertainty (Hansen and Sargent, 2008; Ilut and Schneider, 2014), imperfect common knowledge (Angeletos and Lian, 2018), level-k thinking (Farhi and Werning, 2019; Iovino and Sergeyev, 2021), and many others study the macroeconomic implications of agents misperceiving the true structural relationships and laws of motion in the economy.[9] These generally assume that agents observe the realisations of all variables in the current period. To my knowledge, this paper is the first to systematically study the combination of information frictions and subjective models, and consider how their interaction shapes aggregate dynamics.[10] The decomposition in the first part of the paper uses arbitrary forms of information frictions and subjective models, while the later parts suggest a novel combination of rational inattention and model uncertainty to explain the empirical results.

In much of this theoretical literature, the departure from FIRE naturally implies heterogeneity in expectations, as has been well-documented empirically (e.g. Carroll, 2003). Noisy information implies agents receive idiosyncratic signals, and recent work has also begun to explore differences in the incentives to acquire information across agents (Broer et al., 2020; Macaulay, 2021). A great many papers have explored heterogeneity in subjective models (e.g. Brock and Hommes, 1997; Branch and Evans, 2006; Farhi and Werning, 2019). However, in relaxing only one aspect of FIRE at a time, these models necessarily miss the narrative heterogeneity channel explored in this paper, as this relies on heterogeneity in both information and subjective models simultaneously.

I also contribute to the literature on the role of narratives in economic decisions. While theoretical models of narratives have been developed for questions in microeconomics and political economy (Bénabou et al., 2018; Eliaz and Spiegler, 2020), most work in macroeconomics has been empirical, concerned with tracking particular narratives and their impacts (Shiller, 2017; Larsen et al., 2021). The framework developed in this paper

[9]See Molavi (2019) for a discussion of how these various theories can all be seen as forms of misspecification in agent subjective models.

[10]Kohlhas and Robertson (2021) build a model in which agents learn over time about the relationship between a set of signals and endogenous variables of interest, but their focus is on this learning process, not on the nature of those signals, and whether there are frictions in observing them.

is a first step towards incorporating economic narratives into macroeconomic models.

The empirical part of the paper also contributes to our understanding of the narratives households use to understand inflation. This therefore relates to early work on the reasons many households dislike inflation Shiller (1997), and more recent work studying how this relates to expectations of other variables and actions (Andre et al., 2019; Candia et al., 2020; Kamdar, 2019). I extend this literature by separating information from subjective models in a survey with a long time series and rich data on household characteristics and other expectations. This allows me to observe information and subjective models about inflation at the individual level, over two decades in which macroeconomic conditions changed a great deal.

The IAS data I use has also been examined in other contexts. Bhandari et al. (2019) use the question on subjective models of inflation as a motivation for a model of time-varying ambiguity aversion. Nunes and Park (2020) study the response of consumption to expected inflation. Both of these, however, interpret the data using representative agent models with full information, so they do not cover the joint distribution of information and subjective models at the core of this paper. Michelacci and Paciello (2020) is more closely related, as they use other unique questions in the IAS to elicit preferences over inflation and interest rates, and explain the heterogeneity in these preferences by wealth using a model with Knightian uncertainty. However, while subjective models and preferences over aggregate variables are closely linked, this paper differs from them in studying the relationship of subjective models with information, rather than wealth.

Finally, the model with rational inattention and Knightian uncertainty also contributes to the literatures on the time-varying transmission of aggregate shocks,[11] selection effects in information (Yang, 2019; Afrouzi and Yang, 2021), and the use of information-treatment experiments in estimating the causal effects of expectations.[12] In particular, the model suggests a novel channel through which inflation expectations may remain high after a transitory inflation shock, particularly among those with the most negative subjective models of inflation, which affects the aggregate reaction to future inflationary shocks.

The rest of the paper is structured as follows. Section 2 derives the novel decomposition of aggregate responses to shocks in a general log-linear model with arbitrary information and subjective models. Section 3 explores information and subjective models about inflation in the IAS data. Section 4 develops the model of rational inattention and Knightian uncertainty that matches the empirical findings, and section 5 explores the implications of that model. Section 6 concludes.

[11]See e.g. Primiceri (2005), Galí and Gambetti (2009, 2015), Paul (2020), among many others.

[12]See e.g. Armona et al. (2019), Coibion et al. (2019a,b, 2020, 2021a), Roth and Wohlfart (2020).

2 A general result

I begin by presenting a decomposition of the aggregate household response to an arbitrary shock, in a general log-linear model. The decomposition highlights the roles played by information and subjective models, and their distribution across households, in determining the strength of aggregate shock transmission. Specifically, I show that the aggregate consumption response to a shock comes through three channels: the average narrative, the model heterogeneity channel, and the narrative heterogeneity channel.

Household $h \in H$ chooses a $N_x \times 1$ vector of choice variables $\boldsymbol{X}_t^h$ in period t. Letting lower case letters be log-deviations of variables from some arbitrary point, a log-linear approximation of their policy function can be written:[13]

$$\boldsymbol{x}_t^h = \boldsymbol{\mu}_t^h \mathbb{E}_t^h \boldsymbol{z}_t^h \tag{1}$$

where $\boldsymbol{z}_t^h$ is a $N_z \times 1$ vector of variables exogenous to the household, and $\boldsymbol{\mu}_t^h$ is a $N_x \times N_z$ matrix of coefficients. For simplicity, I will refer to these coefficients as representing the household's preferences, though note they could also come from any constraints the household faces.

The vector $\boldsymbol{z}_t^h$ may be arbitrarily large, and may include both aggregate and idiosyncratic variables. Some elements of $\boldsymbol{z}_t^h$ may be known or observed by household h, but for the unknown elements, the household-specific expectations operator $\mathbb{E}_t^h$ may or may not coincide with rational expectations. The elements of $\boldsymbol{z}_t^h$ may also be realised in any period: the indexation at time t simply reflects that they are the variables that matter for period t choices. Restricting $\boldsymbol{z}_t^h$ to only include variables the household takes as given is without loss of generality, as anything endogenous to the household can also be expressed as a linear function of other variables. Substituting out for that function in the $\boldsymbol{x}_t^h$, and continuing to do so for any remaining endogenous variables, gives a policy function only in terms of variables exogenous to the household.

This setup encompasses a wide range of models. For example, in the common textbook setup where households have CRRA utility over consumption, and can trade one-period risk-free bonds for intertemporal consumption smoothing, the consumption function log-

[13]This linearisation does not need to be taken about a steady state, or about the same point for each household. If two households have different idiosyncratic state variables, they can therefore have different responses to aggregate variables and expectations, just as they would in a fully non-linear model. This is why the coefficients $\boldsymbol{\mu}_t^h$ are indexed by household and by period, as the linearisation could be taken about different points each period.

linearised around the steady state can be expressed as:

$$c_t^h = (1-\beta)\sum_{s=0}^{\infty}\beta^s \mathbb{E}_t^h y_{t+s}^h - \sigma\beta\sum_{s=0}^{\infty}\beta^s(\mathbb{E}_t^h i_{t+s} - \mathbb{E}_t^h \pi_{t+s+1}) \tag{2}$$

where y_t^h is real income in period t, i_t is the nominal interest rate, and π_t is gross inflation. See Appendix A for the derivation.

This is the familiar result that consumption depends on the present value of future income and all future real interest rates. Within the framework of equation 1, $\boldsymbol{z}_t^h$ contains all current and future realisations of y_t^h, i_t, and π_{t+1}. The coefficients $\boldsymbol{\mu}_t^h$ contain the relevant combinations of the discount factor β and the inverse elasticity of intertemporal substitution σ. In this case all unconstrained households have the same consumption function, so the same $\boldsymbol{\mu}_t^h$ parameters. If we also introduce hand-to-mouth consumers, their consumption function is simply $c_t^h = y_t^h$. They can be written as having the same set of variables $\boldsymbol{z}_t^h$ as unconstrained households, but the coefficients in $\boldsymbol{\mu}_t^h$ are all set to zero, except the coefficient on y_t^h, which is 1.

I now consider a shock to one of the variables in the policy function z_{nt}^h. The reaction of household choices is determined by the changes in the expectation of each element of the policy function:

$$\frac{d\boldsymbol{x}_t^h}{dz_{nt}^h} = \boldsymbol{\mu}_t^h \frac{d\mathbb{E}_t^h \boldsymbol{z}_t^h}{dz_{nt}^h} \tag{3}$$

Applying the chain rule to the derivative of each element of $\mathbb{E}_t^h \boldsymbol{z}_t^h$ leads to a simple expression for the household's response to the shock.

Proposition 1 *For any household with policy function described by equation 1, the response to a shock to z_{nt}^h is given by:*

$$\frac{d\boldsymbol{x}_t^h}{dz_{nt}^h} = \boldsymbol{\mu}_t^h(\boldsymbol{I} - \boldsymbol{\mathcal{M}}_t^h)^{-1}\boldsymbol{\delta}_{nt}^h \tag{4}$$

where:

$$\boldsymbol{\mathcal{M}}_t^h = \begin{pmatrix} 0 & \mathcal{M}_{12,t} & \dots & \mathcal{M}_{1N_z,t} \\ \mathcal{M}_{21,t} & 0 & \dots & \mathcal{M}_{2N_z,t} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{M}_{N_z1,t} & \mathcal{M}_{N_z2,t} & \dots & 0 \end{pmatrix}, \qquad \mathcal{M}_{ij,t}^h \equiv \frac{\partial \mathbb{E}_t^h z_{it}^h}{\partial \mathbb{E}_t^h z_{jt}^h} \tag{5}$$

$$\boldsymbol{\delta}_{nt}^h = \left(\frac{d\mathbb{E}_t^h z_{1t}^h}{dz_{nt}^h}\bigg|_{\mathbb{E}_t^h z_{j\neq 1,t}}, \frac{d\mathbb{E}_t^h z_{2t}^h}{dz_{nt}^h}\bigg|_{\mathbb{E}_t^h z_{j\neq 2,t}}, \dots, \frac{d\mathbb{E}_t^h z_{N_zt}^h}{dz_{nt}^h}\bigg|_{\mathbb{E}_t^h z_{j\neq N_z,t}} \right)'$$

Proof. The derivative of the expectation of each element z_{it}^h of $\boldsymbol{z}_t^h$ can be decomposed using the chain rule:

$$\frac{d\mathbb{E}_t^h z_{it}^h}{dz_{nt}^h} = \frac{d\mathbb{E}_t^h z_{it}^h}{dz_{nt}^h}\bigg|_{\mathbb{E}_t^h z_{j\neq i,t}} + \sum_{j\neq i}^{N_z} \frac{\partial \mathbb{E}_t^h z_{it}^h}{\partial \mathbb{E}_t^h z_{jt}^h} \frac{d\mathbb{E}_t^h z_{jt}^h}{dz_{nt}^h} \tag{6}$$

Stacking this expression over all elements of $\boldsymbol{z}_t^h$ and rearranging gives:

$$\frac{d\mathbb{E}_t^h \boldsymbol{z}_t^h}{dz_{nt}^h} = (\boldsymbol{I} - \boldsymbol{\mathcal{M}}_t^h)^{-1} \boldsymbol{\delta}_{nt}^h \tag{7}$$

which substituted into equation 3 gives the result. ■

Equation 4 is useful because it distinctly highlights the separate roles played by the household's information, subjective model, and preferences. When the shock occurs, households receive some information that some of the variables that matter for their choices have changed. This is the $\boldsymbol{\delta}_{nt}^h$ component of equation 4. If they observe z_{it}^h, then the i^{th} element of $\boldsymbol{\delta}_{nt}^h$ is simply the derivative dz_{it}^h/dz_{nt}^h. However, if there is imperfect or noisy information, $\boldsymbol{\delta}_{nt}^h$ will not necessarily be equal to this full-information update, and indeed if the household gets no direct information about an element z_{it}^h then the i^{th} element of $\boldsymbol{\delta}_{nt}^h$ will be zero.

Direct information about z_{it}^h causes the household to update their expectation $\mathbb{E}_t^h z_{it}^h$. The next step in the household's belief formation is that this update to $\mathbb{E}_t^h z_{it}^h$ may in turn cause them to update their expectations of other variables, if the other variables are linked to z_{it}^h in their subjective model. The elements of $\boldsymbol{\mathcal{M}}_t^h$ capture these subjective model linkages, along with how the household weights the information in $\mathbb{E}_t^h z_{it}^h$ relative to other information they have about the variable in question. If the household is Bayesian, this weighting of information will be determined by the relative uncertainties in the prior belief about z_{jt}^h and the signal provided about z_{jt}^h by $\mathbb{E}_t^h z_{it}^h$. While I will assume households behave in this way in the model in section 4 below, it is not necessary for the results in this section.[14]

Importantly, $\boldsymbol{\mathcal{M}}_t^h$ only captures the direct effect of expectations of one variable on another. However, in all but the simplest subjective models updating may go through several steps. An update to $\mathbb{E}_t^h z_{it}^h$ may affect $\mathbb{E}_t^h z_{jt}^h$ directly, but also indirectly through its effect on the expectation of some other variable $\mathbb{E}_t^h z_{kt}^h$, which is linked in the household's subjective model to both z_{it}^h and z_{jt}^h. The matrix $(\boldsymbol{I} - \boldsymbol{\mathcal{M}}_t^h)^{-1}$ captures all such linkages, direct (in $\boldsymbol{\mathcal{M}}_t^h$) and indirect. From here, it will be convenient to work directly with this,

[14]For example, if households engage in cognitive discounting as in Gabaix (2020), they shrink the expectations implied by their model of the economy towards the steady state. That shrinkage would also be captured in the $\boldsymbol{\mathcal{M}}_t^h$ matrix.

which I refer to as the cross-learning matrix:[15]

$$\boldsymbol{\chi}_t^h \equiv (\boldsymbol{I} - \boldsymbol{\mathcal{M}}_t^h)^{-1} \tag{8}$$

where the $(i,j)^{th}$ element of $\boldsymbol{\chi}_t^h$ will be denoted $\chi_{ij,t}^h$. It is these values that are measured in the empirical literature on cross-learning (e.g. Roth and Wohlfart, 2020).

Finally, having updated all of their expectations using their information, and filtering it through their subjective model, household choices are determined by their reaction to each of those expectations, which is contained in the preference matrix $\boldsymbol{\mu}_t^h$. The information, subjective model, and response components of the household's economic narrative are therefore represented by $\boldsymbol{\delta}_{nt}^h$, $\boldsymbol{\chi}_t^h$, and $\boldsymbol{\mu}_t^h$ respectively.

Notice that full information rational expectations is nested in this framework, as with those assumptions expectations are still determined by information and subjective models. Consider, for example, a positive shock to future inflation π_{t+1} in the standard consumption function in equation 2. If the household receives news about the shock, their expected future inflation rises. This will lead them to update their expectations of future interest rates, according to their perception of the monetary policy reaction function. Both of these expectations will lead them to update their expectations over future real income, and perhaps in turn their inflation and interest rate expectations, again according to their model. Full information rational expectations is simply the special case of this where all current variables are observed, and the subjective model coincides with the true model in equilibrium.[16]

I now turn to aggregate behaviour. Consider a unit mass of the households modeled above. The vector of aggregate choices is given by:

$$\bar{\boldsymbol{x}}_t = (\bar{x}_{kt}), \quad \bar{x}_{kt} = \int_0^1 \omega_{kt}^h x_{kt}^h dh \tag{9}$$

[15]This has a parallel in the literature on production networks (see Carvalho and Tahbaz-Salehi, 2019, for a review). The direct links in $\boldsymbol{\mathcal{M}}_t^h$ are analogous to the elements of the input-output matrix, and $\boldsymbol{\chi}_t^h$ is the corresponding Leontief inverse. As in the production network literature, it is this Leontief inverse that regulates the transmission of shocks.

[16]This is therefore different from models that use Directed Acyclic Graphs (DAGs) to represent narratives, as there variables in narratives are assumed to have a causal ordering (see Spiegler, 2020, for a review). Most general equilbrium models do not have such a recursive ordering, in which case the true equilibrium laws of motion cannot be expressed as a DAG.

where ω^h_{kt} denotes a weighting applied to household h's choice x^h_{kt}, such that:[17]

$$\bar{\boldsymbol{x}}_t = \mathbb{E}_H \boldsymbol{x}^h_t \tag{10}$$

where $\mathbb{E}_H$ denotes the expected value across households.

Again, consider a shock to z_{nt}. The only change to the shock considered in Proposition 1 is that now the shock is assumed to be to an aggregate variable, so the h superscript is dropped.[18] Proposition 1 and the properties of covariances lead us to the decomposition of the aggregate choice response:

Proposition 2 *The response of aggregate choice $\bar{x}_{kt}$ to a shock to z_{nt} is given by:*

$$\frac{d\bar{x}_{kt}}{dz_{nt}} = \sum_{i=1}^{N_z}\sum_{j=1}^{N_z}\left[\bar{\mu}_{ki,t}\bar{\chi}_{ij,t}\bar{\delta}_{jn,t} + \bar{\delta}_{jn,t}Cov_H(\mu^h_{ki,t},\chi^h_{ij,t}) + Cov_H(\delta^h_{jn,t},\mu^h_{ki,t}\chi^h_{ij,t})\right] \tag{11}$$

where $\delta^h_{jn,t}$ and $\mu^h_{ki,t}$ denote the j^{th} element of $\boldsymbol{\delta}^h_{nt}$ and the $(k,i)^{th}$ element of $\boldsymbol{\mu}^h_t$ respectively, $\bar{\delta}_{jn,t}$ and $\bar{\mu}_{ki,t}$ are their aggregate counterparts, and $\bar{\chi}_{ij,t}$ is the aggregate value of $\chi^h_{ij,t}$.

Proof. From the definition of $\bar{x}_{kt}$ (equation 10), we have:

$$\frac{d\bar{x}_{kt}}{dz_{nt}} = \mathbb{E}_H \frac{dx^h_{kt}}{dz_{nt}} \tag{12}$$

The k^{th} row of equation 4 can be written as:

$$\frac{dx^h_{kt}}{dz_{nt}} = \sum_{i=1}^{N}\sum_{j=1}^{N}\mu^h_{ki,t}\chi^h_{ij,t}\delta^h_{jn,t} \tag{13}$$

Substituting this into equation 12 gives:

$$\frac{d\bar{x}_{kt}}{dz_{nt}} = \sum_{i=1}^{N}\sum_{j=1}^{N}\mathbb{E}_H\mu^h_{ki,t}\chi^h_{ij,t}\delta^h_{jn,t} \tag{14}$$

From the definition of covariance, $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y) + Cov(X,Y)$ for any X,Y.

[17]This weight is equal to the ratio of X^h_{kt} to aggregate $\bar{X}_{kt}$, evaluated at the point about which the log-linear approximation was taken. If the policy function of all households is approximated about the same point, $\omega^h_{kt} = 1$ for all h. If the approximation is taken about steady state, then the weights would reflect any steady state heterogeneity.

[18]This still allows for shocks with heterogeneous effects across households. For example, we can specify that the income of household h is given by $y^h_t = \alpha^{h\prime}\boldsymbol{z}^h_t$. A heterogeneous shock to incomes can therefore be represented by an aggregate shock to some latent z_{nt}, and all heterogeneity will be accounted for in the relevant element of the α^h vector, which in turn will enter the $\boldsymbol{\chi}^h_t$ cross-learning matrix.

Applying this to equation 14 implies:

$$\frac{d\bar{x}_{kt}}{dz_{nt}} = \sum_{i=1}^{N}\sum_{j=1}^{N}\left[\bar{\delta}_{jn,t}\mathbb{E}_H(\mu^h_{ki,t}\chi^h_{ij,t}) + Cov_H(\delta^h_{jn,t}, \mu^h_{ki,t}\chi^h_{ij,t})\right] \tag{15}$$

Applying the covariance formula again to the first term inside the sum in equation 15 implies equation 11. ■

This decomposition shows the three channels determining the aggregate response to shocks. The first term is the *average narrative*: changes in the average information, subjective model, and preferences over any variable will affect the aggregate response to shocks. This is the term that most empirical work on narratives in macroeconomics have focused on (e.g. Shiller, 2017; Larsen et al., 2021). The second term is the *model heterogeneity channel*, and the third is the *narrative heterogeneity channel*. These show that along with average narratives, aggregate responses to shocks may be strongly influenced by the distribution of information and subjective models across households.

To see the intuition for these channels, consider again the textbook consumption function in equation 2, and a positive shock to future inflation π_{t+1}. If all households start to believe that higher inflation is associated with lower real incomes, then the average $\chi^h_{y\pi,t}$ falls, and the aggregate consumption response to the future inflation will fall. This is the effect of the average narrative. If, however, this pessimistic subjective model of the effects of inflation only takes hold among hand-to-mouth households, then aggregate consumption will respond much more positively to the shock than the average would suggest, because the households who update their expected future real incomes down are also those who do not react strongly to $\mathbb{E}^h_t y^h_{t+1}$. This is the effect of the model heterogeneity channel. Finally, if all households are unconstrained, but the pessimistic model of inflation is concentrated among households who do not obtain any information about future inflation, then this again raises $d\bar{c}_t/d\pi_{t+1}$. Those households who would update $\mathbb{E}^h_t y^h_{t+1}$ down and reduce consumption accordingly if they learned that inflation was about to rise are also the households who do not observe the shock, and so do not update $\mathbb{E}^h_t \pi^h_{t+1}$. This is the effect of the narrative heterogeneity channel.

It is important to be clear that this is a decomposition, not a solution for aggregate actions. $\boldsymbol{\delta}^h_{nt}$ captures direct information received by the household, but the information received depends on the true reaction of $\boldsymbol{z}^h_t$ to the shock. That reaction will include general equilibrium effects, which may in turn depend on aggregate actions. This is nonetheless a useful exercise, as for a given movement in aggregate variables it highlights the channels through which shocks transmit to aggregate household behaviour. In this way it is similar to the decomposition in Auclert (2019), which takes movements in several aggregate variables as given to find the transmission to aggregate consumption.

Indeed, in this log-linear framework the earnings heterogeneity channel in Auclert (2019) is a special case of the narrative heterogeneity channel, in which all households have correct subjective models and full information. Households receive heterogeneous information about their earnings because the shock affects them differently, and the earnings heterogeneity channel operates when these differences correlate with differences in marginal propensities to consume out of income, which is one of the μ parameters in the model presented here.

In the following section I go on to provide survey evidence on these effects when considering a shock to inflation, which suggests that there is important heterogeneity in information and subjective models beyond that implied by fundamental heterogeneity in shock exposure, as studied in Auclert (2019) and Bilbiie (2019). I then build a model with rational inattention and subjective models based on ambiguity aversion over the implications of inflation. The empirical evidence and the model will suggest particular signs and time-series behaviours for the three terms in the context of inflation shocks. The results in this section, however, are more general. If we take seriously the notion that agents have heterogeneous information and heterogeneous subjective models of the economy, understanding aggregate behaviour requires understanding these three channels.

3 Survey evidence on information and subjective models of inflation

In this section I document three stylised facts about household information and subjective models. Specifically, the facts refer to the information households obtain about inflation, and their subjective models of how inflation is related to aggregate economic performance. These facts will inform the model in section 4.

First, I find that more households hold subjective models in which inflation damages the real economy when realised inflation has recently been high. Second, households with subjective models in which inflation has no real effects obtain less direct information about inflation. Third, households with subjective models in which inflation damages the economy have persistently higher inflation perceptions and expectations than those with other models.

To study the joint behaviour of information and subjective models, we need data that is informative about each separately. This is a challenge, as most papers studying evidence for either information frictions or subjective models use data on realised expectations, which are the result of a combination of both information and subjective models (as shown in section 2). Expectation data alone cannot therefore separate these two components of narratives.

As an example, Coibion and Gorodnichenko (2015) measure information frictions by regressing forecast errors on past forecast revisions of the same variable. They use inflation expectations in their application, as it has the most comprehensive survey data. They show that this recovers the degree of information frictions if agents form expectations using either a sticky or noisy information model. However, in both of those models, agents are assumed to know the true equilibrium law of motion for inflation. The same regression results could alternatively be consistent with agents who are fully informed about the current state of the world, but who use a misspecified or incorrect model to form forecasts of the future (Hajdini, 2020), or indeed with a combination of information frictions and subjective models.[19]

Standard expectations data is therefore insufficient for this exercise. To disentangle information and subjective models I use survey data from the Bank of England, which contains several hypothetical questions which explicitly ask respondents about their subjective model of the economy. These are similar in spirit to the vignettes used by Andre et al. (2019) to study cross-learning, which is denoted by $\chi^h_{ij,t}$ in section 2, and is critical in Propositions 1 and 2. The survey also contains questions on the information sources households use to arrive at their inflation expectations. Importantly, these questions are embedded in a survey with a long time series, and which contains rich data on household characteristics and expectations.

3.1 Data

The Inflation Attitudes Survey (IAS) has been fielded quarterly since 2001 to a repeated cross-section of UK households. In the first quarter of each year approximately 4000 households are surveyed, while in other quarters approximately 2000 are surveyed. I use the individual-level response data from 2001-2019, omitting the quarters conducted after the onset of the Covid-19 pandemic, as the implementation of the survey had to be changed substantially at this time (see Bank of England, 2020).

Alongside questions on expected inflation, interest rates, and other macroeconomic and personal variables, respondents are asked several questions which do not commonly appear in other household surveys. These questions are helpful in disentangling information and subjective models about inflation.

The first of these is Question 3, which asks households about their subjective model of the relationship between inflation and the 'strength of the economy'. It has been asked

[19] Coibion and Gorodnichenko (2015) consider one potential type of model misspecification, that agents shrink rational forecasts towards heterogeneous long-run priors, and find no evidence for it. However, other forms of subjective model distortions could offer alternative explanations for their results, as Gabaix (2020) shows for cognitive discounting.

in every wave of the survey.

Question 3 *If prices started to rise faster than they are now, do you think Britain's economy would end up stronger, or weaker, or would it make little difference?*

This differs from standard questions on expected future economic outcomes because it does not invoke the use of information about the state of the world. The answers to this question are therefore informative about cross-learning, denoted $\chi^h_{ij,t}$ in section 2.[20] In the analysis below, I will refer to a respondent answering that inflation would make the economy stronger/no difference/weaker as having a positive/neutral/negative subjective model of inflation respectively.

There are two possible interpretations of this question. Households may view it as asking about the causal effects of inflation on the economy (as in the model of Spiegler, 2021). Alternatively, they could see it as asking about the most likely source of a rise in inflation, if they believe supply- and demand-driven inflationary shocks are associated with different real outcomes (Kamdar, 2019). For the purposes of this paper, this distinction does not matter, as $\chi^h_{ij,t}$ in the decomposition of aggregate actions (Proposition 2) is simply the degree to which households update their expectations of one variable when their expectation of another changes. In this case, it is the updating of expectations about the strength of the economy when expected inflation rises. This is captured by the question, whether the updating occurs because of a perceived direct causal link from inflation to the real economy, or a belief about the kinds of shocks hitting the economy.

The second set of novel questions concern the information households use to arrive at their inflation expectations, without asking what those expectations are. This allows us to learn about household information ($\delta^h_{jn,t}$) without contamination from cross-learning ($\chi^h_{ij,t}$).

Question 2f *What were the most important factors in getting to your expectation for how prices in the shops would change over the next 12 months?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*

[20] In section 2 I noted that $\chi^h_{ij,t}$ comprised subjective models and any weighting the household put on expectations of z^h_{jt}. However, these weights do not change the sign of $\chi^h_{ij,t}$, which is all that is measured by the qualitative responses to Question 3, as long as no household perfectly observes the 'overall state of the economy'. I therefore interpret the responses to Question 3 as the sign of the cross-learning from expected inflation to expectations of the state of the real economy.

4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

We can divide the possible answers into four categories. First, options 1 and 2 concern past experienced price rises. Options 3 and 4 are direct information about inflation. Options 5-8 concern other macroeconomic variables, either current or expected, and options 9 and 10 are extras. A rational household may well use the information sources in options 1,2 and 5-8 to forecast inflation, but in the decomposition in Proposition 2 this would represent cross-learning from information about other variables. The only answers that represent the use of direct information about inflation are options 3 and 4.

Question 2f was only asked in 2016Q1, but very similar questions were asked at other times. In each, the respondent is asked about the information sources they used to arrive at their expectation of inflation (either over the next 12 months or a longer horizon), or the information sources that led them to change their expectation over the last 12 months, again for expectations of 12-month ahead inflation or longer-term inflation. For each of these questions I construct a dummy variable equal to 1 if the respondent reports using direct information about inflation, and equal to 0 if they do not. The full details of each question, and the options representing direct information on inflation, are in Appendix B.

I combine these different dummy variables to give one overall indicator for if the respondent used direct information on inflation in forming their expectations, that is whether $\delta^h_{\pi n,t} > 0$. This indicator is observed in 8 different quarters between 2009Q1-2019Q1.

In Appendix C.1 I confirm that these measures of information and subjective models correlate with planned household choices, and that the signs of these correlations are consistent with the measures picking up the desired elements of household beliefs. For that I use further questions from the survey, which concern how consumption plans respond to expected inflation. For details of these auxiliary questions see Appendix C.1.

The other questions used in this section are much more standard, asking households to give point estimates for "how prices have changed over the last twelve months" and "how much would you expect prices in the shops generally to change over the next twelve months". For each question, respondents choose from a list of ranges, and follow-up questions may then asked with more precise ranges, until the respondent has selected an inflation rate bin from the set:

$$\mathbb{E}_t^h \pi_s \in \{\leq -5\%, (-5\%, -4\%], (-4\%, -3\%], (-3\%, -2\%], (-2\%, -1\%], (-1\%, 0\%), 0\%, (0\%, 1\%), [1\%, 2\%), [2\%, 3\%), [3\%, 4\%), [4\%, 5\%), [5\%, 6\%), [6\%, 7\%), [7\%, 8\%), [8\%, 9\%), [9\%, 10\%), \geq 10\%\}$$

For the exercises in section 3.4, I code perceptions and expectations at the midpoint of the selected bin, with the lowest and highest bins coded as -5.5% and 10.5% respectively. In the sections below I refer to these answers as perceived and expected inflation respectively.

3.2 Fact 1: time-series properties of subjective beliefs about inflation

The first stylised fact concerns the time-series behaviour of subjective models of the effects of inflation. Figure 1 shows the proportions answering Question 3 with each subjective model of inflation over time.[21]

Figure 1: Proportions giving each answer to Question 3: "If prices started to rise faster than they are now, do you think Britain's economy would end up stronger, or weaker, or would it make little difference?"

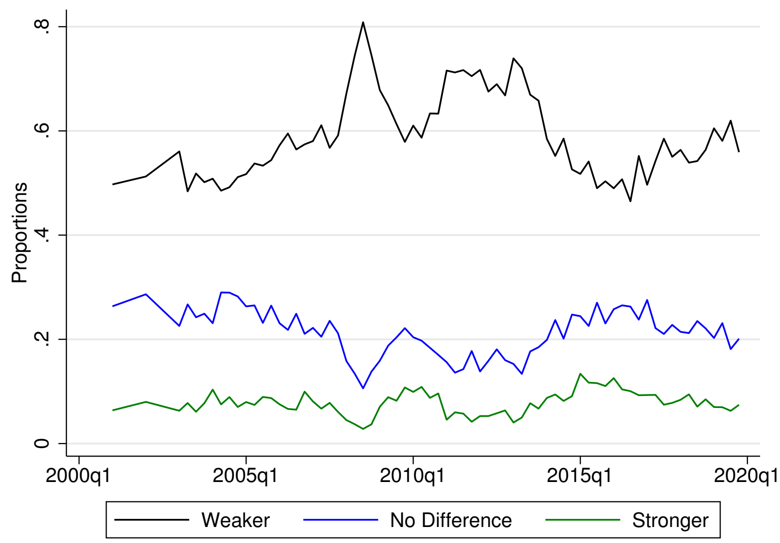

The majority of households answer that inflation would make the economy weaker in all quarters, in keeping with the findings in Shiller (1997), and more recently the

[21]The proportions shown do not add up to 1 because the proportion answering 'don't know' is omitted for figure clarity.

experimental evidence of Andre et al. (2019). However, there is substantial variation in the distribution of answers over time, which is not observable in those previous studies.

Much of this time-series variation can be explained by recent inflation experiences. UK inflation rose sharply in both 2008 and 2011, and so did the proportion of respondents saying that inflation would make the economy weaker. As inflation fell from 2013 to 2015, so did the proportion with this negative model of inflation. In fact the correlation between annual CPI inflation and the proportion of respondents with negative models of inflation is extremely high, at 0.799. Tests in Appendix C.2 show that this correlation is robust to the addition of various macroeconomic controls, which themselves explain far less of the variation in the distribution of responses than realised inflation.

This leads me to my first stylised fact:

Fact 1 *A greater proportion of households believe inflation weakens the economy when realised inflation is high.*

Note that this is not what we would expect to see in a rational expectations model. The question is about the effect of one aggregate variable (inflation) on the aggregate performance of the economy. Even if households are differentially exposed to the shock, if they all had model-consistent beliefs they would all give the same answer to this question. The fact that there is heterogeneity at all is evidence that at least some household subjective models are inconsistent with rational expectations.

These patterns also suggest that the majority of households are not using New Keynesian-style models. In a textbook New Keynesian model, a rise in inflation causes the central bank to raise the nominal interest rate. If the Taylor Principle is satisfied, the real interest rate rises, so output falls. If it is not, the real rate falls, and output falls. If this is the model used by households in the survey, we should therefore see the majority of households responding that inflation would make the economy weaker before interest rates reach the Zero Lower Bound, and the majority responding that inflation would make the economy stronger once we reach the ZLB in 2009. There is little evidence for this in Figure 1, and indeed statistical tests in Appendix C.2 find no evidence of such a shift.[22]

[22]Of course, after 2009 a New Keynesian model would predict that a sufficiently large rise in inflation would lift the economy away from the ZLB and lead to higher real interest rates and lower output. However, in 2013 the Bank of England issued forward guidance that interest rates would not rise until unemployment fell below 7% (Bank of England, 2013), so it is unlikely that households were expecting them to contract in response to small rises in inflation at this time.

3.3 Fact 2: information and subjective models in the cross-section

I next turn to the cross-sectional distribution of information. Table 1 shows the results of a probit regression of the information indicator defined in section 3.1 on the respondent's subjective model of inflation, represented by their answer to Question 3. The first column shows this with quarter fixed effects but without controls, while the second also includes a range of household characteristics as controls.[23]

Table 1: Information correlates with subjective models

	(1)	(2)
end up stronger	-0.0102	-0.00827
	(0.0191)	(0.0192)
make little difference	-0.0356***	-0.0315**
	(0.0128)	(0.0129)
dont know	-0.0627***	-0.0605***
	(0.0172)	(0.0172)
Controls	None	All
Time FE	Yes	Yes
Observations	8270	8270

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the average marginal effects from estimating a probit regression of the information indicator on the responses to Question 3. The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

In both specifications, those answering that inflation makes no difference to the aggregate economy, or who don't know the effect of inflation, are significantly less likely to use information about inflation than someone who believes inflation makes the economy weaker. There is no significant difference in the probability of using direct inflation information between those holding this view and those with positive subjective models of inflation. The coefficients displayed are average marginal effects, so the probability of using direct inflation information is 3-3.5 percentage points lower for those with a neutral model of the effects of inflation than those who believe inflation weakens the economy. Over the whole population 23% of respondents use direct inflation information, so this difference is substantial.

[23]These are gender, age, class, employment status, income, education, region, and home-ownership status. Age, class, income and education are all reported in bands, and included as categorical variables. Note that even in column 1 I exclude respondents for whom income or education is missing, so the sample is the same across columns.

This leads to my second stylised fact:

Fact 2 *Households who believe inflation makes no difference to the economy acquire less information about inflation than households who believe inflation does affect the economy (in either direction).*

The information indicator is composed of answers to several different types of question. In particular, some questions concern information used to arrive at the respondent's point forecast for inflation, while others concern the information they used to update their inflation expectations over the last year. Some questions concern expected inflation over the next 12 months, while others ask about a longer horizon. In Appendix C.3 I repeat the regressions of Table 1 on subsets of the questions, and find that the results are qualitatively robust to these alternatives. The results also remain significant for all such splits that maintain a substantial sample size.

This is not consistent with models with exogenous information, as there would be no reason for information to be systematically correlated with household subjective models. It is, however, consistent with models of endogenous information acquisition, as the value of inflation information is lower for households who believe inflation makes little difference to other variables that matter for their decisions. The (broadly defined) strength of the aggregate economy is such a variable as long as households believe there is some relationship between the aggregate economy and their personal decisions, which is supported by evidence in Roth and Wohlfart (2020), among others. The implications of this link from subjective models to information acquisition are discussed further in section 4.

3.4 Fact 3: inflation expectations and perceptions vary with subjective models

Finally, I compare perceived and expected inflation across households with different subjective model beliefs. Figures 2a and 2b show the time series of perceived and expected inflation by qualitative subjective model of inflation.

There is a persistent wedge between the perceptions and expectations of the different groups. Respondents who believe inflation weakens the economy systematically perceive that inflation has been higher, and expect it to be higher over the next year, than those who believe inflation makes no difference to the economy. They, in turn, perceive and expect higher inflation than those with a positive subjective model of inflation.

The differences are large: Table 2 shows that even after controlling for the full set of available household characteristics, those with a negative model of inflation perceive that inflation has been 54 basis points higher than those with a neutral view, and 70 basis

Figure 2: Inflation perceptions and expectations by subjective model.

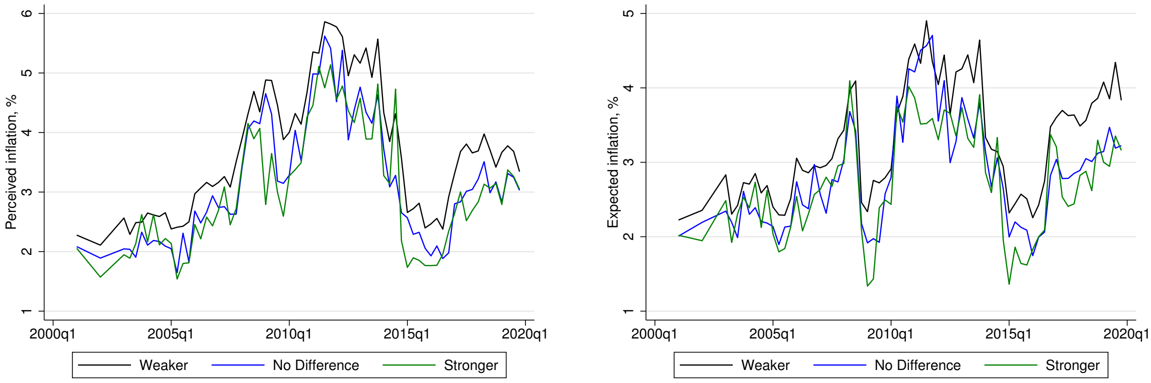

(a) Perception over the past 12 months $\mathbb{E}_t\pi_{t,t-12}$ (b) Expectation for next 12 months $\mathbb{E}_t\pi_{t+12,t}$

points higher than those with a positive view. The gaps are similarly large and strongly significant for expected inflation.

Table 2: Perceived and expected inflation are higher for those with more negative subjective models.

	(1) Perceived inflation	(2) Expected inflation
end up stronger	-0.696***	-0.565***
	(0.0371)	(0.0353)
make little difference	-0.543***	-0.466***
	(0.0226)	(0.0207)
dont know	-0.462***	-0.413***
	(0.0315)	(0.0294)
Controls	Yes	Yes
Time FE	Yes	Yes
R-squared	0.179	0.113
Observations	85803	85201

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing perceived and expected inflation on the responses to Question 3. The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

This leads to my third stylised fact:

Fact 3 *Households who believe inflation weakens the economy perceive higher inflation, and expect higher future inflation, than those with less negative subjective models.*

This is not driven by the households using different kinds of information to arrive at their perceptions and expectations: Table 1 shows that the households with positive subjective models use similar information sources to those with negative models. It is, however, consistent with subjective models being determined as worst-case beliefs under model uncertainty about the effect of inflation, as in the model in section 4. Intuitively, whenever inflation is high, the worst case model is that high inflation damages the economy.

4 Rational inattention and ambiguity aversion over inflation

In this section I present a simple model that can rationalise the empirical findings documented in section 3. The key elements are that households face Knightian uncertainty over the effect of inflation on their real wages, and that it is costly for them to process information about inflation. These features imply a two-way relationship between information and subjective models. In section 5 I explore the implications of this for aggregate macroeconomic fluctuations.

Empirical Fact 1 is explained by the effect of information on subjective models. When faced with Knightian uncertainty, ambiguity averse households make choices as if they face the worst case model (Hansen and Sargent, 2008). In this model, the worst-case belief varies with expected inflation, and so with realised information. When inflation is expected to be high, then the worst case model is one where high inflation is bad for welfare, which in this case occurs if it erodes real wages. However, when expected inflation is low, the worst case is that high inflation is good, that is if real wages increase with inflation. High current inflation implies high expected inflation on average, and so more households believe that inflation weakens real wages, and thus the aggregate economy.

Empirical Fact 2 is explained by the effect of subjective models on the information choices of rationally inattentive households facing information processing costs. A household who believes inflation has a strong effect on their real wages will adjust their consumption by more in response to inflation news than one who does not believe inflation affects their real wages. Information has a greater expected value to households who will condition their actions more strongly on its realisations,[24] so the households with neutral subjective models of inflation choose to process less information.

The link from information to subjective models that explains Fact 1 also ensures that negative models of inflation are associated with higher expected inflation, as in empirical Fact 3, since it is high expected inflation that leads to such negative models. This is strengthened further in section 4.5, where the Knightian uncertainty is extended to

[24]See Afrouzi and Yang (2021) for a similar argument, in which firms only pay for information when they are given the opportunity to adjust prices.

include the amount of bias in inflation information. The worst case for households with negative models of inflation is that the information is biased downwards, and inflation is in fact higher than the signals would suggest. This belief pushes the expectations of the households holding negative models of inflation even higher. Importantly, the extent of this confirmation bias is dependent on household attention to inflation, as bias in signals is irrelevant for expectations if the signals are extremely noisy. This effect therefore involves both directions of effects between information and subjective models.

4.1 Model setup

Time is discrete, and there are two periods. In each period, household h chooses consumption C_t^h and labour supply N_t^h. Their problem in the first period is:

$$\max_{C_t^h, N_t^h, A^h} \mathbb{E}_1^h U^h = \log(C_1^h) + \chi \log(1 - N_1^h) + \beta \mathbb{E}_1^h (\log(C_2^h) + \chi \log(1 - N_2^h)) \tag{16}$$

subject to

$$C_2^h = W_2 N_2^h + \frac{1+i}{\Pi_2} A^h, \quad C_1^h + A^h = W_1 N_1^h \tag{17}$$

where W_t is the real wage faced by all households in period t, i is the nominal interest rate on assets A^h bought in period 1, and Π_2 is gross inflation between periods 1 and 2. All wages and prices are observed in period 1, but period 2 wages and prices are unknown in period 1. While the price level in period 1 is observed, in the sections below I will denote π_1 as the rate of inflation between some initial period 0 and period 1, and I will assume that this is not necessarily observed perfectly. This assumption appears in models with agents who use a Kalman filter to update their inflation expectations (e.g. Coibion and Gorodnichenko, 2015), and is consistent with the evidence in Macaulay and Moberly (2021), who find substantial uncertainty over the rate of inflation over the past year.[25]

The first order conditions consist of two intratemporal labour supply conditions and a consumption Euler equation:

$$\frac{\chi C_t^h}{1 - N_t^h} = W_t \tag{18}$$

$$\frac{1}{C_1^h} = \beta \mathbb{E}_1^h \frac{1+i}{\Pi_2} \frac{1}{C_2^h} \tag{19}$$

To proceed, I take a log-quadratic approximation to utility, as is common in the rational inattention literature (e.g. Maćkowiak and Wiederholt, 2009). Specifically, I

[25]Since agents do observe current prices, the assumption can be thought of as embodying limited memory of past prices, as in Angeletos and Lian (2021).

substitute the labour supply conditions (18) and the present value budget constraint into expected utility $\mathbb{E}_1^h U^h$, then take a log-quadratic approximation about a steady state with $\Pi_2 = 1, 1+i = \beta^{-1}, W_1 = W_2$. Letting lower-case letters be log-deviations of the corresponding variables from steady state, the log-quadratic approximation of expected utility is:

$$\mathbb{E}_1^h \tilde{U}^h = w_1 + \beta \mathbb{E}_1^h w_2 + \frac{(1+\chi)}{\beta} \mathbb{E}_1^h \Bigg[-\frac{1+\beta}{2}(c_1^h)^2 + c_1^h(w_1 + \beta w_2 + \beta \pi_2) - \frac{1-\beta}{2}(w_1)^2 - \beta w_1 w_2 - \beta w_1 \pi_2 \Bigg] + \text{ constant} \quad (20)$$

where I have assumed that the nominal interest rate i is in steady state in period 1. See Appendix D.1 for the derivation.

A household that is fully informed about both w_2 and π_2 would choose consumption according to the first order condition:

$$c_1^{h*} = \frac{1}{1+\beta}(w_1 + \beta w_2 + \beta \pi_2) \quad (21)$$

With imperfect information, the household maximises expected utility by setting:

$$c_1^h = \mathbb{E}_1^h c_1^{h*} = \frac{1}{1+\beta}(w_1 + \beta \mathbb{E}_1^h w_2 + \beta \mathbb{E}_1^h \pi_2) \quad (22)$$

The expectations of future real wages and inflation are therefore critical in determining consumption choices. The remainder of this section studies how these expectations are formed when information processing is costly and the household is ambiguity averse over the effect of inflation on real wages.

The timing of expectation formation is as follows. The households start period 1 with some prior subjective model, which they use in making their information decisions (see section 4.3). Once the household observes the realisation of their chosen signals, they will update the parameters in their subjective model to the worst case, which will depend their newly acquired information (sections 4.4 and 4.5). The realised signals and updated subjective model will then be used to form the expectations the household will use to choose consumption (equation 22).

4.2 Subjective models

To form their expectations, households take information on current real wages and inflation and forecast forward, using their subjective models. Suppose that household h

believes real wages and inflation are determined according to:

$$w_t = \alpha^h \pi_t + \rho_w^h w_{t-1} + u_{wt} \tag{23}$$

$$\pi_t = \gamma^h w_t + \rho_\pi^h \pi_{t-1} + u_{\pi t} \tag{24}$$

where the shocks $u_{wt}, u_{\pi t}$ are both i.i.d., $N(0, \sigma_{w2}^2)$ and $N(0, \sigma_{\pi 2}^2)$ respectively. In Appendix D.2 I show that in a model completed with standard firm conditions, this form of the household's model nests the laws of motion consistent with rational expectations equilibrium.

With this subjective model, expected period-2 real wage and inflation are:

$$\mathbb{E}_1^h w_2 = \frac{\rho_w^h w_1 + \alpha^h \rho_\pi^h \mathbb{E}_1^h \pi_1}{1 - \alpha^h \gamma^h} \tag{25}$$

$$\mathbb{E}_1^h \pi_2 = \frac{\gamma^h \rho_w^h w_1 + \rho_\pi^h \mathbb{E}_1^h \pi_1}{1 - \alpha^h \gamma^h} \tag{26}$$

where the current inflation rate appears as an expectation because the household may be imperfectly informed about current inflation. I refer to $\mathbb{E}_1^h \pi_1$ as perceived inflation in the sections below, in keeping with the evidence in section 3.

For given information, households with different subjective models (different $\alpha^h, \gamma^h, \rho_w^h, \rho_\pi^h$) expect different future real wages and inflation, as shown in the general model in section 2.

4.3 Optimal information processing

The household chooses the structure and precision of the signals they will receive, s^h, to maximise expected utility. Once the signals are realised, they update their expectations of π_1, π_2, w_2, and choose consumption according to equation 22. Before the signals are realised, the expected utility loss from limited information is given by (derivation in Appendix D.3):

$$\mathbb{E}_1^h \tilde{U}^{h*} - \mathbb{E}_1^h \tilde{U}^h = \frac{\beta(1+\chi)}{2(1+\beta)} \mathbb{E}_1^h \big(w_2 + \pi_2 - \mathbb{E}_1^h (w_2 + \pi_2 | s^h)\big)^2 \tag{27}$$

That is, utility losses are proportional to the expected squared error in forecasts of $w_2 + \pi_2$, which will depend on the precision of the household's signals. In common with many papers in the rational inattention literature, I limit households to processing signals about current variables only (see Jurado (2021) for a discussion of this assumption). That is, I assume that the household cannot process any information about π_2 or w_2 directly, but can only obtain signals about variables already realised in period 1. This ensures that as the cost of information approaches zero the household information set in period 1 contains realised values of all period 1 variables, but not realisations of variables in future periods, as in standard full-information DSGE models. Since the only unobserved variable in period 1 is current inflation π_1, it is this that households process information about.

I also assume for simplicity that the household chooses information as if they are certain about the parameters of their subjective model. Similarly, they ignore that they will update those parameters after receiving information. This is akin to the anticipated utility assumption in many models with least-squares learning, where agents do not consider that their perceived law of motion will change as they observe new periods of data in the future (see Bullard and Suda (2016) for a discussion of this in the learning literature).

Substituting the subjective model-based forecasts of w_2 and π_2 (25 and 26) into the expected utility loss (27) gives expected utility losses in terms of the squared error in the household's perception of inflation:

$$\mathbb{E}_1^h \tilde{U}^{h*} - \mathbb{E}_1^h \tilde{U}^h = \frac{\beta(1+\chi)(\rho_\pi^h)^2(1+\alpha^h)^2}{2(1+\beta)(1-\alpha^h\gamma^h)^2} \mathbb{E}_1^h \big(\pi_1 - \mathbb{E}_1^h(\pi_1|s^h)\big)^2 \tag{28}$$

Importantly, the parameters $\alpha^h, \gamma^h, \rho_\pi^h$ from the household's subjective model affect the expected utility loss from inattention to inflation, because they determine how errors in $\mathbb{E}_1^h(\pi_1|s^h)$ translate into errors in forecasted w_2 and π_2, and therefore to errors in consumption. This is why subjective models will affect the household's information choices. Specifically, note that from the consumption function (22) and the forecasts of w_2 and π_2 (25 and 26) we have:

$$\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)} = \frac{\beta\rho_\pi^h(1+\alpha^h)}{(1+\beta)(1-\alpha^h\gamma^h)} \tag{29}$$

The utility loss from squared errors in perceived inflation is therefore proportional to the squared elasticity of consumption to perceived inflation:

$$\mathbb{E}_1^h \tilde{U}^{h*} - \mathbb{E}_1^h \tilde{U}^h = \frac{(1+\beta)(1+\chi)}{2\beta} \left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^2 \mathbb{E}_1^h \big(\pi_1 - \mathbb{E}_1^h(\pi_1|s^h)\big)^2 \tag{30}$$

That is, the household suffers a greater utility loss from limited information about current inflation when consumption responds more strongly to that information.

More precise information in s^h will reduces the expected squared error in the inflation perception, but following the rational inattention literature I assume that increasing signal precision is costly to the household. Specifically, the cost of a signal s^h is proportional to the Shannon mutual information between s^h and π_1:[26]

$$\psi \mathcal{I}(\pi_1; s^h) = \psi\big(H(\pi_1) - \mathbb{E}_1^h H(\pi_1|s^h)\big) \tag{31}$$

where ψ is the constant marginal cost of information, and $H(x) = -\int f(x)\log(f(x))dx$ is the entropy of a variable $x \sim f(x)$. As entropy measures the dispersion or uncertainty in a distribution, the cost is therefore increasing in the amount by which the household expects the signal realisation will reduce their uncertainty about π_1.

Household h's prior belief about current inflation is assumed to be Gaussian:

$$g^h(\pi_1) \sim N(\bar{\pi}^h, \sigma_\pi^{h2}) \tag{32}$$

With Gaussian priors, and a quadratic household objective function (28), the optimal signal about current inflation takes the form (Maćkowiak and Wiederholt, 2009):

$$s^h = \pi_1 + \varepsilon_\pi^h, \quad \varepsilon_\pi^h \sim N(0, \sigma_\varepsilon^{h2}) \tag{33}$$

With this signal structure, the information cost and posterior perception of inflation are given by:

$$\psi \mathcal{I}(\pi_1; s^h) = \frac{\psi}{2}\log_2\left(\frac{1}{1-\tau^h}\right) \tag{34}$$

$$\mathbb{E}_1^h(\pi_1|s^h) = (1-\tau^h)\bar{\pi}^h + \tau^h(\pi_1 + \varepsilon_\pi^h) \tag{35}$$

where $\tau^h = \sigma_\pi^{h2}/(\sigma_\pi^{h2} + \sigma_\varepsilon^{h2})$ is the signal-to-noise ratio of signal s^h.

The household's optimal information choice is given in Proposition 3:

Proposition 3 *The utility-maximising signal structure is as in equation 33, with σ_ε^{h2} such that:*

$$\tau^h = \begin{cases} 0 & \text{if } \left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^2 \leq \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}} \\ 1 - \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}}\left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^{-2} & \text{if } \left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^2 > \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}} \end{cases} \tag{36}$$

Proof. In Appendix D.3. ■

That is, if the elasticity of consumption to perceived inflation is close to 0, the household pays no attention to inflation. Once that elasticity is sufficiently positive or negative, perceived inflation affects decisions enough to warrant paying for some information, and

[26] The use of Shannon mutual information originates with Sims (2003), but in that paper he assumes a hard constraint on the mutual information, rather than the constant marginal cost used here and in many other applications of rational inattention (see Maćkowiak et al., 2020).

$\tau^h > 0$. As the consumption elasticity to perceived inflation grows further, attention rises and τ^h approaches 1, at which point perceived inflation is equal to the true realised π_1.[27]

This elasticity of consumption to perceived inflation is a function of the household's subjective model (equation 29). Household h processes no information if:

$$-\sqrt{\frac{\psi(1+\beta)}{\beta(1+\chi)\log(2)\sigma_\pi^{h2}}} \leq \frac{\rho_\pi^h(1+\alpha^h)}{1-\alpha^h\gamma^h} \leq \sqrt{\frac{\psi(1+\beta)}{\beta(1+\chi)\log(2)\sigma_\pi^{h2}}} \tag{37}$$

As long as $\psi > 0$, a household with $\alpha^h = -1$ will never pay attention to inflation. In such a subjective model, real wages fall one-for-one with inflation, so where nominal wages are independent of prices. In this case the substitution and income effects of higher expected inflation exactly cancel out, and period-1 consumption is unresponsive to perceived inflation.

The no-attention region is wider if prior inflation uncertainty σ_π^{h2} is lower. As has been argued elsewhere (e.g. Cavallo et al., 2017), we therefore have that households pay less attention to inflation when it is not expected to be volatile. A higher information cost also implies a wider no-attention region. Similarly, outside of the no-attention region household attention is increasing in σ_π^{h2} and decreasing in ψ.

Information choices are therefore determined by the household's subjective model, and this naturally implies the model matches empirical Fact 2. A simple proxy for 'the strength of the economy' might be aggregate consumption. If households believe others hold beliefs similar to their own, then the households who report in the survey that inflation makes no difference to the economy are those with subjective models such that $dc_1^h/d\mathbb{E}_1^h(\pi_1|s^h)$ is close to zero. They therefore process less information about inflation than those with stronger positive or negative $dc_1^h/d\mathbb{E}_1^h(\pi_1|s^h)$.

4.4 Worst case subjective models

The household faces Knightian uncertainty about the parameters of their subjective model. After the realisation of s^h is observed, the household then updates their subjective model to reflect this: following the literature on ambiguity aversion they make decisions using worst-case beliefs (Hansen and Sargent, 2008). For simplicity, in this section I assume that households are only uncertain about α^h, the effect of inflation on real wages.

Formally, the household selects beliefs and actions as if they are playing a game with an 'evil agent', who distorts α^h to minimise expected utility, while the household

[27] Although τ^h does not equal $\mathcal{I}(\pi_1; s^h)$, it is a strictly increasing function of information processed, so for ease of exposition I refer to τ^h and attention interchangeably.

simultaneously chooses c_1^h to maximise expected utility. The maximisation problem is solved by the consumption function in equation 22. Substituting this into equation 20 gives expected utility in terms of w_1 and expectations over w_2 and π_2 only:

$$\mathbb{E}_1^h \tilde{U}^h = w_1 + \beta \mathbb{E}_1^h w_2 + \frac{\beta(1+\chi)}{1+\beta} \mathbb{E}_1^h \left(\frac{1}{2} w_1^2 + \frac{1}{2} w_2^2 + \frac{1}{2} \pi_2^2 - w_1 w_2 - w_1 \pi_2 + w_2 \pi_2 \right) + \text{constant} \tag{38}$$

The expectation terms are formed by the household filtering their information on w_1 and π_1 through their subjective model. To simplify notation, I use $\mathbb{E}_1^h \pi_1$ below to mean the posterior inflation perception after information processing $\mathbb{E}_1^h(\pi_1|s^h)$. The evil agent modifies the wage equation of this model to:

$$w_t = (\alpha^h + \hat{\alpha}^h)\pi_t + \rho_w^h w_{t-1} + u_{wt} \tag{39}$$

The household does not consider every possible distortion to their subjective model. Rather, they only consider distortions in the set $\hat{\alpha}^h \in [-\alpha^*, \alpha^*]$ for some $\alpha^* > 0$, the magnitude of which refects household condifence in their central model.[28]

The minimisation problem is therefore:

$$\min_{\mathbb{E}_1^h w_2, \mathbb{E}_1^h \pi_2, \hat{\alpha}^h \in [-\alpha^*, \alpha^*]} w_1 + \beta \mathbb{E}_1^h w_2 + \frac{\beta(1+\chi)}{1+\beta} \mathbb{E}_1^h \left(\frac{1}{2} w_1^2 + \frac{1}{2} w_2^2 + \frac{1}{2} \pi_2^2 - w_1 w_2 - w_1 \pi_2 + w_2 \pi_2 \right) \tag{40}$$

subject to

$$\mathbb{E}_1^h w_2 = \frac{\rho_w^h w_1 + (\alpha^h + \hat{\alpha}^h)\rho_\pi^h \mathbb{E}_1^h \pi_1}{1 - (\alpha^h + \hat{\alpha}^h)\gamma^h} \tag{41}$$

$$\mathbb{E}_1^h \pi_2 = \frac{\gamma^h \rho_w^h w_1 + \rho_\pi^h \mathbb{E}_1^h \pi_1}{1 - (\alpha^h + \hat{\alpha}^h)\gamma^h} \tag{42}$$

The solution is summarised in Proposition 4:

Proposition 4 *The evil agent is indifferent between all $\hat{\alpha}^h \in [-\alpha^*, \alpha^*]$ when:*

$$\mathbb{E}_1^h \pi_1 = -\frac{\gamma^h \rho_w^h w_1}{\rho_\pi^h} \equiv \pi^{h*} \tag{43}$$

If $w_1 \in (-w^*, w^*)$ and $\alpha^* < \bar{\alpha}$ (constants defined in Appendix D.5),[29] *then in the region*

[28] This simple constraint allows the main insights to be explored clearly. However, it does not strictly correspond to the entropy-based distortion set in Hansen and Sargent (2008). The specific form of the permissable set does not however alter the qualitative results, as shown in Appendix D.4.

[29] These parameter restrictions are weak, since this is a linearised model, so it only applies in the region close to steady state, where w_1 is small. See Appendix D.5 for details.

around $\mathbb{E}_1^h \pi_1 = \pi^{h*}$*, the evil agent selects:*

$$\hat{\alpha}^h = \begin{cases} \alpha^* & \text{if } \mathbb{E}_1^h \pi_1 < \pi^{h*} \\ -\alpha^* & \text{if } \mathbb{E}_1^h \pi_1 \geq \pi^{h*} \end{cases} \tag{44}$$

Proof. In Appendix D.5. ■

That is, when perceived inflation is low, the worst-case belief involves distorting the effect of inflation on real wages α^h upwards: the worst case is that low inflation is associated with low real incomes. However, when perceived inflation is higher, the reverse is true. The worst case is that inflation erodes real wages, and so the ambiguity averse household distorts their subjective model in that direction, with a negative $\hat{\alpha}^h$.

This immediately matches empirical Fact 3. Equation 44 implies that those with higher inflation perceptions will bias their subjective model towards negative α, so will be more likely to answer that inflation makes the economy weaker. These higher perceptions also translate to higher expected future inflation for those households.[30]

It also matches empirical Fact 1. Recall from section 4.3 that:

$$\mathbb{E}_1^h(\pi_1 | s^h) = (1 - \tau^h)\bar{\pi}^h + \tau^h(\pi_1 + \varepsilon_\pi^h) \tag{45}$$

For all households with $\tau^h > 0$, higher realised inflation π_1 will, on average, imply higher perceived inflation. Through equation 44 this will imply that more households bias their beliefs about α down, and so more households will believe that inflation weakens the economy when realised π_1 is high.

However, these results are strengthened if the household also faces Knightian uncertainty over the structure of their signals. This is studied in the next section.

4.5 Adding uncertainty over signal bias

In this section I add a second source of Knoghtian uncertainty: the household is uncertain if the signals they have processed about inflation are biased. To maintain simplicity, I assume that they believe the data generating process for s^h is:

$$s^h = \pi_1 + b^h + \varepsilon_\pi^h, \quad \varepsilon_\pi^h \sim N(0, \sigma_\varepsilon^{h2}) \tag{46}$$

That is, they know the distribution of the noise term ε_π^h, as this was chosen by them at the start of the period. Only the mean of s^h is uncertain. As in the previous section, I assume that the household makes their information choice as if they are certain of their

[30] Proof in Appendix D.5.

central model, not taking into account that this model will be updated. In this case any bias b^h is irrelevant for information choices, as the household can always construct $\tilde{s}^h = s^h - b^h$, which is unbiased. I therefore assume without loss of generality that the household's central model has $b^h = 0$.[31]

However, if they are uncertain about b^h, then after observing the realisation of s^h they update their beliefs about b^h away from this central model to $\hat{b}^h$, towards the worst case signal bias. With this distortion, the household's perceived inflation becomes:

$$\mathbb{E}_1^h(\pi_1|s^h,\hat{b}^h) = (1-\tau^h)\bar{\pi}^h + \tau^h(s^h - \hat{b}^h) = \mathbb{E}_1^h(\pi_1|s^h,0) - \tau^h\hat{b}^h \tag{47}$$

That is, distorting b^h distorts perceived inflation in proportion with τ^h, the signal-to-noise ratio of signals chosen by the household at the start of period 1. The first result from this is therefore that for households with $\tau^h = 0$, uncertainty about signal bias has no effect on their perceived inflation, and so has no effect on their expected utility.

For households who pay some attention to inflation, however, the worst-case belief will involve some distortion to b^h. I make the simple assumption here that as with the bias to α^h, the household considers a set range of distortions, in this case that $\hat{b}^h \in [-b^*, b^*]$. In this case the distortion to b^h only affects the worst-case belief about α^h through $\mathbb{E}_1^h\pi_1$, so Proposition 4 continues to hold.[32]

The evil agent problem is therefore as in equations 40-42, except that the evil agent now also minimises expected utility with respect to $\mathbb{E}_1^h\pi_1$ and $\hat{b}^h$, subject to equation 47 and $\hat{b}^h \in [-b^*, b^*]$.

The solution is given by Proposition 5:

Proposition 5 *For $w_1 < \bar{w}$ and $\alpha^* \in (\underline{\alpha}, \bar{\alpha})$ (constants defined in Appendix D.6), there exists a $\pi^{hb} \in (\pi^{h*} - \tau^h b^*, \pi^{h*} + \tau^h b^*)$ such that in the neighbourhood of $\mathbb{E}_1^h(\pi_1|s^h,\hat{b}^h) = \pi^{h*}$ the evil agent selects:*

$$(\hat{\alpha}^h, \hat{b}^h) = \begin{cases} (\alpha^*, b^*) & \text{if } \mathbb{E}_1^h(\pi_1|s^h,0) < \pi^{hb} \\ (-\alpha^*, -b^*) & \text{if } \mathbb{E}_1^h(\pi_1|s^h,0) \geq \pi^{hb} \end{cases} \tag{48}$$

Proof. In Appendix D.6. ■

[31] Strictly, the form of s^h in equation 33 is the optimal signal structure given the quadratic objective and Gaussian prior belief, so the optimal signal is unbiased. Uncertainty over b^h could arise if the household constructs this optimal signal from various information sources, but is uncertain about the bias in some of those disparate sources.

[32] Note that in the entropy constraint used by Hansen and Sargent (2008), the constraint would be over a combination of distortions to b^h and to α^h, and so the evil agent would trade off these two margins of distortion. I abstract from this here to maintain tractability, but Appendix D.4 shows that numerically the qualitative results continue to hold with this alternative constraint.

That is, if the perceived inflation under the central model is sufficiently high, then the worst case belief involves biasing α^h downwards, so that high inflation is associated with low real wages, as in Proposition 4. In addition, the worst case compounds this by choosing $\hat{b}^h = -b^*$, so the household believes that their inflation signals are biased downwards, which would imply that inflation is in fact even higher than their signals initially suggested. Conversely, if perceived inflation under the central model is sufficiently low, the worst case involves $\hat{\alpha}^h = \alpha^*$ and $\hat{b}^h = b^*$, so low inflation is associated with low real wages, and perceived inflation is biased downwards.

This implies a kind of confirmation bias: households with high prior beliefs about inflation have higher $\mathbb{E}_1^h(\pi_1|s^h, 0)$ on average (equation 35), and so are more likely to believe that their signals are biased downwards, and inflation is higher than the signals suggest.

The first implication of this for the distribution of beliefs is that it strengthens the relationship between subjective models and perceived (and expected) inflation. Figure 3 shows perceived inflation against the worst-case subjective model $\alpha^h + \hat{\alpha}^h$ for a simulation of 10000 households, comparing the model with α^h distortions only with the model including b^h distortions.[33]

Figure 3: Binned scatter plot of perceived inflation without any $\hat{b}^h$ distortion to signal bias (black) and with $\hat{b}^h$ distortions (blue), against the worst-case subjective model $\alpha^h + \hat{\alpha}^h$.

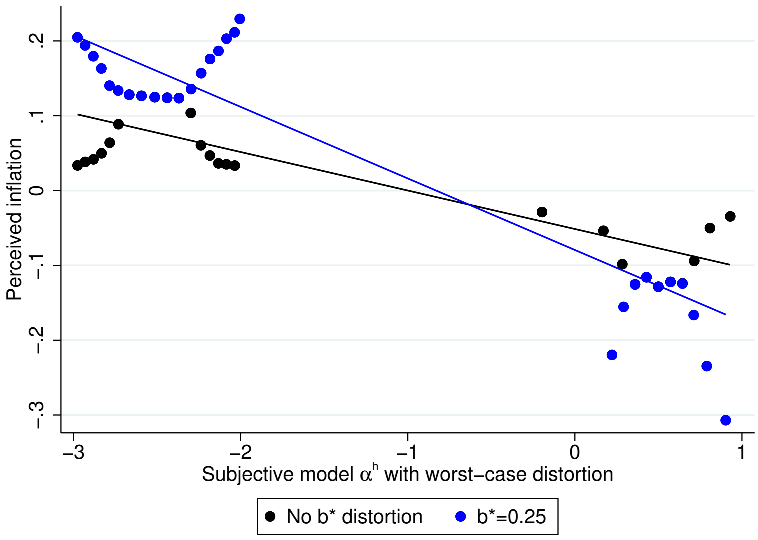

Note: Data from a simulation of 10000 households with parameters listed in Appendix E.

Even in the case without the b^* distortion (plotted in black), the households with

[33] Expected inflation naturally follows similar patterns, as it is heavily influenced by perceived inflation in this model.

more negative subjective models of inflation have higher inflation perceptions on average, through the mechanisms discussed in section 4.4. This relationship is substantially steeper, however, in the case with Knightian uncertainty over signal bias b^h (plotted in blue). Specifically, the simulation has $\alpha^* = 1.5$, so the households with distorted models close to $\alpha^h + \hat{\alpha}^h = -2.5$ and $\alpha^h + \hat{\alpha}^h = 0.5$ are those who start out with central models close to $\alpha^h = -1$. Under those central models, the optimal response to inflation is very small, as the households expect the income and substitution effects of inflation to offset each other. Those households therefore pay little attention to inflation, and beliefs about signal bias have little to no impact on perceptions. Away from those central models, however, households are more attentive, and in that case households with high perceived inflation distort their beliefs higher, while those with low perceptions distort them even lower.

The second implication is that perceived and expected inflation is increasing in household attention among households with $\hat{\alpha}^h = -\alpha^*$, because a higher τ^h implies a greater distortion to $\mathbb{E}_1^h \pi_1$ from uncertainty over the bias of household signals b^h. Conversely, perceived and expected inflation are decreasing in attention among those with $\hat{\alpha}^h = \alpha^*$, for the same reason. This relationship is shown in Figure 4. In Appendix C.4, I test this implication in the IAS survey data, and find support: perceived and expected inflation are increasing in information, but only among those with a negative subjective model of inflation.

This relationship between attention, subjective models, and expectations is important for the way the economy responds to shocks, through the narrative heterogeneity channel identified in section 2. More specifically, in section 5.2 I explore the effects of a transitory rise in realised inflation. The relationships discussed here imply that high inflation expectations may become 'baked in' (as often discussed by policymakers and the media). More importantly, they are most likely to become 'baked in' specifically among those who react most negatively to such high expectations.

Figure 4: Binned scatter plots of perceived inflation with $\hat{b}^h$ distortions against attention τ^h, by worst-case model beliefs.

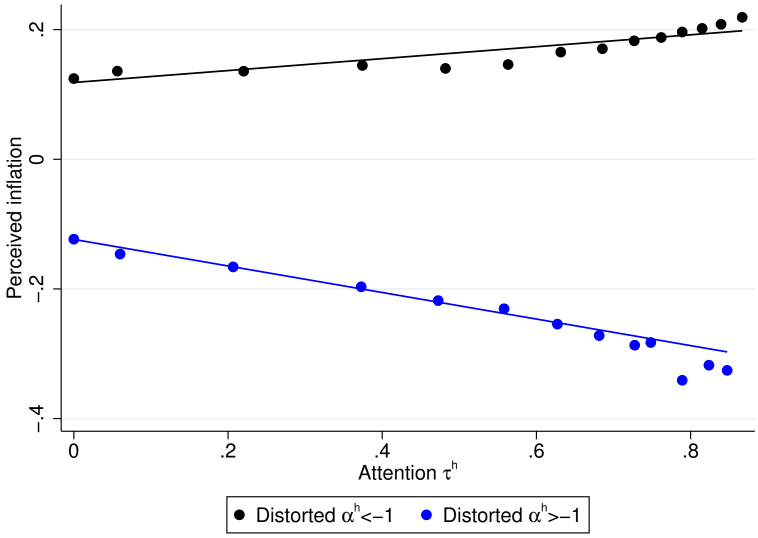

Note: Households with worst-case beliefs such that $\alpha^h + \hat{\alpha}^h < -1$, who decrease consumption when perceived inflation rises, are plotted in black. Households with $\alpha^h + \hat{\alpha}^h > -1$, who increase consumption when perceived inflation rises, are plotted in blue. Data from a simulation of 10000 households with parameters listed in Appendix E. The calibration is such that $\alpha^h + \hat{\alpha}^h < -1$ if and only if $\hat{\alpha}^h < -\alpha^*$.

5 Implications of narrative behaviour

In this section I show that the feedback between information and subjective models has important implications for macroeconomic behaviour. I first show that when information choices depend on subjective models of inflation, as explored in section 4.3, there arises a selection effect because the informed households are the households who respond to information. This reduces the impact of information frictions at the aggregate level, much as the selection effect in menu cost pricing models reduces the aggregate effect of price stickiness (Golosov and Lucas, 2007). This also has implications for the interpretation of information-based survey experiments (e.g. Coibion et al., 2019a).

Second, I show that in a simple dynamic extension of the model in section 4, transitory increases in inflation can become 'baked into' expectations, particularly among those with negative subjective models of inflation. This implies that shock transmission is history-dependent: the same shock will have different effects if it arrives when inflation has recently been high, compared to arriving at steady state.

5.1 Selection in attention

First, consider the implications of the effect of subjective models on information choice, as discussed in section 4.3. To isolate this, assume to start with that $\alpha^* = b^* = 0$, so the only heterogeneity in subjective models is that present at the start of period 1, when the households choose their information.

Consider a shock that increases inflation in period 1. The partial equilibrium effect of this on the consumption of household h is given by:

$$\frac{\partial c_1^h}{\partial \pi_1} = \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \frac{\partial \mathbb{E}_1^h \pi_1}{\partial \pi_1} = \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \tau^h \tag{49}$$

where the second equality follows from equation 35.

The partial equilibrium response of aggregate consumption is therefore:

$$\frac{\partial \bar{c}_1}{\partial \pi_1} = \int_0^1 \omega^h \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \tau^h dh = \int_0^{q_0} \omega^h \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \tau^h dh \tag{50}$$

where ω^h is a weight on household h as in equation 9. The fraction of households who pay no attention ($\tau^h = 0$) is $1 - q_0$, and they are assumed to be indexed by $h \in [q_0, 1]$.

To see how the relationship between information and subjective models affects aggregate outcomes, compare this to a model in which all households have the same signal-to-noise ratio $\bar{\tau}$, equal to the average τ^h from the baseline model:

$$\bar{\tau} = \mathbb{E}_H(\tau^h) = \mathbb{E}_H(\tau^h | \tau^h > 0) \cdot q_0 \tag{51}$$

This, for example, could reflect an economist calibrating a model with homogeneous information frictions to micro-level evidence on household information. In such a model the aggregate partial equilibrium response of consumption to the inflation shock can be split into two integrals:

$$\begin{aligned} \left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{\tau^h = \bar{\tau}} &= \int_0^1 \omega^h \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \bar{\tau} dh \\ &= \int_0^{q_0} \omega^h \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \tau^h \frac{\bar{\tau}}{\tau^h} dh + \int_{q_0}^1 \omega^h \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \bar{\tau} dh \end{aligned} \tag{52}$$

The first of these two integrals is identical to the expression for $\partial \bar{c}_1 / \partial \pi_1$ in the baseline model with endogenous attention, except that each household's response is weighted by $\bar{\tau}/\tau^h$. Relative to the baseline model, the consumption responses of more attentive households receive a lower weight, while less attentive households are over-weighted. The definition of $\bar{\tau}$ (equation 51) implies that the re-weighting for the mean household within

$h \in [0, q_0]$ is $q_0 \leq 1$. The down-weighting of attentive households is therefore stronger if a greater proportion of the population is inattentive (q_0 is lower), as this brings down $\bar{\tau}$.

The second integral overweights the consumption responses of inattentive households. In the baseline model, their response to perceived inflation does not appear, because their inattention means their inflation perceptions are unaffected by the shock. Here, however, they do appear, as it is assumed that their perceptions react to the shock with elasticity $\bar{\tau}$. Relative to the baseline model with endogenous attention, this alternative with homogeneous attention therefore under-weights the most attentive households, and over-weights the least attentive.

This leads to systematic differences in aggregate consumption responses, because the most attentive households have such high τ^h in the baseline model because they have the largest consumption responses to perceived inflation. Formally, the difference between the aggregate consumption responses in the model with homogeneous attention and the endogenous-τ^h baseline is:

$$\left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{\tau^h=\bar{\tau}} - \frac{\partial \bar{c}_1}{\partial \pi_1} = \mathbb{E}_H\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}(\bar{\tau} - \tau^h)\right) = -Cov\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}, \tau^h\right) \tag{53}$$

The difference therefore depends on the covariance of information and subjective models: by making attention exogenous, the homogenous-τ^h model presented here omits the narrative heterogeneity channel of shock transmission discussed in section 2. This covariance depends here on the distribution of subjective models, as τ^h is increasing in the absolute value of $\partial c_1^h / \partial \mathbb{E}_1^h \pi_1$. Among households with $\partial c_1^h / \partial \mathbb{E}_1^h \pi_1 > 0$, the covariance of consumption responses and τ^h is positive, but among those with $\partial c_1^h / \partial \mathbb{E}_1^h \pi_1 < 0$ it is negative.

This implies that the partial equilibrium aggregate consumption response to the shock is always (weakly) closer to zero in the homogeneous-τ^h model. If most households increase consumption when perceived inflation rises, then the baseline aggregate consumption response to a π_1 increase is positive. At the same time, the narrative heterogeneity channel is positive, so expression 53 is negative. Conversely, if most households have strong negative subjective models of inflation, and so decrease consumption in response to higher perceived inflation, then the baseline aggregate response is negative. The covariance of consumption responses and attention is also negative, so expression 53 is positive.

Figure 5 shows this effect graphically. It plots the consumption response of an individual household to a shock to π_1 against the same household's response to an increase in perceived inflation $\mathbb{E}_1^h \pi_1$. If households observed inflation precisely, this would simply be the 45° line (red dashed line).

Figure 5: Consumption response to a change in actual inflation against response to perceived inflation.

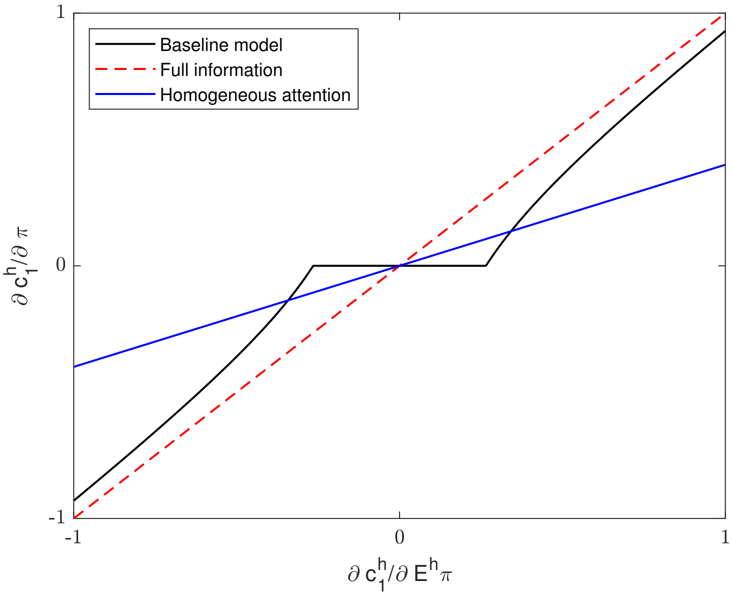

Parameters: $\beta = 0.95, \chi = 1, \psi = 0.2, \sigma_\pi^{h2} = 1$, distribution of subjective models such that $\bar{\tau} = 0.4$.

The black solid line shows this relationship in the baseline model with endogenous τ^h. Households with $\partial c_1^h/\partial \mathbb{E}_1^h \pi_1$ close to zero pay no attention to current inflation, and so their perceptions of inflation do not change when the shock hits. They therefore do not react. Households with greater $\partial c_1^h/\partial \mathbb{E}_1^h \pi_1$ pay more attention, and so their perceptions are more sensititve to the inflation change, and their elasticity of consumption to π_1 is therefore closer to the 45° line.

If the endogenous τ^h is replaced by a fixed $\bar{\tau}$ for all households, the elasticity of c_1^h to π_1 is instead given by the blue solid line. The consumption response to the shock is larger in magnitude than in the baseline case for all households with $\partial c_1^h/\partial \mathbb{E}_1^h \pi_1$ such that $\tau^h < \bar{\tau}$ in the baseline model. Conversely, the consumption response is reduced towards zero for all those who are more attentive than average in the baseline model. The narrative heterogeneity channel amplifies the effect of the shock on aggregate consumption, so removing it by assuming all households have the same information frictions weakens the effect of the shock.

As a simple example, consider the case where all attentive households $h \in [0, q_0)$ have subjective models such that they all pay the same amount of attention, $\tau^h = \tau$.[34] Suppose also that all inattentive households $h \in [q_0, 1]$ have $\partial c_1^h/\partial \mathbb{E}_1^h \pi_1 = 0$. In this case this final group do not respond to the shock whatever assumption is made about their

[34] From Proposition 3, these households have the same τ as long as they all have the same $(\partial \bar{c}_1/\partial \mathbb{E}_1^h \pi_1)^2$, so they may be split between those with positive and negative consumption responses to $\mathbb{E}_1^h \pi_1$.

information, so we have:

$$\frac{\partial \bar{c}_1}{\partial \pi_1} = \tau \int_0^{q_0} \omega^h \frac{\partial \bar{c}_1}{\partial \mathbb{E}_1^h \pi_1} dh = \tau \cdot \left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{FI} \tag{54}$$

$$\left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{\bar{\tau}} = \bar{\tau} \int_0^{q_0} \omega^h \frac{\partial \bar{c}_1}{\partial \mathbb{E}_1^h \pi_1} dh = q_0 \tau \cdot \left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{FI} \tag{55}$$

where $\left.\frac{\partial \bar{c}_1}{\partial \pi_1}\right|_{FI}$ is the aggregate consumption response if all households had full information about π_1. The more inattentive households there are (i.e. the lower is q_0), the more this reduces the average attention $\bar{\tau}$. This reduces the consumption response in the baseline endogenous-τ model somewhat, as there are fewer households with a non-zero elasticity of consumption to perceived inflation. However, in the homogeneous-τ model, there is an extra effect, because the increase in inattentive households reduces the average attention $\bar{\tau}$, reducing the predicted reactions of the households with large $\partial \bar{c}_1 / \partial \mathbb{E}_1^h \pi_1$.

This effect is analogous to the selection effect in menu cost models of price setting (Caplin and Spulber, 1987; Golosov and Lucas, 2007). In those models, the aggregate price level is much less sticky than the average of firm-level stickiness, because the firms who change their prices in equilibrium do so because their current prices are a long way from the optimal price. Price adjustments are therefore disproportionately drawn from firms desiring large price changes.

In the model presented here, households obtaining information about inflation are disproportionately drawn from those who would react strongly to that information. Just as the price level in a menu cost is more flexible than the average flexibility at the firm level, this implies that aggregate consumption is more responsive to inflation than is implied by micro-level estimates of household attention. The narrative heterogeneity channel of inflation can therefore explain why representative agent models typically require only small information frictions to match aggregate data (Maćkowiak and Wiederholt, 2015), while micro-level studies find very large degrees of inattention (Link et al., 2021).[35]

Implications for information treatment experiments: a growing number of recent papers have used information treatments in surveys to uncover the causal effect of expectations on a variety of outcomes (see Candia et al., 2020, for a review). These experiments involve providing a random subset of the respondents to a survey with some information, then tracking how the expectations and/or actions of those respondents change relative to those who were not provided the information. Experiments of this kind in macroeconomics have used, for example, information on current aggregate or local

[35] Afrouzi and Yang (2021) explore a related mechanism to this, finding that firms acquire information only in periods when they are changing prices, which is when their expectations matter for the dynamics of the price level.

variables (Armona et al., 2019; Coibion et al., 2020), professional forecasts of aggregate variables (Roth and Wohlfart, 2020; Coibion et al., 2021a), details of central bank policy (Coibion et al., 2019b), and the average expectations of other respondents (Coibion et al., 2018, 2021b). The key advantage of these randomised controlled trials is that the variation in treatment is completely exogenous, so the resulting estimates reveal the causal effect of the information.

However, the narrative heterogeneity channel explored here provides a caveat to the application of these estimates to the study of aggregate behaviour. The standard approach to estimate the causal effect of an information treatment is to regress the outcome variable on the expectation of interest, instrumented using an indicator for whether the respondent was in the treatment or control group.[36] The estimate is therefore consistent for the local average treatment effect on those who update their expectations as a result of the information provision, and is most influenced by those who update the furthest.

The dependence of information choice on subjective models implies that these compliers are not representative of the population as a whole. Those who respond the most strongly to the provision of extra information are those with the most uncertain beliefs, that is the respondents who have been least attentive to that variable. At the extreme, attentive households will already have seen the information provided, so will not update at all. Proposition 3 implies that these inattentive compliers, who drive the estimates from the survey experiment, are inattentive because their actions are particularly inelastic to information.

The survey estimates will therefore recover the causal effect of expectations among those who respond the least to expectations (see Appendix F for a formal treatment). Importantly, this section has shown that this is not the effect that will matter when shocks hit the economy. Rather, it is the most attentive households who matter most for aggregate responses. This therefore provides a note of caution for those using the results of survey experiments to calibrate models of aggregate behaviour.[37]

In addition, this offers an explanation for why, when studying the effects of inflation expectations on consumption, the survey experiment in Coibion et al. (2019a) finds very different effects than the natural experiment studied in D'Acunto et al. (2021), even to the extent of obtaining different signs of the main estimates. The natural experiment

[36] It is also common to use a second instrument, the interaction of the treatment indicator with the agent's prior expectation (e.g. Coibion et al., 2019a). This does not substantially change the intuition discussed here.

[37] It is worth noting that in some cases the response of inattentive agents is precisely what the researcher needs, and so the response to aggregate shocks would be inappropriate. For example, when studying central bank communication with the general public the researcher is typically interested in how previously inattentive households will respond to the provision of more accessible information (Coibion et al., 2019b; Haldane et al., 2021).

studied is a pre-announced increase in VAT that occurred in Germany in 2005. As a shock hitting future inflation, the estimated effects of the policy will have been driven by the most attentive households,[38] while the information experiment in Coibion et al. (2019a) estimates the response among the inattentive. With attention choices driven by differences in subjective models, that imply differences in the elasticity of consumption to expectations, it is unsurprising that the two studies arrive at different conclusions.

5.2 History-dependent shock transmission and baked-in expectations

I now return to the two-way feedback between information and subjective models. The interaction implies that the transmission of inflation shocks to aggregate consumption depends on the recent history of realised inflation, as this alters the distribution of subjective models and information choices going into each period. After a transitory spike in inflation, as was arguably seen with supply bottlenecks as the US and other coutnries relaxed coronavirus restrictions in mid-2021 (Powell, 2021), subjective models shift unequally across households, implying substantial changes in the effects of any subsequent shocks.

To explore these effects, I begin by showing how the aggregate consumption response to an inflation shock depends on the distribution of prior beliefs about inflation, before extending the two-period model to a dynamic setting, in which past realised inflation affects that distribution.

The effect of prior beliefs: suppose that households are divided between G discrete groups. Each group g comprises a fraction q_g of households, with $\sum_{g\in G} q_g = 1$. Within each group all households also share the same σ_π^{h2}, γ^h, ρ_π^h and ρ_w^h. Since G can be arbitrarily large, this is without loss of generality.

In section 5.1 I assumed that $\alpha^* = b^* = 0$, but here I restore the role of worst-case beliefs by assuming both of these maximum distortions are strictly positive. With the households split into groups, the partial equilibrium response of aggregate consumption to an inflation shock is:

$$\frac{\partial \bar{c}_1}{\partial \pi_1} = \int_0^1 \omega^h \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| \hat{\alpha}^h, \hat{b}^h \right) \tau^h dh = \sum_{g\in G} q_g \tau^g \mathbb{E}_H \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| g, \hat{\alpha}^h, \hat{b}^h \right) \tag{56}$$

where the second equality uses the fact that τ^h is constant within each household group (equation 36). Notive that now the response of c_1^h to perceived inflation is dependent

[38] D'Acunto et al. (2021) argue that as a highly-publicised policy announcement the majority of German households were aware of the shock. However, even if the shock was universally observed, it would still imply different estimates from survey experiments, which are driven by inattentive households.

on the worst-case model distortions $\hat{\alpha}^h, \hat{b}^h$, as households decide consumption using their distorted models.[39]

The consumption response to perceived inflation is not however constant within groups, because the households receive idiosyncratic realisations of their signals about current inflation, and have possibly different prior means $\bar{\pi}^h$. Both of these mean that the households will have heterogeneous inflation perceptions under their central models $\mathbb{E}_1^h(\pi_1|s^h, 0)$, and so some will distort their subjective models towards inflation increasing real wages and upward bias in signals (α^*, b^*), while others will distort their subjective models in the other direction $(-\alpha^*, -b^*)$ (see Proposition 5). Incorporating distorted expectations (equations 41 and 42) into the consumption function (22) and differentiating shows that such differences in $\hat{\alpha}^h$ affect the consumption response to perceived inflation:

$$\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle|\hat{\alpha}^h, \hat{b}^h\right) = \frac{\beta\rho_\pi^h(1+\alpha^h+\hat{\alpha}^h)}{(1+\beta)(1-(\alpha^h+\hat{\alpha}^h)\gamma^h)} \tag{57}$$

Specifically, those with $\mathbb{E}_1^h(\pi_1|s^h, 0)$ sufficiently high that they distort α^h downwards have lower (or more negative) consumption responses to perceived inflation than those with the same central model who distort in the opposite direction:[40]

$$\begin{aligned}\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle| -\alpha^*, \hat{b}^h\right) &= \frac{\beta\rho_\pi^h(1+\alpha^h-\alpha^*)}{(1+\beta)(1-(\alpha^h-\alpha^*)\gamma^h)} \\ &< \frac{\beta\rho_\pi^h(1+\alpha^h+\alpha^*)}{(1+\beta)(1-(\alpha^h+\alpha^*)\gamma^h)} = \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle|\alpha^*, \hat{b}^h\right)\end{aligned} \tag{58}$$

Intuitively, a lower $\hat{\alpha}^h$ means that any rise in perceived inflation is associated with less of a rise in expected future real wages, which implies less of an increase in consumption.

The consumption response to perceived inflation is constant among households with the same central model (in the same group) and the same worst-case distortion $\hat{\alpha}^h$. Denoting P_g as the proportion of households in group g with $\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb}$, the average consumption elasticity to perceived inflation in household group g is therefore:

$$\mathbb{E}_H\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle|g, \hat{\alpha}^h, \hat{b}^h\right) = \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle|g, -\alpha^*, \hat{b}^h\right)P_g + \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle|g, \alpha^*, \hat{b}^h\right)(1-P_g) \tag{59}$$

Substituting this into equation 56, the response of aggregate consumption to an in-

[39]This was also true in section 5.1, but was irrelevant because α^* and b^* had been set to 0, so there were no worst-case distortions to subjective models.

[40]This inequality holds as long as $\gamma^h > -1$. See Appendix D.6 for a discussion of this weak restriction.

flation shock is:

$$\frac{\partial \bar{c}_1}{\partial \pi_1} = \sum_{g \in G} q_g \tau^g \Bigg(\bigg(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \bigg| g, -\alpha^*, \hat{b}^h \bigg) P_g + \bigg(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \bigg| g, \alpha^*, \hat{b}^h \bigg) (1 - P_g) \Bigg) \tag{60}$$

The transmission of the shock to aggregate consumption therefore depends on the proportions within each group of households with high perceived inflation. It is through these P_g terms that the transmission of the shock can become history dependent.

For an individual household h, the probability that their inflation perception exceeds π^{hb} is:

$$\Pr(\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb}) = \Pr((1 - \tau^h)\bar{\pi}^h + \tau^h(\pi_1 + \varepsilon_\pi^h) > \pi^{hb}) \tag{61}$$

If the household is completely inattentive ($\tau^h = 0$), this is equal to 1 if $\bar{\pi}^h > \pi^{hb}$, and 0 otherwise. For a household with $\tau^h > 0$, however, the noise in s^h affects perceptions, so using the definition of τ^h to write the variance of ε_π^h in terms of σ_π^{h2} and τ^h:

$$\begin{aligned} \Pr(\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb}) &= \Pr(\varepsilon_\pi^h > \frac{\pi^{hb} - (1 - \tau^h)\bar{\pi}^h - \tau^h \pi_1}{\tau^h}) \\ &= 1 - \Phi\bigg(\frac{\pi^{hb} - (1 - \tau^h)\bar{\pi}^h - \tau^h \pi_1}{\sigma_\pi^h \sqrt{\tau^h(1 - \tau^h)}} \bigg) \end{aligned} \tag{62}$$

where $\Phi(.)$ is the CDF of the standard normal distribution.

This is increasing in both realised inflation π_1 and household priors $\bar{\pi}^h$. Defining an indicator $\mathbb{1}_{hb}$ equal to 1 if household h has $\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb}$, and equal to 0 otherwise, for any group with $\tau^g > 0$ we have:[41]

$$P_g = \mathbb{E}_H(\mathbb{1}_{hb}|g) = \int_{\bar{\pi}^h} \bigg(1 - \Phi\bigg(\frac{\pi^{gb} - (1 - \tau^g)\bar{\pi}^h - \tau^g \pi_1}{\sigma_\pi^g \sqrt{\tau^g(1 - \tau^g)}} \bigg) \bigg) f(\bar{\pi}^h|g) d\bar{\pi}^h \tag{63}$$

Equations 60 and 63 therefore characterise the aggregate consumption response to an inflation shock in terms of the central models of each group, and the distribution of prior means $\bar{\pi}^h$ within groups. From equation 62 we can see that a distribution featuring more households with higher prior means $\bar{\pi}^h$ implies a higher P_g, because $\Pr(\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb})$ is increasing in $\bar{\pi}^h$.[42]

If the distribution of prior means shifts up in all groups simultaneously, all of the P_g

[41] Since π^{hb} is defined as the $\mathbb{E}_1^h(\pi_1|s^h, 0)$ such that expected utility is the same between distortions $(-\alpha^*, -b^*)$ and (α^*, b^*), it is independent of the household's realised $\mathbb{E}_1^h(\pi_1|s^h, 0)$, so is equal across households within a group g.

[42] Formally, consider two distributions such that $f_A(\bar{\pi}^h|g)$ has first order stochastic dominance over $f_B(\bar{\pi}^h|g)$. Since $\Pr(\mathbb{E}_1^h(\pi_1|s^h, 0) > \pi^{hb})$ is strictly increasing in $\bar{\pi}^h$, this implies that $P_g(f_A) > P_g(f_B)$.

terms increase, and the consumption response to inflation shocks changes according to:

$$d\left(\frac{\partial \bar{c}_1}{\partial \pi_1}\right) = \sum_{g\in G} q_g \tau^g \left[\left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle| g, -\alpha^*, \hat{b}^h\right) - \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1}\middle| g, \alpha^*, \hat{b}^h\right)\right] dP_g < 0 \tag{64}$$

where the inequality follows from equation 58.

That is, if the distribution of prior means shifts upwards, the distribution of perceived inflation shifts up to match, for any level of realised inflation. That, in turn, means that the worst case belief for a greater proportion of households is that inflation harms real wages. When realised inflation rises, this more negative distribution of subjective models implies a lower aggregate consumption response.

It should be noted that b^* has not featured in these equations. This is because, while it affects the level of c_1^h through perceived inflation, fed through the household's subjective models to expected future inflation and real wages, this effect does not vary with the true inflation rate π_1. It therefore has no effect on the contemporaneous reaction of consumption to an inflation shock. It will, however, have an effect in a dynamic setting, because perceived inflation in one period will affect prior beliefs in the next.

Dynamic model: to study the time-variation in shock transmission brought about by the interaction of endogenous information and subjective model beliefs, I study a simple dynamic extension of the model developed above. The model is made dynamic by repeating the two-period model many times. That is, define a period of the dynamic model as composed of two sub-periods, which correspond to the two periods of the model in section 4. At the end of each period, all households therefore consume all of their savings from the first sub-period. The only thing connecting periods is therefore the households' information.

Specifically, if household h believes that the process driving inflation is persistent between periods, then their prior belief about inflation in the first sub-period of t will depend on information they processed in previous periods. Since this is a linear model, the perceived law of motion for inflation in the first sub-period can be written:

$$\pi_{1t+1} = \tilde{\rho}^h \pi_{1t} + u_{t+1} \tag{65}$$

where u_{t+1} is independent of π_{t+1} and has mean 0.[43] To ensure that household priors remain Gaussian, as assumed in equation 33 in the two-period model, I impose further that $u_{t+1} \sim N(0, \sigma_u^{h2})$. The mean and variance of period-$t+1$ priors are therefore:

$$\bar{\pi}_{t+1}^h = \mathbb{E}_{1t}^h \pi_{1t+1} = \tilde{\rho}^h \mathbb{E}_{1t}^h(\pi_{1t} | s_t^h, \hat{b}_t^h) \tag{66}$$

[43] Restricting $\mathbb{E}u_{t+1} = 0$ is equivalent to assuming that all households know that inflation is stable about $\bar{\Pi}$, the point about which the model is linearised.

$$\sigma_{\pi t+1}^{h2} = (\tilde{\rho}^h)^2 Var_t^h(\pi_{1t}|s_t^h, \hat{b}_t^h) + \sigma_u^{h2} = (\tilde{\rho}^h)^2 \sigma_{\pi t}^{h2}(1-\tau_t^h) + \sigma_u^{h2} \tag{67}$$

Substituting out for τ_t^h using equation 36, the prior variance becomes:

$$\sigma_{\pi t+1}^{h2} = \frac{(\tilde{\rho}^h)^2 \psi \beta}{(1+\beta)(1+\chi)\log(2)} \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \bigg| 0, \hat{b}^h \right)^{-2} + \sigma_u^{h2} \tag{68}$$

This is therefore constant over time, and only varies across households due to variation in their perceived noise in the inflation process σ_u^{h2}, and their central subjective models. Through equation 36, this in turn implies that τ_t^h is also constant over time. I therefore drop the t subscript from these objects in the remaining equations.

The prior mean, however, may vary over time and with the state of the world. Substituting out for posterior perceived inflation using equation 47, we have:

$$\bar{\pi}_{t+1}^h = \tilde{\rho}^h \left(\bar{\pi}_t^h + \tau^h(\pi_{1t} - \bar{\pi}_t^h + \varepsilon_{\pi t}^h - \hat{b}_t^h) \right) \tag{69}$$

This shows that a high realisation of inflation in period t will increase prior expected inflation in the next period, with the strongest effects on the most attentive households. This implies that the transmission of shocks is history-dependent: high inflation in period t feeds into higher prior beliefs about inflation in period $t+1$, which as shown in equation 64 above implies a lower response of aggregate consumption. Shocks to inflation therefore have less positive, and more negative, effects on aggregate consumption when inflation has recently been high.

In a simple way this mechanism therefore generates the phenomenon often discussed by policymakers and commentators of temporary high inflation leading to high expectations becoming 'baked in', and it demonstrates one consequence of that for the transmission of aggregate shocks. However, this direct impact of π_1 only generates persistence in inflation priors, and so inflation perceptions and expectations, equal to $\tilde{\rho}^h \tau^h$. If we interpret expectations becoming 'baked in' to mean that they remain high for longer than would be suggested by the true inflation process,[44] then this direct effect cannot generate that unless households use a subjective law of motion in which $\tilde{\rho}^h$ is much higher than the true persistence of inflation.[45]

However, there is a second effect in equation 69 which also leads to further history-dependence of perceived inflation, coming through the worst-case distortions b^h. From Proposition 5 we know that households select a positive distortion to b^h if their perceived

[44]This interpretation fits with the mentions of 'baking in' during the rise in inflation in the US in 2021, when policymakers largely saw the inflation increase as transitory, but worried expectations might remain persistently high despite this (Waller, 2021).

[45]See Angeletos et al. (2020) and Macaulay and Moberly (2021) for evidence on this and further discussion of its implications.

inflation is sufficiently low, and a negative distortion if it is high. Households with high perceived inflation therefore see their perceptions rise further, following the interpretation of b^h as implying a kind of confirmation bias. For these households, this leads to higher prior means in the next period. Similarly, those with low perceptions reduce them further, and so reduce $\bar{\pi}^h_{t+1}$. In this way the b^h distortion spreads out the distribution of prior means in the following period, and this effect is larger for more attentive households, who rely more on their signals for forming their perceptions of inflation, so who are more affected by confirmation bias. Conditional on realised inflation and the period-t prior the distribution of period-$t+1$ prior means is:

$$\Pr(\bar{\pi}^h_{t+1} = \pi | \pi_{1t}, \bar{\pi}^h_t) = \begin{cases} \phi\left(\frac{\pi - \tilde{\rho}^h((1-\tau^h)\bar{\pi}^h_t + \tau^h(\pi_{1t} - b^*))}{\tilde{\rho}^h \sigma^h_\pi \sqrt{\tau^h(1-\tau^h)}}\right) & \text{if } \pi \leq \tilde{\rho}^h(\pi^{hb} - \tau^h b^*) \\ 0 & \text{if } \tilde{\rho}^h(\pi^{hb} - \tau^h b^*) < \pi \leq \tilde{\rho}^h(\pi^{hb} + \tau^h b^*) \\ \phi\left(\frac{\pi - \tilde{\rho}^h((1-\tau^h)\bar{\pi}^h_t + \tau^h(\pi_{1t} + b^*))}{\tilde{\rho}^h \sigma^h_\pi \sqrt{\tau^h(1-\tau^h)}}\right) & \text{if } \pi > \tilde{\rho}^h(\pi^{hb} + \tau^h b^*) \end{cases} \tag{70}$$

where $\phi(.)$ is the PDF of the standard normal distribution.

From this, we see that more households will adjust their perceived inflation upwards in period t, and so will have higher prior means in period $t+1$, when perceived inflation is high. This effect therefore amplifies the direct effect of realised π_{1t} on future prior means. Temporarily high realised inflation implies that more households select worst-case beliefs which feature even higher inflation, as a result of the confirmation bias effect. This in turn increases their priors about inflation in the next period, implying a 'baking in' of high inflation perceptions and expectations beyond the perceived persistence of inflation $\tilde{\rho}^h$. Equation 64 shows that this has knock-on effects on the transmission of shocks. In a general equilibrium model, it could of course also change the future dynamic of inflation. [To add: section with general equilibrium implications]

These effects can be seen in Figures 6a and 6b, which plot the evolution of the mean $\bar{\pi}^h$ and $\partial \bar{c}_1 / \partial \pi_1$ after a one-off positive shock to realised inflation, in a simulation of 10000 households. In both plots I compare the case with only direct effects of higher realised inflation ($b^* = 0$, in black) with the case including the confirmation bias effect ($b^* > 0$, in blue).

The average prior belief shifts up when the shock hits both with and without the indirect effect through $\hat{b}^h$ distortions. However, prior means increase by more on average with the indirect effect, and importantly they stay persistently higher than before the shock, unlike the case with direct effects only. This implies that the elasticity of aggregate consumption to further inflation shocks also remains persistently lower than before the shock. This confirmation bias mechanism therefore implies that temporarily high inflation

Figure 6: Time series of mean $\bar{\pi}^h$ and $\partial \bar{c}_1 / \partial \pi_1$ after a one-off positive shock to realised inflation in period 0.

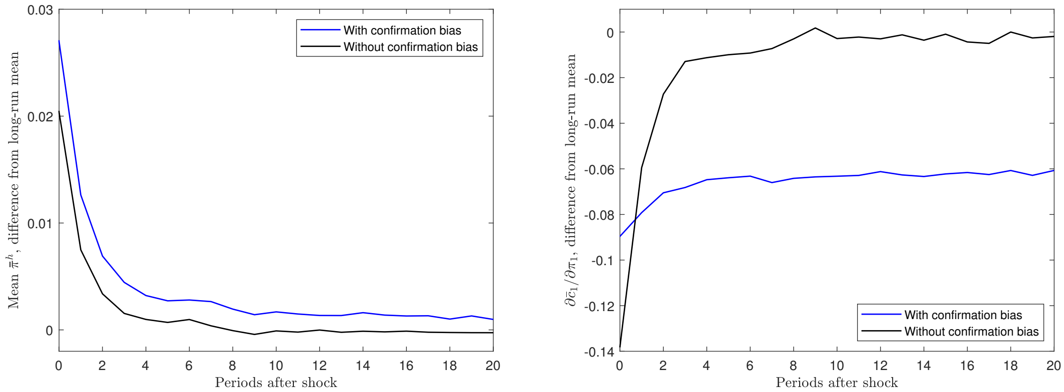

(a) Mean $\bar{\pi}^h$ (b) $\dfrac{\partial \bar{c}_1}{\partial \pi_1}$

Note: Data from a simulation of 10000 households with parameters listed in Appendix E. All variables are plotted relative to the mean value obtained after a burn-in of 100 periods with no shocks.

can become 'baked in' to expectations and perceptions, and this lowers the response of aggregate consumption to any future inflation shocks.

In each period, the response of consumption to inflation $\partial \bar{c}_1 / \partial \pi_1$ can be decomposed into the average narrative and narrative heterogeneity channels as follows:

$$\begin{aligned} \frac{\partial \bar{c}_1}{\partial \pi_1} &= \mathbb{E}_H \left(\tau^h \cdot \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| \hat{\alpha}^h, \hat{b}^h \right) \\ &= \mathbb{E}_H(\tau^h) \mathbb{E}_H \left(\frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| \hat{\alpha}^h, \hat{b}^h \right) + Cov_H \left(\tau^h, \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| \hat{\alpha}^h, \hat{b}^h \right) \end{aligned} \tag{71}$$

Figures 7a and 7b show the contributions of these channels to the path of $\partial \bar{c}_1 / \partial \pi_1$ after the same one-off inflation shock, for the cases with and without the confirmation bias effect.

Both channels fall more persistently when households distort their perceived inflation ($b^* > 0$). Importantly, the narrative heterogeneity channel remains low because the confirmation bias distortion has most impact on households who are the most attentive. The baking in of higher inflation expectations is therefore strongest amongst those who are most attentive, which is the households who react strongly to inflation. With the shock pushing more households toward negative subjective models of how inflation affects real incomes, this implies a larger fall in the transmission of further inflation shocks than implied by the average narrative. For a policymaker concerned about future positive inflation surprises, high inflation gets baked into the expectations of the worst possible households.

Figure 7: Time series of total $\partial\bar{c}_1/\partial\pi_1$, the average narrative, and the narrative heterogeneity channel, after a one-off positive shock to realised inflation in period 0.

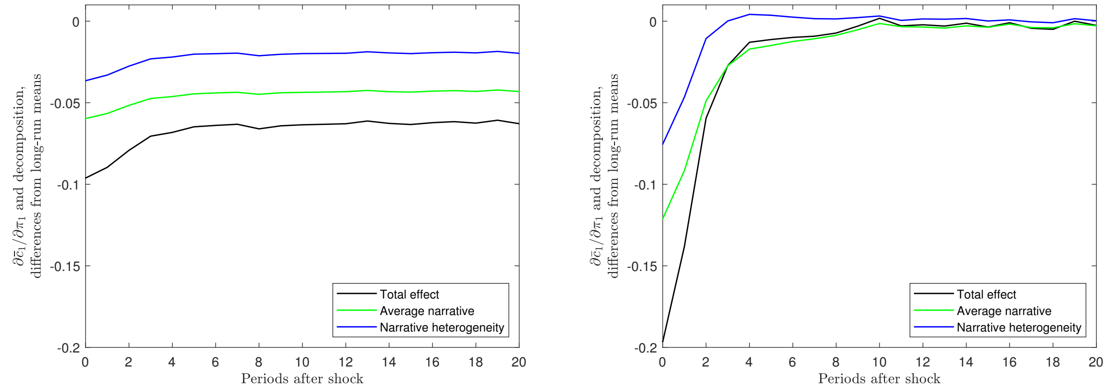

(a) With confirmation bias ($b^* > 0$) (b) Without confirmation bias ($b^* = 0$)

Note: Data from a simulation of 10000 households with parameters listed in Appendix E. All variables are plotted relative to the mean value obtained after a burn-in of 100 periods with no shocks.

For simplicity, in this dynamic setting I have assumed that household central models remain constant over time. If this was relaxed, we would see a further channel through which this 'baking in' of inflation expectations impacts the transmission of shocks. If a household's posterior model in period t affects their central model in period $t+1$, then high inflation, which leads more households to distort α^h down, would lead to lower average α^h in central models in the next period. Households who started out with $\alpha^h > -1$ would on average increase consumption by less after future increases in perceived inflation, and so would pay less attention. Conversely, households starting with $\alpha^h < -1$ would react more negatively to perceived inflation, and so would pay more attention. This would naturally make the narrative heterogeneity term $Cov_H\left(\tau^h, \frac{\partial c_1^h}{\partial \mathbb{E}_1^h \pi_1} \middle| \hat{\alpha}^h, \hat{b}^h\right)$ more negative. It would also mean that the confirmation bias becomes stronger in households with negative subjective models of inflation, and weaker in those with positive models. Even more so than in Figures 7a and 7b, expectations of high inflation would therefore get baked in the most among households with the strongest negative reaction to inflation.[To add: include this in the formal model]

6 Conclusion

This paper explores the interaction of the two components of expectation formation, information and subjective models, which previous literature has tended to treat separately.

I show that in a general log-linear model, shocks pass through to aggregate responses along three channels. The first, the average narrative, is the transmission that would be seen in a representative agent model. The second and third, however, relate to the heterogeneity in subjective models and in information. In particular, the narrative heterogeneity channel operates when information and subjective models covary systematically within the population of agents.

I use unique features of the Bank of England Inflation Attitudes Survey to document that subjective models and information about inflation to covary systematically with each other, and with inflation perceptions and expectations. Subjective models also vary systematically with realised inflation. These results suggest that the transmission of shocks to aggregate consumption is affected by the interaction of information and subjective models, and changes over time as a result.

Finally, I propose a model with rational inattention and model uncertainty as one way to explain the empirical findings. The model generates a selection effect on information, history-dependent shock transmission, and the possibility that expectations of high inflation may become 'baked in' after large transitory inflation spikes.

References

Afrouzi, H. and Yang, C. (2021). Selection in Information Acquisition and Monetary Non-Neutrality.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2019). Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample.

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

Angeletos, G.-M. and La'O, J. (2010). Noisy Business Cycles. *NBER Macroeconomics Annual*, 24(1):319–378.

Angeletos, G.-M. and Lian, C. (2018). Forward Guidance without Common Knowledge. *American Economic Review*, 108(9):2477–2512.

Angeletos, G.-M. and Lian, C. (2021). Determinacy without the Taylor Principle. *NBER Working Paper Series*, (28881).

Angeletos, G.-M. and Sastry, K. (2021). Inattentive Economies.

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

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

Bank of England (2013). *Monetary policy trade-offs and forward guidance*. London.

Bank of England (2020). Inflation Attitudes Survey Methodology and notes - online survey.

Bénabou, R., Falk, A., and Tirole, J. (2018). Narratives, Imperatives, and Moral Reasoning. *National Bureau of Economic Research*.

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

Bilbiie, F. O. (2019). The New Keynesian cross. *Journal of Monetary Economics*.

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

Brock, W. A. and Hommes, C. H. (1997). A Rational Route to Randomness. *Econometrica*, 65(5):1059–1095.

Broer, T., Kohlhas, A., Mitman, K., and Schlafmann, K. (2020). Heterogenous Information Choice in General Equilibrium.

Bullard, J. and Mitra, K. (2002). Learning about monetary policy rules. *Journal of Monetary Economics*, 49(6):1105–1129.

Bullard, J. and Suda, J. (2016). The stability of macroeconomic systems with Bayesian learners. *Journal of Economic Dynamics and Control*, 62:1–16.

Candia, B., Coibion, O., and Gorodnichenko, Y. (2020). Communication and the Beliefs of Economic Agents. *NBER Working Paper Series*, 27800.

Caplin, A. S. and Spulber, D. F. (1987). Menu Costs and the Neutrality of Money. *The Quarterly Journal of Economics*, 102(4):703.

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *The Quarterly Journal of Economics*, 118(1):269–298.

Carvalho, V. M. and Tahbaz-Salehi, A. (2019). Production Networks: A Primer. *Annual Review of Economics*, 11:635–663.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

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

Coibion, O., Georgarakos, D., Gorodnichenko, Y., and van Rooij, M. (2019a). How Does

Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial. *National Bureau of Economic Research Working Paper Series.*

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

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

Coibion, O., Gorodnichenko, Y., Kumar, S., and Ryngaert, J. (2021b). Do You Know that I Know that You Know…? Higher-Order Beliefs in Survey Data. *The Quarterly Journal of Economics*, 136(3):1387–1446.

Coibion, O., Gorodnichenko, Y., and Ropele, T. (2020). Inflation Expectations and Firm Decisions: New Causal Evidence. *The Quarterly Journal of Economics*, 135(1):165–219.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2019b). Monetary Policy Communications and their Effects on Household Inflation Expectations. *NBER Working Paper Series*, 25482.

D'Acunto, F., Hoang, D., and Weber, M. (2021). Managing Households' Expectations with Unconventional Policies. *The Review of Financial Studies.*

Del Negro, M., Giannoni, M., and Patterson, C. (2015). The forward guidance puzzle. *Staff Reports 574, Federal Reserve Bank of New York.*

Domm, P. (2021). Markets could be challenged in the week ahead as the Fed prepares to reverse easy policy. *CNBC, 2021, Oct. 29.*

Eliaz, K. and Spiegler, R. (2020). A Model of Competing Narratives. *American Economic Review*, 110(12):3786–3816.

Eusepi, S. and Preston, B. (2011). Expectations, Learning, and Business Cycle Fluctuations. *American Economic Review*, 101(6):2844–72.

Eusepi, S. and Preston, B. (2018). The science of monetary policy: An imperfect knowledge perspective. *Journal of Economic Literature*, 56(1):3–59.

Farhi, E. and Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, 109(11):3887–3928.

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

Galí, J. and Gambetti, L. (2009). On the sources of the great Moderation. *American Economic Journal: Macroeconomics*, 1(1):26–57.

Galí, J. and Gambetti, L. (2015). The Effects of Monetary Policy on Stock Market Bubbles: Some Evidence. *American Economic Journal: Macroeconomics*, 7(1):233–257.

Gibbons, R. and Prusak, L. (2020). Knowledge, Stories, and Culture in Organizations. *AEA Papers and Proceedings*, 110:187–92.

Golosov, M. and Lucas, R. E. (2007). Menu Costs and Phillips Curves. *Journal of Political Economy*, 115(2):171–199.

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

Haldane, A. (2020). Avoiding Economic Anxiety. *Speech given at the Cheshire and Warrington LEP Economic Summit Webinar*, 30/09/2020.

Haldane, A., Macaulay, A., and McMahon, M. (2021). The Three E's of Central-Bank Communication with the Public. In Pasten, E. and Reis, R., editors, *Independence, Credibility, and Communication of Central Banking*, pages 279–342. Central Bank of Chile.

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press.

Hubert, P. and Ricco, G. (2018). Imperfect Information in Macroeconomics. *Revue de l'OFCE*, 137:181–196.

Ilut, C. L. and Schneider, M. (2014). Ambiguous Business Cycles. *American Economic Review*, 104(8):2368–99.

Iovino, L. and Sergeyev, D. (2021). Central Bank Balance Sheet Policies without Rational Expectations.

Jurado, K. (2021). Rational Inattention in the Frequency Domain.

Kamdar, R. (2019). The Inattentive Consumer: Sentiment and Expectations.

Kohlhas, A. and Robertson, D. (2021). A Theory of Rational Caution.

Larsen, V. H., Thorsrud, L. A., and Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics*, 117:507–520.

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

Macaulay, A. (2021). The Attention Trap : Rational Inattention , Inequality , and Fiscal Policy. *European Economic Review*, 135:103716.

Macaulay, A. and Moberly, J. (2021). Subjective laws of motion for inflation.

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

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

Maćkowiak, B. and Wiederholt, M. (2015). Business Cycle Dynamics under Rational Inattention. *Review of Economic Studies*, 82(4):1502–1532.

Michelacci, C. and Paciello, L. (2020). Ambiguity Aversion and Heterogeneity in Households' Beliefs.

Molavi, P. (2019). Macroeconomics with Learning and Misspecification : A General Theory and Applications.

Nunes, R. and Park, D. (2020). Inflation Expectations, Interest Rates, and Consumption Behavior.

Paul, P. (2020). The time-varying effect of monetary policy on asset prices. *Review of Economics and Statistics*, 102(4):690–704.

Powell, J. H. (2021). Monetary Policy in the Time of COVID. *Speech given at the "Macroeconomic Policy in an Uneven Economy," economic policy symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming (via webcast), 27/08/2021.*

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, 72(3):821–852.

Reis, R. (2006). Inattentive Producers. *The Review of Economic Studies*, 73(3):793–821.

Roth, C. and Wohlfart, J. (2020). How do expectations about the Macroeconomy affect personal expectations and Behavior? *Review of Economics and Statistics*, 102(4):731–748.

Shiller, R. J. (1997). Why Do People Dislike Inflation? In Romer, C. D. and Romer, D. H., editors, *Reducing Inflation: Motivation and Strategy*, pages 13–70. University of Chicago Press.

Shiller, R. J. (2017). Narrative Economics. *American Economic Review*, 107(4):967–1004.

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

Spiegler, R. (2020). Behavioral implications of causal misperceptions. *Annual Review of Economics*, 12:81–106.

Spiegler, R. (2021). A Simple Model of Monetary Policy under Phillips-Curve Causal Disagreements.

Waller, C. J. (2021). The Economic Outlook and a Cautionary Tale on "Idiosyncratic" Price Changes and Inflation. *Speech given at the Stanford Institute for Economic Policy Research Associates Meeting, Stanford, California (via webcast), 19/10/2021.*

Yang, C. (2019). Rational Inattention, Menu Costs, and Multi-Product Firms: Micro Evidence and Aggregate Implications. *Working paper.*

A Consumption function in a standard household problem

This derivation closely follows that in Bilbiie (2019) appendix A, and is also similar to consumption functions derived in Farhi and Werning (2019) and others.

Suppose that a household maximises:

$$\mathbb{E}_t^h \sum_{s=0}^{\infty} \beta^s \frac{(C_{t+s}^h)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} \text{ s.t. } C_{t+s}^h + B_{t+s}^h = R_{t+s-1} B_{t+s-1}^h + Y_{t+s}^h \tag{72}$$

Where C_t^h is consumption, σ is the intertemporal elasticity of substitution, B_t^h are real one-period bonds bought in period t, R_t is the gross interest rate on such a bond, and Y_t^h is real income.

The first order condition is the standard Euler equation, which when log-linearised about steady state becomes:

$$c_t^h = \mathbb{E}_t^h c_{t+1}^h - \sigma r_t \tag{73}$$

Substituting forward we obtain:

$$c_t^h = \mathbb{E}_t^h c_{t+s}^h - \sigma \sum_{k=0}^{s-1} r_{t+k} \tag{74}$$

Assuming that $b_t^h = 0$ (as it is in equilibrium in a standard representative-agent or two-agent New Keynesian model), the present value budget constraint is:

$$\mathbb{E}_t^h \sum_{s=0}^{\infty} C_{t+s}^h \prod_{k=0}^{s-1} R_{t+k}^{-1} = \mathbb{E}_0^h \sum_{t=0}^{\infty} Y_{t+s}^h \prod_{k=0}^{s-1} R_{t+k}^{-1} \tag{75}$$

Log-linearising (recalling that in steady state $R = \beta^{-1}$):

$$\sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h (c_{t+s}^h - \sum_{k=0}^{s-1} r_{t+k}) = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h (y_{t+s}^h - \sum_{k=0}^{s-1} r_{t+k}) \tag{76}$$

Use the Euler equation to substitute out for $\mathbb{E}_t^h c_{t+s}^h$ to obtain:

$$\sum_{s=0}^{\infty} \beta^s (c_t^h - (1-\sigma)\mathbb{E}_t^h \sum_{k=0}^{s-1} r_{t+k}) = \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h (y_{t+s}^h - \sum_{k=0}^{s-1} r_{t+k}) \tag{77}$$

Rearranging:

$$\begin{aligned} \frac{1}{1-\beta} c_t^h &= \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h y_{t+s}^h - \sigma \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h (\sum_{k=0}^{s-1} r_{t+k}) \\ &= \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h y_{t+s}^h - \frac{\sigma\beta}{1-\beta} \sum_{s=0}^{\infty} \beta^s \mathbb{E}_t^h r_{t+s} \end{aligned} \tag{78}$$

Multiplying through by $1-\beta$, and applying the Fisher equation $r_t = i_t - \pi_{t+1}$, we obtain equation 2.

B Defining the direct information indicator in the IAS

The full set of questions used to construct the information dummy is set out below, along with the dates at which each was asked and how the answers are mapped into the information indicator used above. All of the questions were only asked in the first quarter of the year(s) indicated. In the main exercises I exclude questions 2e and 2g from the total information variable, to ensure that there are no periods in which two questions are asked. I remove these rather than the short run questions in those periods to keep the majority of questions as short run expectations. The cross-sectional results are robust to including these extra questions, but the time series results are not, because including 2e and 2g means a much higher proportion of people are recorded as being informed in those periods, which affects the time series estimates. Only using short run questions (i.e. dropping 2018 and 2019 when Question 2civ is used) does not affect the cross-sectional results.

Question 2aiv *What were the most important factors that led you [to change (insert their response to how expectation has changed)] your expectation of prices in the shops over the next 12 months?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *The level of the exchange rate (the value of sterling)*
10. *Other factors*
11. *None*

Asked: 2017
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2civ *What were the most important factors that led you to change/not change your expectation of prices in the shops in the longer term?*

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *The level of the exchange rate (the value of sterling)*
10. *Other factors*
11. *None*

Asked: 2018, 2019
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2d *When you said prices would go up in the next 12 months, how important were the following things in getting to that answer?*

For each option, possible answers are:

- *Very important*
- *Fairly important*

- *Not very important*
- *Not at all important*
- *Don't know*
- *Refused*

Options:

1. *How prices have changed in the shops in your most recent visits (i.e. the last 1 to 6 months).*
2. *How prices have changed in the shops over the longer term (i.e. the last 12 months or more)*
3. *The current level of interest rates.*
4. *The current strength of the British Economy.*
5. *The inflation target set by the government.*
6. *Reports on inflation outlook in the media.*
7. *Reports of VAT changes in the media.*
8. *Other factor(s).*

Asked: 2009, 2010, 2011, 2013
Information indicator: =1 if 'very important' selected for option 6, =0 otherwise.

Question 2e *And which, if any, of the same factors were important in getting to your expectation of how prices will change over the longer term (say in 5 years time)?*

1. *How prices have changed in the shops in your most recent visits (i.e. the last 1 to 6 months).*
2. *How prices have changed in the shops over the longer term (i.e. the last 12 months or more)*
3. *The current level of interest rates.*
4. *The current strength of the British Economy.*
5. *The inflation target set by the government.*
6. *Reports on inflation outlook in the media.*
7. *Reports of VAT changes in the media.*
8. *Other factor(s).*

Asked: 2011, immediately after Question 2d
Information indicator: =1 if item 6 selected, =0 otherwise.

Question 2f *What were the most important factors in getting to your expectation for how prices in the shops would change over the next 12 months?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*

2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

Asked: 2016
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

Question 2g *And what were the most important factors in getting to your expectation for how prices in the shops would change over the longer term (say in 5 years' time)?*

Please select up to 4:

1. *How prices have changed in the shops recently, over the last 12 months*
2. *How prices have changed in the shops, on average, over the longer term i.e the last few years*
3. *Reports of current inflation in the media*
4. *Discussion of the prospects for inflation in the media*
5. *The level of interest rates*
6. *The inflation target set by the government*
7. *The current strength of the UK economy*
8. *Expectations about how economic conditions in the UK are likely to evolve*
9. *Other factors*
10. *None*

Asked: 2016
Information indicator: =1 if items 3 or 4 selected, =0 otherwise.

C Extra empirical results using the IAS data

C.1 The relationship of planned consumption with measured information and subjective models

To confirm that the survey measures of information and subjective models uncover meaningful aspects of household beliefs, I show how they correlate with planned consumption behaviour.

To this end, I use a question in the survey that concerns the household's consumption response to expected inflation.

Question 17.2 *Which, if any, of the following actions are you taking, or planning to take, in the light of your expectations of price changes over the next twelve months?*

- *Cut back spending and save more.*

The crucial feature of this question is that it asks about consumption choices which are explicitly driven by expected inflation.[46] A household answering 'yes' to this question, and who reports elsewhere in the survey that they expect prices to rise in the next year, is therefore indicating that $dc_t^h/d\mathbb{E}_t^h p_{t+1} < 0$.[47] A question that only asked about consumption or consumption changes, without reference to the cause of the behaviour, would conflate this with reactions to expectations of other variables, which might also be influenced by the same shocks as expected inflation, either directly or through cross-learning.

To see how the consumption response to expected price changes fits into the framework of section 2, notice that the household response to a change in the expectation of a given variable is determined by two things: the impact that has on expectations of other variables through cross-learning, and the response to each of those updated expectations.

$$\frac{dx_{kt}^h}{d\mathbb{E}_t^h z_{jt}} = \sum_{i=1}^{N} \mu_{ki,t}^h \chi_{ij,t}^h \tag{79}$$

Substituting this into equation 15, after some rearrangement, yields:

$$\frac{d\bar{x}_{kt}}{dz_{nt}} = \sum_{j=1}^{N} \left[\bar{\delta}_{jn,t} \mathbb{E}_H \left(\frac{dx_{kt}^h}{d\mathbb{E}_t^h z_{jt}} \right) + Cov_H \left(\delta_{jn,t}^h, \frac{dx_{kt}^h}{d\mathbb{E}_t^h z_{jt}} \right) \right] \tag{80}$$

Question 17.2 is informative about the sign of $\frac{dc_t^h}{d\mathbb{E}_t^h p_{t+1}}$. If current prices are assumed to be fixed by the household, then this is the same as the sign of $\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}}$. We can therefore use the relationship between answers to Question 17.2 and the inflation information indicator developed in section 3.1 as evidence on the narrative heterogeneity channel for shocks to inflation, and their effects on aggregate consumption.

[46] Question 17.1 in the survey is also about consumption, asking if the respondent will "bring forward major purchases such as furniture or electrical goods". I do not use this in my main measure of consumption responses for two reasons. First, as Nunes and Park (2020) note, the question refers specifically to durable goods, which may not respond to prices in the same way as aggregate consumption, which is the object of interest here. Second, it is very rarely chosen: just 6% of respondents said they would bring forward major purchases. In contrast, 40% report that they will cut back spending and save more. Any estimation on this variable will therefore be heavily influenced by a small subset of agents.

[47] This interpretation is discussed in more detail at the end of this appendix.

Specifically, the vast majority of respondents (98%) expect positive inflation over the next 12 months.[48] For these households, yes and no responses to Question 17.2 respectively indicate that:

$$\frac{dc_t^h}{d\mathbb{E}_t^h p_{t+1}} \begin{cases} < 0 & \text{if answer yes} \\ \geq 0 & \text{if answer no} \end{cases} \tag{81}$$

For the minority who expect deflation, these inequalities are reversed: responding with 'yes' indicates consumption is being cut because of an expected fall in prices. I therefore define the following indicator:

$$\frac{\widetilde{dc_t^h}}{d\mathbb{E}_t^h p_{t+1}} = \begin{cases} 1 & \text{if Q17.2='no' and } \mathbb{E}_t^h \pi_{t+1} > 0 \\ 0 & \text{if Q17.2='yes' and } \mathbb{E}_t^h \pi_{t+1} > 0 \\ 1 & \text{if Q17.2='yes' and } \mathbb{E}_t^h \pi_{t+1} < 0 \\ 0 & \text{if Q17.2='no' and } \mathbb{E}_t^h \pi_{t+1} < 0 \end{cases} \tag{82}$$

For the large majority who expect inflation, this is equal to 1 if $\frac{dc_t^h}{d\mathbb{E}_t^h p_{t+1}} \geq 0$, and equal to 0 if the reaction to expected price rises is strictly negative. The same is true of the minority who expect deflation, except that any household with $\frac{dc_t^h}{d\mathbb{E}_t^h p_{t+1}} = 0$ would respond 'no' to Question 17.2, and so is counted as if their response to expected price rises is strictly negative. The mislabeling is not a large issue, as less than 1% of respondents to Question 17.2 both expect deflation and answer 'no'. The results below are robust to removing the few households who expect deflation (see Table 3 column 2).

Table 3 shows how this is related to the information indicator and the subjective models (responses to Question 3). Column 1 shows the results from estimating a probit regression of $\frac{\widetilde{dc_t^h}}{d\mathbb{E}^h{}_{+1}}$ on the information indicator interacted with subjective models (Question 3), plus the standard household controls and time fixed effects used above. The coefficient on information is significantly negative for those with negative subjective models of inflation, despite the fact that we would expect $\frac{dc_t^h}{d\mathbb{E}_t^h p_{t+1}} \geq 0$ in a standard New Keynesian model. Being informed is therefore associated with a *lower* probability of responding positively to expected inflation for these households.

However, for those who believe inflation makes no difference to the economy, information is not associated with any change in the consumption response to expected inflation.

[48] The analysis in this section excludes any households who report expecting zero inflation over the next 12 months, or who do not answer the inflation expectation question, as Question 17.2 is difficult to interpret for these households. I discuss the appropriate counterfactual implicit in the question below. Including these people, 79% of respondents to Question 17.2 expect positive inflation, 7% expect zero inflation, 2% expect deflation, and 12% do not answer.

For those who believe inflation makes the economy stronger, being informed is associated with a significantly higher $\Pr(\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}} \geq 0)$.

Table 3: Consumption response to inflation correlates with information, by subjective model

	(1) c response to Eπ	(2) c response to Eπ
c response to Eπ		
information indicator=1	-0.213***	-0.224***
	(0.0611)	(0.0613)
end up stronger	0.0108	0.0392
	(0.0891)	(0.0906)
information indicator=1 × end up stronger	0.348*	0.313*
	(0.185)	(0.186)
make little difference	0.130**	0.157***
	(0.0594)	(0.0600)
information indicator=1 × make little difference	0.0240	-0.0149
	(0.126)	(0.128)
dont know	0.0958	0.0978
	(0.0833)	(0.0846)
information indicator=1 × dont know	-0.0158	-0.0342
	(0.186)	(0.187)
Expected Inflation	All	Exclude Deflation
Controls	All	All
Time FE	Yes	Yes
Observations	4940	4871

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of probit regressions of the $\widetilde{\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}}}$ indicator on the information indicator, interacted with responses to Question 3. The omitted category is a household with information indicator=0 who holds the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

This is consistent with individuals filtering information through their subjective models of the economy. If a household who believes inflation weakens the economy gets more information about future positive inflation, their subjective model implies that they should cut consumption, because bad times lie ahead. If instead a household believes inflation strengthens the economy, then they will react in the opposite way to the same inflation. The overall correlation of information and consumption response is negative because the majority of households believe inflation makes the economy weaker.

This therefore supports the claim that the information indicator and answers to Question 3 reflect the information and subjective models used by households in making their consumption decisions.

Since the majority of households in the survey hold negative subjective models of inflation, the overall correlation between the information indicator and the consumption response is negative, and substantial. The average marginal effect of the information indicator across the population is -7.1 percentage points. More information, at least on average, does not bring people more into line with the standard prescriptions of macroeconomic models for consumption-saving behaviour. The narrative heterogeneity channel therefore acts to decrease the aggregate consumption response to an inflation shock.

The analysis here assumes that when asked whether they will cut back consumption and save more, households are comparing their actions to a counterfactual in which there are no price rises over the next 12 months. An alternative possibility is that they are comparing with a consumption plan made in the past, in which case the relevant counterfactual is where expected inflation is unchanged from the level expected when the plan was made. I consider this in two ways, and find that the qualitative patterns in reported consumption responses to inflation are the same for households expecting inflation to increase or decrease relative to the previous year. It does not therefore appear that past inflation is the relevant counterfactual for most respondents.

First, column 2 of Table 3 re-runs the regression in column 1, excluding any respondent who reports expecting prices to fall over the next year. All results are qualitatively the same as over the full sample, showing that the few respondents expecting deflation are not driving the results.

Second, I split the sample by the sign of the respondent's expected change in inflation, computed as the sign of the difference between 12-month ahead inflation forecast and their perception of inflation over the previous 12 months. The results are in Table 4.[49] The sample sizes in each group are substantially smaller than over the full sample, so some significance is lost, but importantly the signs of the key coefficients remain the same. In each group, households who believe inflation makes the economy weaker are less likely to have $\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}} \geq 0$ when they get inflation information. For households who believe inflation makes the economy stronger, this effect is reversed. The similarity of these patterns suggests that most respondents use 'no price change' as the counterfactual when answering Question 17.2, not 'no inflation change'. If the latter was used, we would expect to see changes of sign across the columns in Table 4, as a household expecting a fall in inflation would be reporting $-1 \times \frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}}$, while one expecting a rise in inflation

[49]For brevity I only include the results of the regression including interaction effects, but repeating column 1 of Table 3 with this split also yields no changes in key coefficient signs.

would report $\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}}$.

Table 4: Consumption response to inflation correlates with information, by subjective model and sign of perceived Eπ change.

	(1) $E\Delta\pi < 0$	(2) $E\Delta\pi = 0$	(3) $E\Delta\pi > 0$
c response to Eπ			
information indicator=1	-0.140	-0.305***	-0.257**
	(0.116)	(0.101)	(0.107)
end up stronger	0.0668	-0.178	0.195
	(0.164)	(0.151)	(0.165)
information indicator=1 × end up stronger	0.586	0.349	0.397
	(0.441)	(0.293)	(0.307)
make little difference	0.165	0.136	0.181
	(0.111)	(0.0957)	(0.112)
information indicator=1 × make little difference	0.129	-0.300	0.113
	(0.241)	(0.211)	(0.216)
dont know	0.156	0.0293	0.0264
	(0.176)	(0.128)	(0.167)
information indicator=1 × dont know	-0.141	0.469	0.117
	(0.354)	(0.359)	(0.325)
Controls	All	All	All
Time FE	Yes	Yes	Yes
Observations	1384	1876	1463

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of probit regressions of the $\widetilde{\frac{dc_t^h}{d\mathbb{E}_t^h \pi_{t+1}}}$ indicator on the information indicator, interacted with responses to Question 3, split by the sign of the respondent's inflation expectations. The omitted category in all cases is a household with information indicator=0 who holds the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

C.2 Extra results on the time series behaviour of subjective models of inflation

Bhandari et al. (2019) also study the time series of responses to Question 3, but they conclude that households are more pessimistic about inflation when *output growth* is low. To explore this, I regress the proportion of households responding 'end up weaker' on realised annual CPI inflation and quarterly GDP growth. The results are in column 2 of Table 5. The coefficient on GDP growth is indeed significantly negative, consistent with Bhandari et al. (2019). However, the R^2 is only slightly higher than that of a regression on inflation only (column 1), so GDP growth does not account for much of the variation in survey answers. Indeed, GDP growth does not have any significant relationship with the proportion of households with a negative view of inflation outside of the four worst months of the Great Recession. This is shown in column 3, which contains the results of repeating the regression omitting 2008Q2-2009Q1, the only quarters in the sample when quarter-on-quarter GDP growth was less than -0.5%.

Table 5: Regressions of the proportion of households answering weaker to question 3 of the Inflation Attitudes Survey on aggregate variables.

	(1) Proportion weaker	(2) Proportion weaker	(3) Proportion weaker
Inflation	0.0568^{***}	0.0517^{***}	0.0501^{***}
	(0.00489)	(0.00479)	(0.00469)
GDP growth		-0.0261^{***}	-0.0110
		(0.00869)	(0.0180)
Constant	0.466^{***}	0.487^{***}	0.482^{***}
	(0.0109)	(0.0123)	(0.0152)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.615	0.647	0.554
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing the proportion of households answering Question 3 that inflation makes the economy weaker on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

Similar patterns in reverse are observed for the other answers. Tables 6-8 below repeat the regressions of Table 5, replacing the dependent variable with the proportion of respondents choosing each of the other possible answers to Question 3. For all of them, inflation accounts for a large share of the variation in survey answers, and higher inflation is associated with significantly lower proportions giving each of those answers. Higher GDP growth is associated with higher proportions on these other answers, and that relationship is marginally significant over the sample but is not significantly different from zero for any answer when excluding the worst of the Great Recession. The greater proportion believing inflation weakens the economy when inflation has been high does not therefore come exclusively from one other group.

Table 9 repeats the regressions of Table 5, replacing each variable with its first differ-

Table 6: Regressions of the proportion of households answering no difference to question 3 of the Inflation Attitudes Survey on aggregate variables.

	(1) Proportion no diff.	(2) Proportion no diff.	(3) Proportion no diff.
Inflation	-0.0292***	-0.0262***	-0.0257***
	(0.00303)	(0.00313)	(0.00314)
GDP growth		0.0150***	0.0106
		(0.00473)	(0.0107)
Constant	0.277***	0.264***	0.266***
	(0.00772)	(0.00883)	(0.0103)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.534	0.569	0.470
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing the proportion of households answering Question 3 that inflation makes no difference to the economy on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

ence. The key result of Table 5 is maintained: the proportion believing inflation weakens the economy rises as inflation rises.

To test if the distribution of beliefs about inflation shifts towards believing inflation makes the economy stronger when the economy reaches the Zero Lower Bound, I estimate an ordered probit regression of subjective models of inflation in the zero lower bound period, and a variety of controls.[50] A response that inflation makes the economy stronger is coded as the highest value, and inflation makes the economy weaker is the lowest value (I exclude the 'don't know' answers). A positive coefficient on the zero lower bound period would therefore imply a shift towards believing inflation makes the economy stronger, as we would expect if households follow a standard New Keynesian model. This is not what the results in Table 10 show: there is no significant shift towards a positive view of inflation in the ZLB period.

C.3 Extra results on the cross-sectional behaviour of information on inflation

The first three columns of Table 11 show the results of probit regressions of the inflation information dummy on subjective models, controls, and period fixed-effects, for three

[50]The first column of Table 10 has no controls, the second includes the set of household-level covariates used throughout this paper, and the third uses those household characteristics and adds inflation and GDP growth as extra controls.

Table 7: Regressions of the proportion of households answering stronger to question 3 of the Inflation Attitudes Survey on aggregate variables.

	(1) Proportion stronger	(2) Proportion stronger	(3) Proportion stronger
Inflation	-0.0123***	-0.0116***	-0.0108***
	(0.00193)	(0.00215)	(0.00221)
GDP growth		0.00346	-0.00392
		(0.00363)	(0.00646)
Constant	0.104***	0.102***	0.104***
	(0.00431)	(0.00550)	(0.00638)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.388	0.395	0.311
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing the proportion of households answering Question 3 that inflation makes the economy stronger on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

subsamples. The first only uses questions about the information used to arrive at the respondent's *change* in expected inflation, and the second uses only questions about information used to form point forecasts. The third column excludes questions relating to forecast horizons longer than 12 months, that is excluding the responses to Question 2civ. The signs of the marginal effects are the same as in the main exercise in Table 1, though they are not significant in the case of the revisions questions, as the sample size is small.[51]

The fourth and fifth columns of Table 11 repeat the regression for broader definitions of the information dummy than that used in Table 1. In the fourth column, the information indicator includes Questions 2e and 2g. Again, the results are robust, but I leave these questions out in the main analysis to ensure consistency in the number of questions per quarter, which is useful when considering the time series of information. In the final column, I extend the criteria for setting the information indicator equal to 1 in Question 2d to account for the fact that some people may be unwilling to select the highest importance box for any information source. I therefore set the information indicator to 1 if in answer to Question 2d, the respondent selects 'very important' for direct inflation information (as before), or if they do not select 'very important' for any option, but do

[51]The coefficient on 'don't know' in the first column is significantly less than the coefficient on 'stronger', so even with the small sample size the 'don't know' group get significantly less information than the most informed group.

Table 8: Regressions of the proportion of households answering no idea to question 3 of the Inflation Attitudes Survey on aggregate variables.

	(1) Proportion no idea	(2) Proportion no idea	(3) Proportion no idea
Inflation	-0.0154***	-0.0139***	-0.0135***
	(0.00249)	(0.00262)	(0.00267)
GDP growth		0.00762*	0.00428
		(0.00423)	(0.00987)
Constant	0.153***	0.147***	0.148***
	(0.00687)	(0.00757)	(0.00884)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.355	0.376	0.282
Observations	70	70	66

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing the proportion of households answering Question 3 that they have no idea how inflation affects the economy on annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

respond that four or fewer options were 'fairly important', and direct inflation information is among them. Again, the results are robust.

C.4 Perceived and expected inflation vary with information

In section 4.5 I showed that Knightian uncertainty about the effect of inflation on real wages and the bias in household signals implies that perceived and expected inflation are increasing in household attention for households whose worst case belief is distorted towards low α, that is that inflation reduces real wages. To test this in the IAS data, I regress perceived and expected inflation on the information indicator described in section 3.1. For each dependent variable, I first run the regression for the households who report negative subjective models of inflation in response to Question 3, as these are most likely to correspond to those with $\hat{\alpha}^h = -\alpha^*$. I then repeat the regression for those reporting non-negative subjective models.[52]

The results are in Table 12. Within the group with negative subjective models of inflation, both perceived and expected inflation are significantly higher among those obtaining direct information about inflation. This relationship turns negative among those with other subjective models, though this is not significant. These results are therefore in line with those of the two-period model presented in section 4.

[52] As the information indicator is not observed every quarter there are too few observations to draw conclusions from regressions on each non-negative subjective model option individually.

Table 9: Regressions of the proportion of households answering weaker to question 3 of the Inflation Attitudes Survey on aggregate variables, first differences.

	(1)	(2)	(3)
	D.Proportion weaker	D.Proportion weaker	D.Proportion weaker
D.Inflation	0.0342***	0.0368***	0.0254**
	(0.00947)	(0.0101)	(0.0116)
D.GDP growth		0.0108	0.0153*
		(0.00777)	(0.00786)
Constant	0.0000325	0.000153	-0.000207
	(0.00445)	(0.00445)	(0.00454)
Omitted quarters	None	None	2008Q2-2009Q1
R-squared	0.179	0.196	0.109
Observations	67	67	63

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing the quarterly change in the proportion of households answering Question 3 that inflation makes the economy weaker on quarterly changes in annual CPI inflation and quarter-on-quarter real GDP growth. Proportions are computed using survey weights.

Table 10: Ordered probit regressions of subjective models of inflation on whether the economy is at the zero lower bound on nominal interest rates.

	(1)	(2)	(3)
	Subjective model	Subjective model	Subjective model
Subjective model			
ZLB	-0.00801	-0.00785	-0.00513
	(0.00937)	(0.00962)	(0.00972)
Controls	None	Household	Household + macro
Observations	83526	83526	83526

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of an ordered probit regression of answers to Question 3 on an indicator for whether the UK economy was at the zero lower bound, defined as the period from 2009Q2 to the end of 2019 (end of the sample). The ordering is: "stronger", "no difference", "weaker". Those answering "no idea" are omitted. All regressions are weighted using the survey weights provided in the IAS.

Table 11: Information correlates with subjective models, split by information question type

	(1) Revision	(2) Point forecast	(3) Short horizon	(4) Extra Qs	(5) Q2d wider
end up stronger	0.0575	-0.0335	-0.0123	0.00114	-0.00126
	(0.0380)	(0.0218)	(0.0206)	(0.0196)	(0.0196)
make little difference	-0.0191	-0.0331**	-0.0392***	-0.0310**	-0.0312**
	(0.0233)	(0.0155)	(0.0141)	(0.0132)	(0.0131)
dont know	-0.0408	-0.0715***	-0.0622***	-0.0663***	-0.0472***
	(0.0297)	(0.0206)	(0.0192)	(0.0174)	(0.0180)
Controls	All	All	All	All	All
Time FE	Yes	Yes	Yes	Yes	Yes
Observations	2364	5906	6848	8306	8270

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the average marginal effects from estimating probit regressions of the information indicators constructed from subsets of the questions listed in Appendix B on the responses to Question 3. The omitted category is the belief that inflation makes the economy weaker. All regressions are weighted using the survey weights provided in the IAS.

Table 12: Information is associated with higher perceived and expected information among those with negative subjective models.

	(1) Perceived	(2) Perceived	(3) Expected	(4) Expected
Information	-0.149	-0.805	0.197	-0.0777
	(0.630)	(1.264)	(0.261)	(0.452)
Subjective Model	Negative	Non-negative	Negative	Non-negative
Controls	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
R-squared	0.0162	0.0113	0.0613	0.0596
Observations	5325	2945	5325	2945

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table reports the results of regressing perceived and expected inflation on the information indicator, split by responses to Question 3. The first and third columns are the results using those who answer that inflation would make the economy weaker, and the second and fourth columns use all other respondents. All regressions are weighted using the survey weights provided in the IAS.

D Two period model: derivations and proofs

D.1 Log-quadratic utility approximation (equation 20)

Substituting the labour supply conditions (18) into expected utility (16) gives:

$$\mathbb{E}_1^h U^h = (1+\chi)\log(C_1^h) + \beta(1+\chi)\mathbb{E}_1^h \log(C_2^h) - \chi\log(W_1) - \beta\chi\mathbb{E}_1^h \log(W_2) \tag{83}$$

The two budget constraints (17) together imply a present value budget constraint. When again the labour supply conditions (18) are used to substitute out for N_t, this becomes:

$$C_2^h = \frac{1}{1+\chi}W_2 + \frac{1}{1+\chi}\frac{1+i}{\Pi_2}W_1 - \frac{1+i}{\Pi_2}C_1^h \tag{84}$$

Substituting out for C_2^h in equation 83, and replacing each variable X_t with $\bar{X}e^{x_t}$, gives:

$$\mathbb{E}_1^h U^h = (1+\chi)c_1^h + \beta(1+\chi)\mathbb{E}_1^h \log\left(\frac{\bar{W}}{\beta(1+\chi)}(e^{w_1-\pi_2}+\beta e^{w_2}) - \frac{\bar{C}_1^h}{\beta}e^{c_1^h-\pi_2}\right) - \chi w_1 - \beta\chi\mathbb{E}_1^h w_2 + \text{constant} \tag{85}$$

where $x_t = \log(X_t/\bar{X})$ is the log-deviation of X_t from its steady state $\bar{X}$. Taking a second order Taylor expansion of this about steady state, with respect to c_1^h, w_1, w_2, π_2, gives:

$$\begin{aligned}\mathbb{E}_1^h \hat{U}^h = \mathbb{E}_1^h\Bigg[&(1+\chi)(1-\frac{\bar{C}_1^h}{\bar{C}_2^h})c_1^h + w_1 + \beta w_2 + \frac{\bar{C}_1^h\bar{W}}{\beta(\bar{C}_2^h)^2}\left(-\frac{(1+\beta)}{2}(c_1^h)^2 + c_1^h w_1 + \beta c_1^h(w_2+\pi_2)\right) \\ &+ \frac{\bar{W}}{(1+\chi)(\bar{C}_2^h)^2}\Bigg(\frac{(\beta\bar{W}-(1+\chi)\bar{C}_1^h)}{2\beta}w_1^2 - \bar{W}w_1(w_2+\pi_2) + \frac{\bar{W}-(1+\chi)\bar{C}_1^h}{2}w_2^2 \\ &+ (\bar{W}-(1+\chi)\bar{C}_1^h)w_2\pi_2\Bigg) + (\bar{W}-(1+\chi)\bar{C}_1^h)\pi_2^2\Bigg] + \text{constant}\end{aligned} \tag{86}$$

In steady state, the consumption Euler equation (19) is $\bar{C}_2^h = \bar{C}_1^h$, and the present value budget constraint (84) is $\bar{C}_1^h = \bar{W}(1+\chi)^{-1}$. Substituting these into equation 86 gives equation 20.

D.2 Subjective model nests rational expectations

Suppose that in period 1 there are two exogenous shocks, v_1 and v_2, and there are no endogenous state variables. The rational expectations solution to any linear model will therefore take the form:

$$w_1 = A_1 v_1 + A_2 v_2 \tag{87}$$

$$\pi_1 = B_1 v_1 + B_2 v_2 \tag{88}$$

And the rational expectations of period-2 variables are:

$$\mathbb{E} w_2 = C_1 v_1 + C_2 v_2 \tag{89}$$

$$\mathbb{E} \pi_2 = D_1 v_1 + D_2 v_2 \tag{90}$$

for some constants $A_x - D_x$.

If households observe π_1 precisely, the expectations derived from their subjective models (equations 23 and 24), substituting out for w_1 and π_1 using equations 87 and 88, are:

$$\mathbb{E}_1^h w_2 = \frac{\rho_w^h (A_1 v_1 + A_2 v_2) + \alpha^h \rho_\pi^h (B_1 v_1 + B_2 v_2)}{1 - \alpha^h \gamma^h} \tag{91}$$

$$\mathbb{E}_1^h \pi_2 = \frac{\gamma^h \rho_w^h (A_1 v_1 + A_2 v_2) + \rho_\pi^h (B_1 v_1 + B_2 v_2)}{1 - \alpha^h \gamma^h} \tag{92}$$

These expectations therefore have the same form as the rational expectations in equations 89 and 90. Matching coefficients between the rational and subjective model-driven expectations gives four equations in $\alpha^h, \gamma^h \rho_w^h, \rho_\pi^h$, which have the unique solution:

$$\alpha^h = \frac{A_1 C_2 - A_2 C_1}{A_1 D_2 - A_2 D_1} \tag{93}$$

$$\gamma^h = \frac{B_1 D_2 - B_2 D_1}{B_1 C_2 - B_2 C_1} \tag{94}$$

$$\rho_w^h = \frac{C_1 D_2 - C_2 D_1}{A_1 D_2 - A_2 D_1} \tag{95}$$

$$\rho_\pi^h = \frac{D_1 C_2 - D_2 C_1}{B_1 C_2 - B_2 C_1} \tag{96}$$

The subjective models in equations 23 and 24 therefore nest the rational expectations equilibrium as long as the model structure is such that:

$$\frac{A_1}{A_2} \neq \frac{D_1}{D_2}, \quad \frac{B_1}{B_2} \neq \frac{C_1}{C_2} \tag{97}$$

[To add: a specific example model with a standard firm side]

D.3 Expected utility loss from limited information (equation 27) and optimal information processing (Proposition 3)

Equation 27:

Equation 20 implies:

$$\mathbb{E}_1^h\tilde{U}^{h*}-\mathbb{E}_1^h\tilde{U}^h=\frac{1+\chi}{\beta}\bigg(-\frac{1+\beta}{2}(c_1^{h*})^2+(c_1^{h*}-c_1^h)(w_1+\beta w_2+\beta\pi_2)+\frac{1+\beta}{2}(c_1^h)^2\bigg) \quad (98)$$

Equation 22 implies:

$$w_1+\beta w_2+\beta\pi_2=(1+\beta)c_1^{h*} \quad (99)$$

Substituting this into equation 98 and rearranging gives:

$$\mathbb{E}_1^h\tilde{U}^{h*}-\mathbb{E}_1^h\tilde{U}^h=\frac{(1+\chi)(1+\beta)}{2\beta}\mathbb{E}_1^h(c_1^{h*}-c_1^h)^2 \quad (100)$$

From equation 22 we have:

$$c_1^{h*}-c_1^h=\frac{\beta}{1+\beta}(w_2+\pi_2-\mathbb{E}_1^h(w_2+\pi_2)) \quad (101)$$

Substituting this into equation 100 gives equation 27.

Proposition 3:

I first show that the expected utility loss from noisy information about inflation (28), net of information costs (31), is given by:

$$\mathbb{E}_1^h\tilde{U}^{h*}-\mathbb{E}_1^h\tilde{U}^h=\frac{(1+\beta)(1+\chi)}{2\beta}\bigg(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\bigg)^2\sigma_\pi^{h2}(1-\tau^h)+\frac{\psi}{2}\log_2\bigg(\frac{1}{1-\tau^h}\bigg) \quad (102)$$

Equation 35 implies:

$$\begin{aligned}\mathbb{E}_1^h\big(\pi_1-\mathbb{E}_1^h(\pi_1|s^h)\big)^2&=\mathbb{E}_1^h\big((1-\tau^h)(\pi_1-\bar{\pi}^h)-\tau^h\varepsilon_\pi^h\big)^2\\&=(1-\tau^h)^2\mathbb{E}_1^h(\pi_1-\bar{\pi}^h)^2+(\tau^h)^2\mathbb{E}_1^h(\varepsilon_\pi^h)^2\\&=(1-\tau^h)^2\sigma_\pi^{h2}+(\tau^h)^2\sigma_\varepsilon^{h2}\end{aligned} \quad (103)$$

where the second equality uses that $\mathbb{E}_1^h\pi_1\varepsilon_\pi^h=0$, and the third equality uses the fact that both $\pi_1-\bar{\pi}^h$ and ε_π^h have zero mean.

Rearranging the definition of τ^h gives:

$$\sigma_\varepsilon^{h2}=\sigma_\pi^{h2}\frac{1-\tau^h}{\tau^h} \quad (104)$$

Substituting this into equation 103 yields:

$$\mathbb{E}_1^h\big(\pi_1-\mathbb{E}_1^h(\pi_1|s^h)\big)^2=\sigma_\pi^{h2}(1-\tau^h) \quad (105)$$

Substituting this into equation 28, and adding the information cost from equation 34, gives equation 102.

The household information choice problem therefore reduces to choosing the signal-to-noise ratio τ^h to minimise equation 28, subject to the constraint that $\tau^h \in (0,1]$.[53]

The first order condition with respect to τ^h is:

$$\tau^h = \begin{cases} 0 & \text{if } \left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^2 \leq \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}} \\ 1 - \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}}\left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^{-2} & \text{if } \left(\frac{dc_1^h}{d\mathbb{E}_1^h(\pi_1|s^h)}\right)^2 > \frac{\psi\beta}{(1+\beta)(1+\chi)\log(2)\sigma_\pi^{h2}} \end{cases} \tag{106}$$

This completes the proof.

D.4 Worst-case model with entropy-based constraint

[To add: write up numerical results]

D.5 Proof of proposition 4

I begin with equation 43. Substituting $\mathbb{E}_1^h\pi_1 = \pi^{h*}$ into equations 41 and 42 gives:

$$\mathbb{E}_1^h w_2 = \rho_w^h w_1 \tag{107}$$

$$\mathbb{E}_1^h \pi_2 = 0 \tag{108}$$

At this level of perceived inflation, distortions to the subjective model do not therefore affect expected inflation or real wages, and so do not affect expected utility. The evil agent is therefore indifferent between all $\hat{\alpha}^h$ when perceived inflation is at π^{h*}.

I now turn to the second part of the proposition. Substituting out for $\mathbb{E}_1^h w_2$ and $\mathbb{E}_1^h\pi_2$ in expected utility (38) using equations 41 and 42 gives expected utility in terms of parameters and period-1 information. Differentiating this with respect to $\hat{\alpha}^h$ gives, after some algebraic manipulation:[54]

$$\frac{d\mathbb{E}_1^h\tilde{U}^h}{d\hat{\alpha}^h} = \frac{\beta\rho_\pi^h(\mathbb{E}_1^h\pi_1 - \pi^{h*})}{(1-(\alpha^h+\hat{\alpha}^h)\gamma^h)^3}\Omega^h \tag{109}$$

where:

$$\Omega^h = (1-(\alpha^h+\hat{\alpha}^h)\gamma^h)(1-\Lambda(1+\gamma^h)(1-\rho_w^h)w_1) \\ + \Lambda\rho_\pi^h(1+\gamma^h)(1+\alpha^h+\hat{\alpha}^h)(\mathbb{E}_1^h\pi_1 - \pi^{h*}) \tag{110}$$

[53]This is equivalent to $\sigma_\varepsilon^{h2} \in [0,\infty)$.

[54]Importantly, this involves substituting out for $\rho_\pi^h\mathbb{E}_1^h\pi_1 + \gamma^h\rho_w^h w_1 = \rho_\pi^h(\mathbb{E}_1^h\pi_1 - \pi^{h*})$.

$$\Lambda = \frac{1+\chi}{1+\beta} \tag{111}$$

The denominator of equation 109 is positive if:

$$\hat{\alpha}^h \in (-\frac{1}{|\gamma^h|} - \alpha^h, \frac{1}{|\gamma^h|} - \alpha^h) \tag{112}$$

A sufficient condition for this is that $\hat{\alpha}^h < \bar{\alpha}$, where:

$$\bar{\alpha} \equiv |\gamma^h|^{-1} - |\alpha^h| \tag{113}$$

Next, note that with the same assumption on $\hat{\alpha}^h$, for small w_1 the first term in Ω^h is positive:

$$\begin{aligned} &(1-(\alpha^h+\hat{\alpha}^h)\gamma^h)(1-\Lambda(1+\gamma^h)(1-\rho^h_w)w_1) > 0 \\ \iff &w_1 \in (-w^*, w^*) \end{aligned} \tag{114}$$

where:

$$w^* = \frac{1}{\Lambda(|1+\gamma^h|)(1-\rho^h_w)} \tag{115}$$

Close to $\mathbb{E}^h_1\pi_1 = \pi^{h*}$, the second term of Ω^h is close to 0. For $\hat{\alpha}^h < \bar{\alpha}$ and $w_1 \in (-w^*, w^*)$ we therefore have that $\Omega^h > 0$ in the region of $\mathbb{E}^h_1\pi_1 = \pi^{h*}$. From equation 109, the sign of $d\mathbb{E}^h_1\tilde{U}^h/d\hat{\alpha}^h$ is therefore equal to the sign of $\mathbb{E}^h_1\pi_1 - \pi^{h*}$.

If $\mathbb{E}^h_1\pi_1 > \pi^{h*}$, then $d\mathbb{E}^h_1\tilde{U}^h/d\hat{\alpha}^h > 0$ and expected utility is increasing in $\hat{\alpha}^h$. The evil agent therefore minimises expected utility by setting $\hat{\alpha}^h = -\alpha^*$, the most negative distortion possible. Similarly, when $\mathbb{E}^h_1\pi_1 < \pi^{h*}$, then $d\mathbb{E}^h_1\tilde{U}^h/d\hat{\alpha}^h < 0$ and the evil agent sets $\hat{\alpha}^h = \alpha^*$.

[To add: discussion of the parameter restrictions. Numerically they are very weak.]

Finally, I show that expected future inflation is unambiguously higher for the households who distort α downward, than it is for the group who set $\hat{\alpha}^h = \alpha^*$. We can rewrite equation 42 as:

$$\mathbb{E}^h_1\pi_2 = \frac{\rho^h_\pi(\mathbb{E}^h_1\pi_1 - \pi^{h*})}{1-(\alpha^h+\hat{\alpha}^h)\gamma^h} \tag{116}$$

The parameter restrictions above imply that $1-(\alpha^h+\hat{\alpha}^h)\gamma^h > 0$. Households with $\mathbb{E}^h_1\pi_1 > \pi^{h*}$, who set $\hat{\alpha}^h = -\alpha^*$, therefore have $\mathbb{E}^h_1\pi_2 > 0$, while those with $\mathbb{E}^h_1\pi_2 < \pi^{h*}$ and $\hat{\alpha}^h = \alpha^*$ have $\mathbb{E}^h_1\pi_2 < 0$.

D.6 Proof of proposition 5

First, define:

$$\bar{w} = \frac{1}{\Lambda(1-\rho_w^h)} \tag{117}$$

$$\underline{\alpha} = \frac{\Lambda(1-\rho_w^h)w_1}{1-\Lambda(1-\rho_w^h)w_1} - \alpha^h \tag{118}$$

where:

$$\Lambda = \frac{1+\chi}{1+\beta} \tag{119}$$

and recall the definition of $\bar{\alpha}$ (equation 113).

The key component of the proof is Lemma 1:

Lemma 1 *If:*

$$w_1 < \bar{w} \tag{120}$$

$$\underline{\alpha} < \alpha^* < \bar{\alpha} \tag{121}$$

then $\mathbb{E}_1^h \tilde{U}^h | \hat{\alpha}^h$ *is continuous in* $\mathbb{E}_1^h \pi_1$*, and in the neighbourhood of* $\mathbb{E}_1^h \pi_1 = \pi^{h*}$*:*

$$\begin{aligned} \left.\frac{d\mathbb{E}_1^h \tilde{U}^h}{d\mathbb{E}_1^h \pi_1}\right| \hat{\alpha}^h > 0 \quad \textit{if } \hat{\alpha}^h = \alpha^* \\ \left.\frac{d\mathbb{E}_1^h \tilde{U}^h}{d\mathbb{E}_1^h \pi_1}\right| \hat{\alpha}^h < 0 \quad \textit{if } \hat{\alpha}^h = -\alpha^* \end{aligned} \tag{122}$$

Proof. As in the proof of Proposition 4 (Appendix D.5), substituting out for $\mathbb{E}_1^h w_2$ and $\mathbb{E}_1^h \pi_2$ in expected utility (38) using equations 41 and 42 gives expected utility in terms of parameters and period-1 information. Differentiating this with respect to $\mathbb{E}_1^h \pi_1$ gives:

$$\left.\frac{d\mathbb{E}_1^h \tilde{U}^h}{d\mathbb{E}_1^h \pi_1}\right| \hat{\alpha}^h = \frac{\beta \rho_\pi^h}{(1-(\alpha^h+\hat{\alpha}^h)\gamma^h)^2} \cdot \Gamma^h \tag{123}$$

where:

$$\begin{aligned} \Gamma^h = {} & \Lambda(1+\alpha^h+\hat{\alpha}^h)^2 \rho_\pi^h \mathbb{E}_1^h \pi_1 + (\alpha^h+\hat{\alpha}^h)(1+\Lambda\gamma^h w_1) \\ & - \Lambda(1-\rho_w^h(1+\gamma^h))(1+\alpha^h+\hat{\alpha}^h)w_1 - (\alpha^h+\hat{\alpha}^h)^2\gamma^h(1-\Lambda w_1) \end{aligned} \tag{124}$$

Using the definition of π^{h*} (equation 43), we can write:

$$\mathbb{E}_1^h \pi_1 = \mathbb{E}_1^h \pi_1 - \pi^{h*} - \frac{\gamma \rho_w^h w_1}{\rho_\pi^h} \tag{125}$$

Substituting this into equation 123 and simplifying, we obtain:

$$\frac{d\mathbb{E}_1^h \tilde{U}^h}{d\mathbb{E}_1^h \pi_1}\Big|\hat{\alpha}^h = \frac{\beta\Lambda(\rho_\pi^h(1+\alpha^h+\hat{\alpha}^h))^2}{(1-(\alpha^h+\hat{\alpha}^h)\gamma^h)^2}\cdot(\mathbb{E}_1^h\pi_1 - \pi^{h*}) + \frac{\beta\rho_\pi^h(\alpha^h+\hat{\alpha}^h-\Lambda(1-\rho_w^h)(1+\alpha^h+\hat{\alpha}^h)w_1)}{1-(\alpha^h+\hat{\alpha}^h)\gamma^h} \quad (126)$$

With $\hat{\alpha}^h$ held constant, this is a linear function $\mathbb{E}_1^h\pi_1$ with constant coefficients, and so $\mathbb{E}_1^h\tilde{U}^h|\hat{\alpha}^h$ is continuous in $\mathbb{E}_1^h\pi_1$.

In the neighbourhood of $\mathbb{E}_1^h\pi_1 = \pi^{h*}$, the first term of this is small, so the derivative has the same sign as that of the second term.

If $\hat{\alpha}^h < |\gamma^h|^{-1} - |\alpha^h|$, the denominator is positive, and so $d\mathbb{E}_1^h\tilde{U}^h/d\mathbb{E}_1^h\pi_1$ inherits the sign of:

$$\alpha^h+\hat{\alpha}^h-\Lambda(1-\rho_w^h)(1+\alpha^h+\hat{\alpha}^h)w_1 = (\alpha^h+\hat{\alpha}^h)(1-\Lambda(1-\rho_w^h)w_1)-\Lambda(1-\rho_w^h)w_1 \quad (127)$$

If restriction 120 is satisfied, we have:

$$\begin{aligned}&(\alpha^h+\hat{\alpha}^h)(1-\Lambda(1-\rho_w^h)w_1)-\Lambda(1-\rho_w^h)w_1 > 0\\ \iff &\hat{\alpha}^h > \frac{\Lambda(1-\rho_w^h)w_1}{1-\Lambda(1-\rho_w^h)w_1} - \alpha^h\end{aligned} \quad (128)$$

If condition 121 holds, this is satisfied at $\hat{\alpha}^h = \alpha^*$, so this proves the first part of statement 122.

Similarly, if restriction 120 is satisfied we also have:

$$\begin{aligned}&(\alpha^h+\hat{\alpha}^h)(1-\Lambda(1-\rho_w^h)w_1)-\Lambda(1-\rho_w^h)w_1 < 0\\ \iff &\hat{\alpha}^h < \frac{\Lambda(1-\rho_w^h)w_1}{1-\Lambda(1-\rho_w^h)w_1} - \alpha^h\end{aligned} \quad (129)$$

If condition 121 holds, this is satisfied at $\hat{\alpha}^h = -\alpha^*$, so this proves the second part of statement 122. ■

That is, expected utility is continuous and increasing in $\mathbb{E}_1^h\pi_1$ in the neighbourhood of $\mathbb{E}_1^h\pi_1 = \pi^{h*}$ when $\hat{\alpha} = \alpha^*$, and is continuous and decreasing in $\mathbb{E}_1^h\pi_1$ when $\hat{\alpha} = -\alpha^*$. From this, it follows that if the worst case belief features $\hat{\alpha} = \alpha^*$ it also features $\hat{b}^h = b^*$, as this ensures $\mathbb{E}_1^h\pi_1$ is minimised. Similarly, if $\hat{\alpha} = -\alpha^*$ then $\hat{b}^h = -b^*$.

If the perceived inflation under the central model $\mathbb{E}_1^h(\pi_1|s^h, 0)$ is sufficiently low that $\mathbb{E}_1^h(\pi_1|s^h, 0) + \tau^h b^* < \pi^{h*}$, then from Proposition 4 we have that the evil agent chooses $\hat{\alpha} = \alpha^*$ for any $\hat{b}^h \in [-b^*, b^*]$, and so the expected utility minimising model distortion is $(\hat{\alpha}^h, \hat{b}^h) = (\alpha^*, b^*)$. Similarly, if $\mathbb{E}_1^h(\pi_1|s^h, 0)$ is such that $\mathbb{E}_1^h(\pi_1|s^h, 0) - \tau^h b^* > \pi^{h*}$, the evil agent always selects $\hat{\alpha} = -\alpha^*$, and so the worst case has $(\hat{\alpha}^h, \hat{b}^h) = (-\alpha^*, -b^*)$.

Finally, the continuity of $\mathbb{E}_1^h \tilde{U}^h | \hat{\alpha}^h$ means that there exists a threshold $\pi^{hb} \in (\pi^{h*} - \tau^h b^*, \pi^{h*} + \tau^h b^*)$ such that:

$$(\hat{\alpha}^h, \hat{b}^h) = \begin{cases} (\alpha^*, b^*) & \text{if } \mathbb{E}_1^h(\pi_1 | s^h, 0) < \pi^{hb} \\ (-\alpha^*, -b^*) & \text{if } \mathbb{E}_1^h(\pi_1 | s^h, 0) \geq \pi^{hb} \end{cases} \tag{130}$$

This completes the proof of Proposition 5.

The region of α^* which satisfies constraint 121 is non-empty if:

$$\frac{1}{|\gamma^h|} - |\alpha^h| > \left| \frac{\Lambda(1-\rho_w^h)w_1}{1 - \Lambda(1-\rho_w^h)w_1} - \alpha^h \right| \geq \left| \frac{\Lambda(1-\rho_w^h)w_1}{1 - \Lambda(1-\rho_w^h)w_1} \right| + |\alpha^h| \tag{131}$$

A sufficient condition for the region of α^* which satisfies constraint 121 to be non-empty is therefore:

$$\frac{1}{|\gamma^h|} > \left| \frac{\Lambda(1-\rho_w^h)w_1}{1 - \Lambda(1-\rho_w^h)w_1} \right| + 2|\alpha^h| \tag{132}$$

This will therefore be satisfied as long as γ^h, α^h, and w_1 are all small. [To add: discussion of these. Numerically the set of permissible α^* is large.]

E Parameters for simulation-based figures in sections 4 and 5

Figures 3 and 4:

Preferences, subjective models, priors, and information costs: $\alpha^h \sim U[-1.5, 0.5], \alpha^* = 1.5, \beta = 0.9, b^* = 0.25, \chi = 1, \gamma^h \sim U[0, 0.3], \bar{\pi}^h \sim U[-0.25, 0.25], \psi = 0.001, \rho_\pi^h = \rho_w^h = 0.75, \sigma_\pi^{2h} = 0.1$

Realised aggregate variables: $w_1 = \pi_1 = 0$

Figures 6 and 7:

Preferences, subjective models, and information costs: $\alpha^h \sim U[-1.5, 0.5], \alpha^* = 1.5, \beta = 0.9, b^* = 0.1, \chi = 1, \gamma^h \sim U[0, 0.3], \psi = 0.001, \rho_\pi^h = \rho_w^h = 0.75, \sigma_\pi^{2h} = 0.1$

Realised aggregate variables: $t = 0$: $w_{1t} = 0, \pi_{1t} = 0.2$. $t > 0$: $w_{1t} = 0, \pi_{1t} = 0$.

Priors: in each simulation, I begin with $\bar{\pi}^h \sim U[-0.25, 0.25]$, then run the simulation for 100 periods with no shocks, so that the distribution of $\bar{\pi}^h$ is at its stationary distribution before the inflation shock hits. Note that this stationary distribution is different for the cases with and without the confirmation bias $b^* > 0$.

F Survey experiments in the presence of the narrative heterogeneity channel

Here I consider a simple version of an information experiment in the model of section 4, with $\alpha^* = b^* = 0$ as in section 5.1. Suppose that as in Coibion et al. (2019a), a randomly chosen subset of households are provided with a free public signal about inflation:

$$s^p = \pi_1 + \varepsilon^p, \quad \varepsilon^p \sim N(0, \sigma_p^2) \tag{133}$$

They are provided with this information after processing their own information, so the post-treatment inflation perception of a household in the treatment group is:

$$\mathbb{E}_1^h(\pi_1 | s^h, s^p) = (1 - \tilde{\tau}^h)\mathbb{E}_1^h(\pi_1 | s^h) + \tilde{\tau}^h s^p \tag{134}$$

where:

$$\tilde{\tau}^h = \frac{\sigma_\pi^{h2}(1 - \tau^h)}{\sigma_\pi^{h2}(1 - \tau^h) + \sigma_p^2} \tag{135}$$

The experimenter is interested in the causal effect of inflation perceptions on consumption. Estimating this effect by regressing c_1^h on $\mathbb{E}_1^h \pi_1$, instrumenting for perceived inflation with the treatment dummy, gives the Wald estimator:

$$\hat{\beta}_W = \frac{\hat{\beta}_c}{\hat{\beta}_\pi} \tag{136}$$

where $\hat{\beta}_c$ and $\hat{\beta}_\pi$ are the differences in average consumption and inflation perceptions between the treatment and control groups:[55]

$$\hat{\beta}_c = \int_0^1 c_1^h(\mathbb{E}_1^h(\pi | s^h, s^p))dh - \int_0^1 c_1^h(\mathbb{E}_1^h(\pi | s^h))dh \tag{137}$$

$$\hat{\beta}_\pi = \int_0^1 \mathbb{E}_1^h(\pi | s^h, s^p)dh - \int_0^1 \mathbb{E}_1^h(\pi | s^h)dh \tag{138}$$

Begin with the difference in inflation perceptions, assuming for simplicity that all households have the same prior $\bar{\pi}^h = 0$:

$$\hat{\beta}_\pi = \int_0^1 \omega^h \tilde{\tau}^h(\varepsilon^p - \tau^h \varepsilon_\pi^h + (1 - \tau^h)\pi_1)dh \tag{139}$$

[55] Note I am assuming that the populations of treated and untreated are both of measure 1 for simplicity here. It is as if each household h in in fact a pair: one who gets treated and one who does not.

The first two terms of this integral are:

$$\int_0^1 \omega^h \tilde{\tau}^h(\varepsilon^p - \tau^h \varepsilon_\pi^h)dh = \mathbb{E}_H(\tilde{\tau}^h \varepsilon^p) - \mathbb{E}_H(\tilde{\tau}^h \tau^h \varepsilon_\pi^h) = 0 \tag{140}$$

where the final equality comes from the fact that both ε^p and ε_π^h are independent of household characteristics and information choices, and both have mean zero. We therefore have:

$$\hat{\beta}_\pi = \pi_1 \int_0^1 \omega^h \tilde{\tau}^h (1 - \tau^h) dh \tag{141}$$

Turning to the consumption difference, note that the consumption function (22), substituting out for expected w_2 and π_2 using equations 25 and 26, can be written:

$$c_1^h = A^h w_1 + B^h \mathbb{E}_1 \pi_1 \tag{142}$$

where A^h and B^h are combinations of household preference and subjective model parameters. Using this, we have:

$$\begin{aligned}\hat{\beta}_c &= \int_0^1 \omega^h B^h (\mathbb{E}_1^h(\pi | s^h, s^p) - \mathbb{E}_1^h(\pi | s^h))dh \\ &= \pi_1 \int_0^1 \omega^h B^h \tilde{\tau}^h (1 - \tau^h) dh\end{aligned} \tag{143}$$

It will be useful here to define:

$$T^h \equiv \tilde{\tau}^h (1 - \tau^h) \tag{144}$$

Using this, and the expressions for $\hat{\beta}_c$ and $\hat{\beta}_\pi$ (143 and 141), the Wald estimator (136) can be written as:

$$\hat{\beta}_W = \int_0^1 \omega^h B^h \cdot \frac{T^h}{\mathbb{E}_H T^h} dh \tag{145}$$

This is therefore equal to a weighted average over the true elasticity of consumption to perceived inflation B^h, with the weights equal to the ratio of T^h to $\mathbb{E}_H T^h$. Since T^h is monotonically decreasing in τ^h, this over-weights inattentive households, and underweights those with high τ^h, relative to the arithmetic mean of B^h. This therefore formally shows the point made in section 5.1, that inattentive households drive the estimated effect of expectations on actions in information-provision experiments. Note that this weighting is in the opposite direction to that used to find the response of aggregate consumption to a shock, as that over-weights the most attentive households relative to the arithmetic mean of B^h.

The effect of this over-weighting of inattentive households could be either positive or negative, depending on the distribution of subjective models. Another way to express the Wald estimator in this simple example is:

$$\hat{\beta}_W = \frac{\mathbb{E}_H(B^h T^h)}{\mathbb{E}_H(T^h)} = \mathbb{E}_H(B^h) + \frac{Cov(B^h, T^h)}{\mathbb{E}_H(T^h)} \tag{146}$$

That is, the Wald estimator gives the arithmetic mean of B^h plus a bias term. If the majority of households have $B^h > 0$, then higher $Cov(B^h, \tau^h) > 0$, so $Cov(B^h, T^h) < 0$ and the Wald estimator over-estimates $\mathbb{E}_H(B^h)$. If, however, the majority of households have $B^h < 0$, this bias term could be positive.