

# Household inflation expectations in the UK: exploiting the cross-sectional dimension\*

Alina Barnett<sup>†</sup>      Haroon Mumtaz<sup>‡</sup>      Matthias Paustian<sup>§</sup>      Silvia Pezzini<sup>¶</sup>  
*Bank of England*      *Bank of England*      *Bank of England*      *Bank of England*

First draft: November 2010

Preliminary and incomplete. Please do not quote without permission!

## Abstract

We test whether survey data of household inflation expectations in the U.K. are consistent with the sticky information framework of Mankiw and Reis (2002). Consistent with the theory, the dispersion of inflation expectations increases around turning points in the macro data such as recession and disinflations. When we fit the empirical mean to match the mean of the surveys we estimate that U.K. households update their information sets about once a year. When fitting the full distribution, we use the Kullback-Leibler distance measure to compare a full information model with variants of the sticky information model and find that the sticky information model outperforms the full information model over the entire sample. Within the class of sticky information models, the model with time varying weights fits best the dynamics of the Barclays Basix survey. Finally, the paper combines the macro approach with a microeconomic view of the sticky information hypothesis. We link inflation forecasts based on different vintages of information sets to individual specific characteristics of respondents in the survey. Our findings show that individuals with a higher level of education are more likely to report inflation expectations that are consistent with frequent updating of their information sets.

Key Words: Inflation expectations, structural VAR.

JEL codes: E5;C1

---

\*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Bank of England or the members of the Monetary Policy Committee. This paper has benefited from helpful comments by George Kapetanios, Tim Taylor and Garry Young.

<sup>†</sup>Monetary Assessment and Strategy Division, Bank of England. Email: [alina.barnett@bankofengland.co.uk](mailto:alina.barnett@bankofengland.co.uk)

<sup>‡</sup>Centre for Central Bank Studies, Bank of England. Email: [haroon.mumtaz@bankofengland.co.uk](mailto:haroon.mumtaz@bankofengland.co.uk)

<sup>§</sup>Monetary Assessment and Strategy Division, Bank of England. Email: [matthias.paustian@bankofengland.co.uk](mailto:matthias.paustian@bankofengland.co.uk)

<sup>¶</sup>Monetary Assessment and Strategy Division, Bank of England. Email: [silvia.pezzini@bankofengland.co.uk](mailto:silvia.pezzini@bankofengland.co.uk)

Information in the world exists “solely as the dispersed bits of incomplete and frequently contradictory knowledge, which all the separate individuals possess”, Hayek(1945, p519)

## 1 Introduction

During the recent recession, we have seen a strong increase in cross-sectional dispersion of household inflation expectations in the U.K. This is at odds with standard rational expectations models that are at the core of most central banks’ forecasting models. These models leave no room for heterogeneity or disagreement. In recent years research has moved away from full information rational expectations models towards a framework in which some form of limited information or bounded rationality is assumed. This paper assesses whether survey measures of U.K. households’ inflation expectations are consistent with one such model, namely the sticky information framework of Mankiw and Reis (2002).

We test alternative theories of household inflation expectations formation where the heterogeneity comes from agents using different information vintages and from using different prediction models. At the centre of the test is the framework by Mankiw and Reis (2002). These authors argue that information acquisition is costly for households and that they will therefore update their expectations only infrequently. In between updates of information sets, households form inflation expectations based on optimal forecasts given their outdated information. When the timing of information updates is not synchronized across households, macroeconomic events pass through to inflation expectations only slowly over time. Such a process of information diffusion has rich implications for the cross sectional distribution of household inflation expectations that we wish to exploit.

In a first step, we follow the methodology of Mankiw et al. (2003). We use vector-autoregressions (VARs) run on small set of macro data as a proxy for the process with which household form expectations. Households differ in their forecasts because they have access to different vintages of information on the macrodata. Thus, for a given distribution of the vintage of information across agents, the model will predict a distribution of household inflation forecasts at any point in time. We calibrate the average age of the information set by matching the model’s prediction to the mean of the survey and then use the amount of dispersion this generates in the cross section as a metric for model evaluation.

We find moderate support for the theory in the Basix survey covering the period 1987-2010. Consistent with the model, cross sectional dispersion in the survey increases around turning points in the macrodata such as during recessions and disinflation periods. But the average amount of dispersion generated by the model is too small relative to the data. Furthermore, the dispersion generated by the model correlates only weakly with the dispersion in the survey with a correlation coefficient of about 30%, much smaller than what Mankiw et al. (2003) find for the US. We estimate that UK households update their information sets about once a year. This is in line with the monthly sticky information model of Mankiw et al. (2003) for inflation expectations in the Michigan survey.

In a second step we build on this analysis and allow for parameter uncertainty. We assume that households

form forecasts 'with a trembling hand'. They draw their forecasts randomly from the posterior distribution as opposed to reporting forecasts based on the posterior mean or mode. This proxies for imperfections on the part of respondents in accurately performing the computations required for the VAR forecasts. This can explain the average amount of dispersion seen in the survey over our sample. At the same time, it gives us a natural and non-trivial competition model for the sticky information hypothesis. Under this trembling hand assumption, although small, there is dispersion of forecasts even in a perfect information world.

We then compute non-parametric measures of model fit for the models discussed above but we split them into two alternative classes of models, one of full information and one of sticky information. The first assumes that everybody updates their information sets regularly and thus make forecasts incorporating the latest information while the second assumes that there are households which do not always update regularly with the latest information. Within the second class we allow the share with which the forecasts produced by different information vintages are weighted to vary. We obtain candidate share of forecasts made with different information vintages by using a beta distribution and choose the shares that minimise a measure of density closeness, namely the Kullback- Leibler (Klic) criterion during the entire sample (static sticky information model as in Mankiw and Reis (2002)) and, then to be more general, we choose those that minimise the Klic at each point in time (time varying sticky information model). For comparison we also present results obtained by assuming that all information vintages lead to forecasts that are uniformly weighted. We find that for the Basix survey of inflation expectations the full information model fits less well than the model of sticky information in which we assume some structure for the information vintages. We also find that for the Great Stability period the full information and the sticky information model are closest while during the recent financial crises, the time varying sticky information model fits the survey distribution best. Although it is likely that the time varying sticky information model is overfitted, this model allows us to track how the share of frequent and rare updaters evolves over time. We find that the share of those who updated in the last quarter has significantly decreased during the Great Stability from about 60% to under 10% while the share of those who update less frequently increased. In other words if before 1995 households were updating their information sets once a year, during the Great Stability these updated their information sets every other year. The main implication of this finding is that the distribution of agents is not always geometric and thus the highest proportion of households update, on average, only infrequently. This is in contrast with Mankiw and Reis (2002) model that predicts that the highest share of agents update each period.

In the final section of the paper, we link the macro approach to evaluating the sticky information framework with a microeconomic angle in a novel way. The aim here is to test whether households with certain micro characteristics are more likely to report inflation expectations that are consistent with macro VAR forecasts based on more recent information sets. If those households were informally judged to have smaller costs of updating their information sets because of their personal characteristics, this would lend support to the sticky information framework. We find some modest evidence for the sticky information framework based on this micro approach. Frequent updaters are more likely to be individuals with higher levels of education. One

would expect that individuals with higher education have smaller costs of updating their information sets and would therefore do so more frequently. It appears reasonable that individuals with a degree have easier and more regular access to newspapers and other media that contain information about current inflation. Current inflation in turn is often used as a naive predictor for future short term inflation. The probit regression result that associates people who hold a degree with frequent updaters thus provides some support for the theory of sticky information.

Our results are specific to the classes of expectation formation models used in this paper and therefore we are unable to make broader claims about how well other classes of models, be it those of inattentiveness or not, fit the UK inflation expectations dynamics. However our results suggest that models of information stickiness are a valuable initial step towards understanding the rich dynamics of inflation expectations but future research that will take into account model uncertainty as well as different assumptions about the basket of goods households could be considering when making expectations will further enhance our understanding of the process behind inflation expectation formation.

The paper is organised as follows. Section 2 discusses some of the most relevant empirical literature that models heterogeneous inflation expectations. Section 3 describes the data. Section 4 provides the main analysis and Section 5 concludes.

## 2 Relation to the literature

The empirical literature on inflation expectations is large. Most of it has focused on the mean or median in the cross section, while the literature on cross sectional dispersion is much more recent. Heterogeneity of inflation expectations can be traced back to at least two different origins. Agents may have access to different information sets or they may use different forecasting models.

The first avenue is pioneered by Carroll (2003) and Mankiw and Reis (2002). Carroll (2003b) provides an epidemiological model of inflation expectations in which expert opinion slowly spreads to households, similar to the spread of disease across the population. He analyzes the evolution of the mean and the standard deviation of inflation expectations in the Michigan Survey. Carroll finds that his model tracks the time series of standard deviation well, but generates far too little dispersion on average.

The paper closest to ours is Mankiw et al. (2003). These authors test whether infrequent updating of information sets can explain the extent of disagreement about inflation expectations in the Livingston Survey and in the Michigan Survey. They construct a monthly VAR on annual CPI inflation, short term interest rate, and the output gap from 1947 to 2001. The full cross sectional distribution of inflation expectations at any point in time is constructed by assuming that agents use this VAR when forming their expectations, but update their information set only infrequently. As in Mankiw and Reis (2002), information updating occurs with exogenous probability such that the population shares of agents are geometrically declining in the time elapsed since the last update. This rate of information updating is then chosen to maximize the correlation

between the interquartile range of inflation expectations in the survey data and the interquartile range predicted by the model.

In Mankiw et al. (2003), the overall assessment of the sticky information model is positive: the model predicted disagreement tracks the disagreement among professionals in the Livingston survey well over time. The level of disagreement among the general public is significantly higher on average than that predicted by the model, but the two are highly correlated. Furthermore, the model predicted median expected inflation tracks the median in the two surveys reasonably well over time. Mankiw et al. (2003) find that the professional economists in the Livingston Survey update their expectations on average about every 10 months, while the general public sampled in the Michigan survey updates their expectations on average every 12.5 months.

Our paper differs from Mankiw et. al. (2003) in several respects. First, we do not construct sticky information forecasts based on VAR parameter estimates over the entire sample. Instead, we use real time data and estimate parameters only on data up to the forecast origin of each information cohort, mimicking closely the process that agents reasonably may be expected to have used in practice <sup>1</sup> Second, we allow agents to use different forecasting algorithms other than ordinary least-squares.

Dovern et al. (2010) provide a comprehensive study of the evolution of disagreement among individual expert forecasts in surveys covering the G-7 countries. They report that disagreement about inflation expectations tends to be large and more sensitive to macroeconomic news in countries whose central banks became independent relatively late in the 1990s, of which the U.K. is an example.

Our paper is also related to the literature examining which simple forecasting algorithms might best describe the expectations formation process of boundedly rational agents. Branch and Evans (2006) compare the forecasting performance of recursive least-squares, constant gain least-squares and time-varying parameter algorithms in bivariate VAR models of inflation and real GDP. They show that the constant gain algorithm forecasts well out of sample and provides the best fit to the median response in the Survey of Professional Forecasters. The paper does not address what accounts for the dispersion in forecasts among professionals.

Branch (2004) uses monthly U.S. data from the Survey of Consumer Attitudes and Behavior to build a dynamic model of predictor selection. Agents choose between VAR, adaptive and naive expectations when forming their inflation expectations. The choice of a particular predictor is allowed to vary over time and depends on its relative mean-squared error. He finds that the VAR predictor accounts for about 48 % of the sample, adaptive expectations account for about 44% and the naive predictor for about 7%. It is not clear how well this set of predictors actually matches the data, because the paper does not provide a measure of model fit such as a plot of predicted and actual distribution of expectations in the cross section. This framework is extended in Branch (2007) to allow for a sticky information model with a time varying distribution structure. Branch finds that such time varying weights are consistent with the Michigan survey of inflation expectations.

A paper similar in spirit is by Lanne et al. (2009) who extend the basic sticky information framework to a richer set of predictors. Agents base their inflation expectations either naively on latest inflation or on

---

<sup>1</sup>This was originally suggested in Williams (2003).

professional forecasts of inflation. When modeling the mean of household inflation expectations in the Michigan survey over the period 1981-2001, the authors find support for this hybrid model. They estimate that about 65% of people use naive forecasts, but this parameter is estimated with a very large posterior standard deviation. When fitting the whole survey distribution, there is strong evidence for this hybrid model. The fraction of people using the naive predictor is estimated much more tightly at about 43%, while about 6% of people are found to draw their expectations randomly. These authors suggest that they are able to match the high degree of cross-sectional dispersion of inflation expectations present in the survey by allowing the speed of updating to be dispersed across agents.

Pfajfar and Santoro (2010) analyze the percentiles of the distribution for inflation expectations in the Michigan Survey. They find that a nearly rational region around the median, a static or highly autoregressive region on the left hand side and a fraction of forecasts on the right hand side accord with adaptive learning and sticky information.

### 3 Data

There are three main surveys of consumer inflation expectations conducted in the UK. Barclays Basix and Bank NOP are conducted quarterly while the CitiGroup/YouGov is conducted monthly. This paper focuses only on the quarterly surveys, because the monthly survey by CitiGroup/YouGov is available only since 2005. Table 1 summarizes some key characteristics of the quarterly surveys.

	Barclays Basix	Bank NOP
Survey population	Cross-section of the general public	Cross-section of the general public
Survey organisation	GfK/ NOP (commissioned by Barclays Basix)	GfK/NOP (commissioned by Bank of England)
No of respondents	Roughly 2000 each quarter	Roughly 2000 three times per year
Starting date	1987	1999
Frequency	Quarterly	Quarterly
Questions asked	Rate of inflation over next 12 and 24 months	Prices in general over next 12 and 24 months

Table 1: Description of the data

Figure 1 shows the evolution of the mean and selected percentiles of the distribution of household’s inflation expectations in the Basix and NOP survey. One apparent difference between the two is that overall the mean in Basix one year ahead expectations is higher than the mean in the NOP one year ahead expectations. This may be driven by the choice of answers presented to respondents. The person answering the Basix survey would be given a choice about the level of inflation over the next 12 months between ‘less than 0%’ and ‘larger than 10%’ in increments of 1%. The person answering the NOP survey, until 2008, was given a choice between ‘less than 0%’ and ‘larger than 5%’ in increments of 1%. Since 2008, if a person happens to say that he or she thinks that prices will increase by more than 5% than a new choice is being offered ranging between ‘5-6%’ to ‘larger than 10%’ in increments of 1%. A similar follow up question is asked for respondents who expect prices to fall.

Thus, the range of answers is much narrower in the NOP than in the Basix survey. Given that the two surveys are conducted by the same institution and have a similar survey population, the question asked and the choice of answers given to the households must explain the level difference between the two survey means.

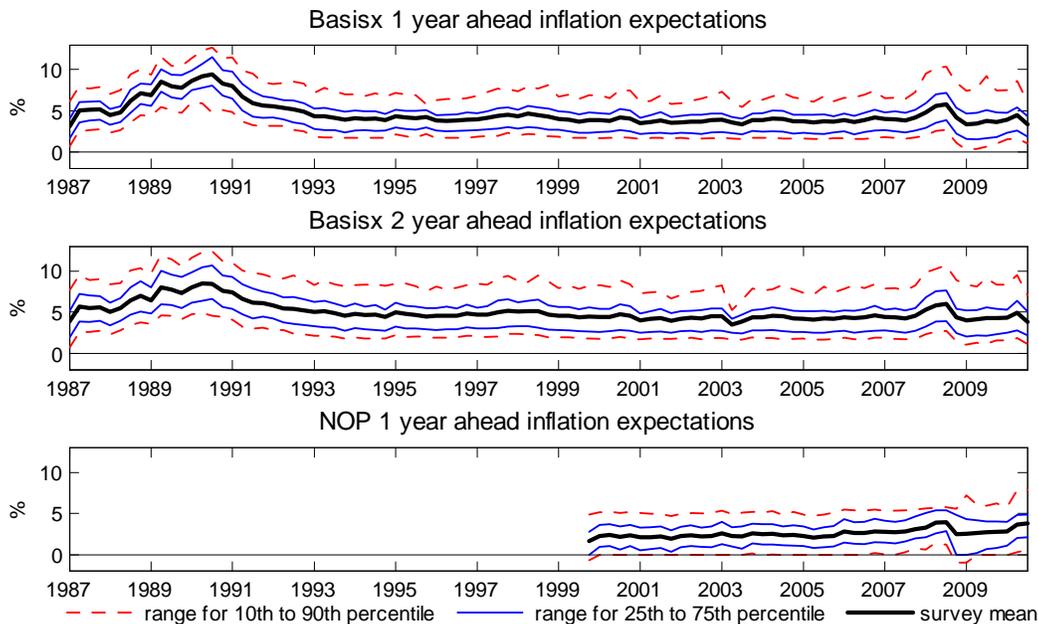


Figure 1: Mean and selected percentiles of inflation expectations in the surveys.

The Basix survey spans a period of time during which two recessions and one major disinflation occurred in the U.K, while the NOP was only started in 1999. Since the sticky information framework makes sharp predictions about the co-movement of dispersion in inflation expectations and time variation in the macroeconomic environment, the Basix survey will be our key survey. However, the NOP survey also asks households about their perceptions of current inflation. The sticky information framework implies that disagreement about perceptions of current inflation and expectations about inflation in the near term should co-vary in times of macroeconomic volatility. We plan to examine this co-movement in future work. Furthermore, the NOP survey contains rich information on individual specific information such as age, education etc. that we exploit in a novel way to provide an additional test of the sticky information framework.

## 4 Model evaluation

Our framework for explaining cross sectional dispersion has two dimensions. In the first dimension the only source of dispersion comes from infrequent updating of information. In the second dimension the dispersion comes from people using different models of expectation formation, allowing for other expectations formation processes, such as a time varying VAR with stochastic volatility and a constant gain updating VAR in the spirit of Branch and Evans (2006). All these VAR models will make different assumptions about how new information

is incorporated in the expectations formation process. This second dimension is not yet implemented in the current draft of the paper due to time constraints, so this version of the paper presents results for the first process only. The richness of our exercise will come from bringing these two dimensions together and using these to best fit the survey of inflation expectations.

## 4.1 Baseline analysis

We take the sticky information model to the data by following closely the approach in Mankiw et al. (2003). Households' inflation expectations are proxied by forecasts based on small scale VARs comprised of real time GDP growth, inflation and Bank rate at a quarterly frequency. We use the Basix survey responses and our measure of inflation is the retail price index excluding mortgage interest payments (RPIX).

We build a time series of the distribution of inflation forecasts predicted by the sticky information model based on rolling window regression of 20 years' length. The forecasts are computed at the posterior mean of the parameters. At any time  $t$ , households that last updated their information set  $j$  periods ago use data covering the period  $t - j - 80$  quarters to  $t - j$  quarters when estimating the VAR. This rolling window crudely captures the idea that households are aware of possible time variation in the macroeconomic interdependencies and generally place little weight on data in the very distant past. Contrary to Mankiw et al. (2003) we do not estimate the VAR over the whole sample, because data about the future was simply not available to households in previous periods.

Conditional on an age distribution of information sets in the cross section, we can use these VAR forecasts to build an artificial distribution of survey responses of one year ahead inflation expectations. This age distribution is constructed based on geometric decay as in Mankiw and Reis (2002). The share of agents using information outdated by  $j$  quarters is given by  $\theta(1 - \theta)^j$ . The parameter  $\theta$  is chosen to minimize the distance between the survey mean of inflation expectations and that of the model over the quarterly observations in the survey. The amount of dispersion the model generates in the cross section is then used as a metric for model evaluation.

### 4.1.1 One year ahead inflation expectations

Figure 2 shows the evolution of the mean and the degree of disagreement in households' inflation expectations from the model and how these compare with the survey. Under the assumption of geometric decay, the best match with the mean of the survey is achieved for  $\theta = 0.28$ . This implies that households update their information sets slightly more often than once a year. This is in line with the monthly sticky information model of Mankiw et al. for inflation expectations in the Michigan survey. The resulting age distribution, plotted in the third panel of the Figure 2, shows that few households use information outdated by much more than three years. The top panel shows that the mean of the model does a decent job of tracking the survey mean, but underpredicts during all of the post 1997 period. However, the survey mean itself is above actual inflation for most of that period, which partly explains why the VAR has a hard time matching the survey.

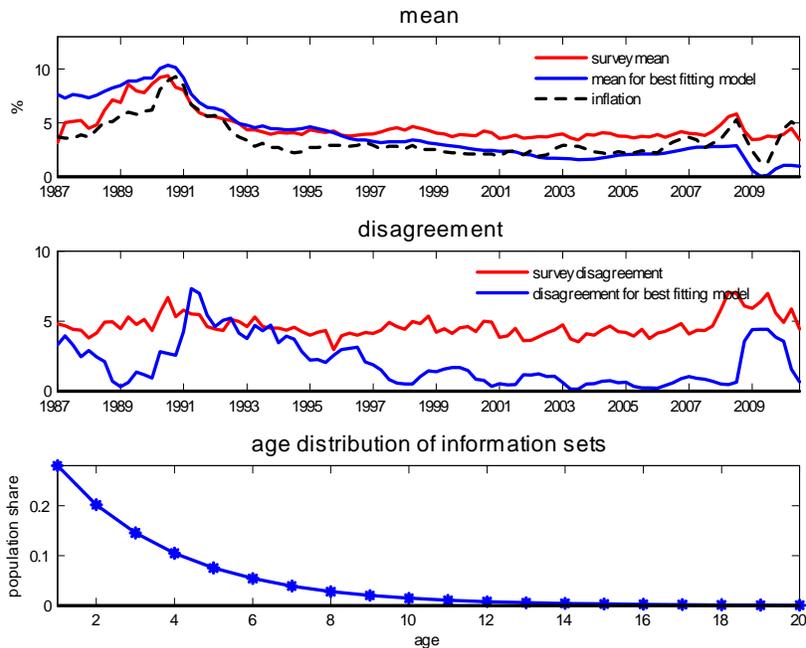


Figure 2: Mean, disagreement, age distribution for one year ahead expectations

The second panel plots dispersion, often referred to as disagreement. The typical measure in the sticky information literature is the interquartile range. But in the survey there is not enough variation over time in this measure, so we choose the range spanned by the middle 75 percent of the distribution as our measure of disagreement. Thus, disagreement is defined as the distance between the forecast at the 87.5th percentile and the forecasts at the 12.5th percentile. The success of the model to match the evolution of disagreement is modest, at best. The second panel in Figure 2 shows that dispersion in the model increases around some key periods such as the early 1990's and over the period 2009 - 2010. The same is true in the survey, although survey disagreement rose a bit earlier in each of these two episodes. On the other hand, the model predicts very little disagreement over the great stability period during early 2000, while disagreement in the survey is still very substantial. Overall, the correlation of disagreement in the model and in the survey is only 30 percent, much smaller than what was found in Mankiw et al. (2003) for the U.S.

Figure 3 illustrates the extent to which the sticky information model can explain the amount of disagreement in the data based on the cross-sectional distribution at three illustrative points in time. When the key macro variables that drive VAR forecasts have undergone turning points such as during times of disinflation or recessions in 1992 and 2010, the model is able to generate a decent amount of cross sectional heterogeneity in short term inflation expectations. In those times, information sets that are outdated by just a few quarters give rise to sufficiently different forecasts because the macro data at the forecast origin is very different. At the same time, in great moderation type of periods such as in 2004, somewhat outdated information sets do not generate much difference in forecast because the macro-economy was very similar in any of the previous quarters. What

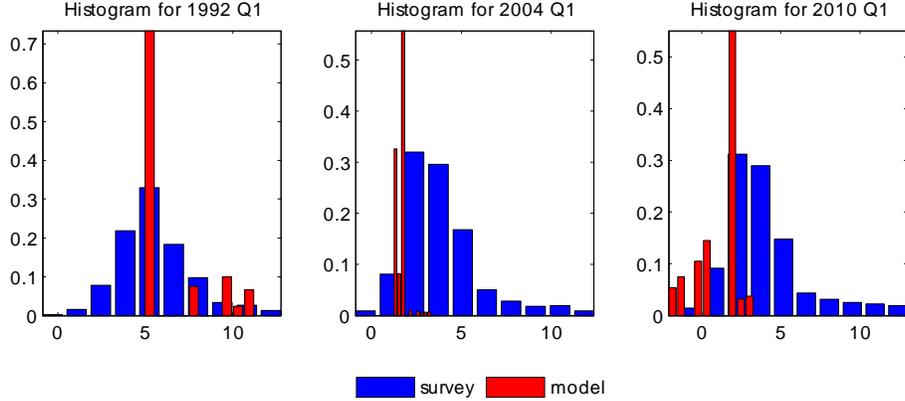


Figure 3: Cross sectional distribution: model and survey.

is striking about the survey data is that the distribution of responses in 2004 becomes less dispersed than in 1992, but not by nearly as much as one would expect given the reduction in volatility of the recent data. This is clearly at odds with the sticky information model.

Could this failure in generating enough dispersion in tranquil times be due to our assumption of geometric decay of the age distribution? To test for this, we allow for a more general age distribution of information sets based on the beta distribution. We vary the two shape parameters of the beta distribution on a fine grid and again choose those values that best match the mean of inflation expectations according to the distance criterion mentioned before. We implement this in the following way. We assume that no information set is outdated by more than 20 quarters. We then pick 20 evenly spaced points in the support of the beta distribution, evaluate the pdf at these points and impose the normalization that these values add up to one. This gives us candidate population shares for the information sets. Values for the two shape parameters on the grid of 0 to 5 allow for quite general distributions, such as uniform, hump shaped, inverted bell shapes and others. We find that the distribution which best fits the mean of the series over time looks very similar to geometric decay. Therefore, if one insists on matching the survey mean over time, there seems little to object against geometric decay.

#### 4.1.2 Two year ahead inflation expectations

Results for the two-year ahead inflation expectations are similar to the previous ones. The speed of information updating is calibrated at 0.23 and thus broadly similar to the value 0.28 obtained previously for the one year ahead expectations model. The fit of the model relative to the mean of the survey is again decent, but the correlation with the disagreement in the survey is only about 25 percent.

We summarize our results for this baseline assessment of the sticky information model as follows. The sticky information model can more than explain the rise in the dispersion of inflation expectations seen during the recent recession and in the early 1990's. But the model fails to generate a large average amount of dispersion during tranquil times.

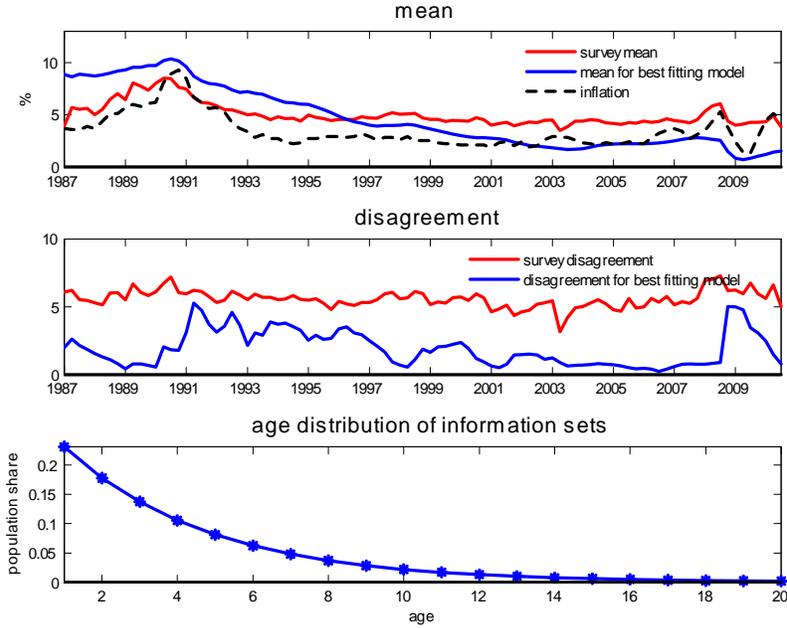


Figure 4: Mean, disagreement, age distribution for two year ahead expectations

One way to explain this large average amount of dispersion is to assume that households may have limited forecasting abilities. Forecasts may be much cruder and simpler functions of relationships observed in the past than the VAR implies. We follow this line of argument in the next subsection, but we need to put bounds on how crude forecasts can be. To this end, we assume that households randomly draw a number from the posterior distribution for their forecasts, rather than report them at the posterior mode as any good econometrician would do. We justify using the full posterior distribution by a 'trembling hand' analogy: households may make idiosyncratic errors when computing the OLS estimator. Our assumption implies that "errors" get smaller when the relationships between the variables in the data become clearer, which is certainly desirable. Another advantage of this assumption is that we can now test the sticky information model against a non trivial alternative - the full information model that generates dispersion which varies endogenously over time due to our trembling hand assumption.

#### 4.1.3 Perceptions of inflation and the sticky information model

Households perceptions of current inflation may serve as a useful cross check of the sticky information model. If information about the macroeconomy passes through the population only slowly, then survey measures of inflation perceptions should be dispersed in the cross section just like inflation expectation should be dispersed. Furthermore, the mean of inflation perceptions should be smooth over time relative to actual inflation as the distributed lag of information sets smoothes the mean. Fortunately, the Bank/NOP survey asks respondents about inflation perceptions. The full distribution of responses is available to us only since 2005. Such a short sample comes with many caveats. Nevertheless, this short period of time contains relatively large swings in

actual inflation and could be quite informative about the theory.

For perceptions, we follow the same approach as before and calibrate the frequency of updating information sets to fit the mean of perceptions. In this subsection, we allow for parameter uncertainty as well as sticky information. To do this, we run a Bayesian VAR on the same rolling window as before. Agents who have current information know actual inflation for sure and there is no dispersion within this cohort. Agents who have information outdated by  $j > 1$  periods scroll the VAR model forward  $j$  times to form an estimate of current inflation. We incorporate parameter uncertainty in the following way. We build an artificial sample survey of responses at any point in time, by drawing  $\theta(1 - \theta)^j 1000$  times from the posterior distribution for the BVAR perceptions of the relevant vintages. The frequency of updating is chosen to minimize the distance between the survey mean for perceptions and the mean predicted by the BVAR model.

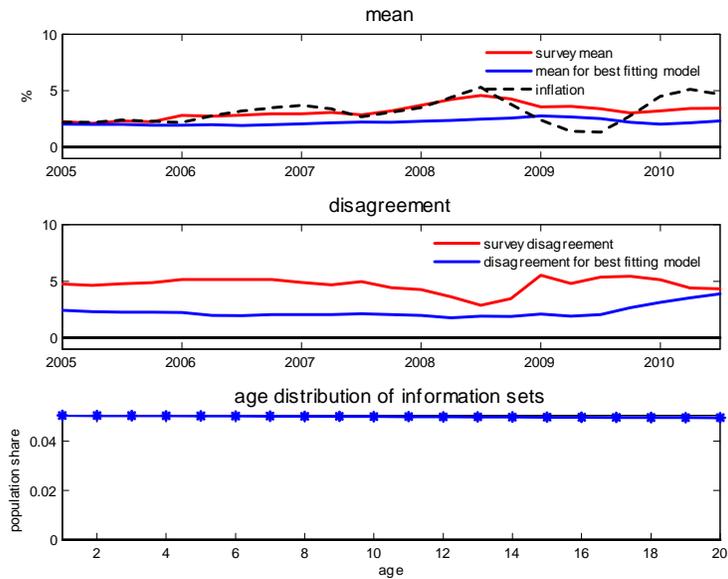


Figure 5: Inflation perceptions: mean, disagreement, age distribution.

Figure 5 summarizes the result from this exercise. The first panel in the figure shows that the mean of inflation perceptions is indeed smooth over time relative to actual inflation as the theory predicts. Turning to the last panel, the calibration procedure prefers a corner solution for the frequency of updating, where the age distribution decays so slowly that it is practically a uniform distribution of vintages. But as before, the amount of dispersion generated by the model is too low relative to the data. Overall, inflation perceptions seem to support the sticky information framework qualitatively, because the mean of inflation perceptions is smooth relative to actual inflation as predicted by the theory.

## 4.2 Comparing sticky and full information models

We now turn a comparison of the sticky information and the full information models. We use a diffuse prior in the Bayesian VAR and estimate the model in the same fashion as before, rolling window regressions based on 20 years of data. As before, we assume geometric decay and choose the rate of information updating to match the mean of inflation expectations in the survey. This is implemented in the following way. We build an artificial sample survey of responses at any point in time  $t$ , by drawing  $\theta(1-\theta)^j$  1000 times from the posterior distribution for the BVAR forecasts based on information outdated by  $j$  periods.<sup>2</sup> The value  $\theta$  is estimated at 0.26, very similar to the results before when only the point forecast was used for each vintage. Figure 6 shows that the main difference with the previous results is that this model generates much more disagreement on average. The average amount of disagreement in the model and the survey are roughly the same, but the model overpredicts disagreement in the early 1990s and underpredicts in the new millennium. The mean is matched in a similar way as before.

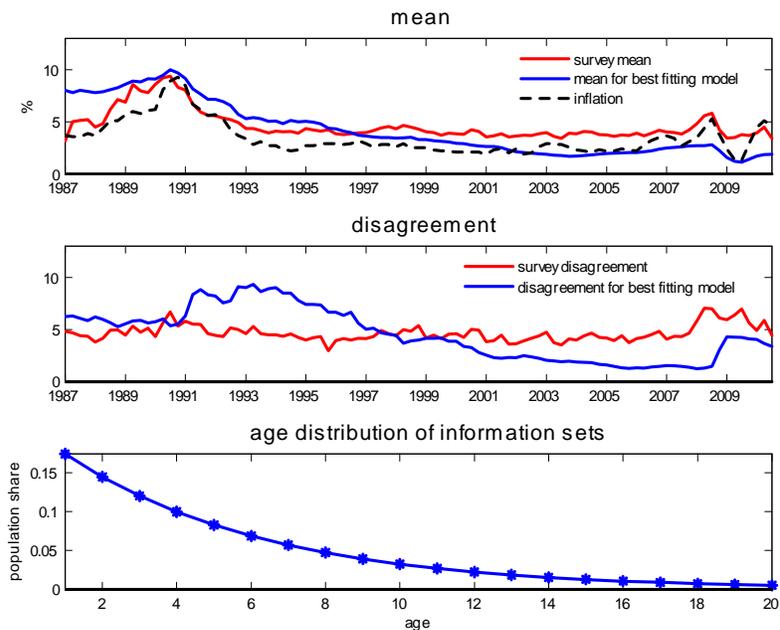


Figure 6: Mean, disagreement, age distribution for BVAR model

Turning to the cross sectional distribution at selected points in time, Figure 7 shows that the model is successful at predicting the full range of responses in the survey in 1992 Q1. In fact, the slow diffusion of information through the population implies that it predicts too many responses in the bucket of inflation expectations larger than 10%. This arises because inflation in the years prior to 1992 was relatively high and older vintage extrapolate this high level into the future. In the great moderation period (for which we choose

<sup>2</sup>In practice, we need to truncate and let  $j$  run from 0 to 20. In line with the survey, we transform each BVAR forecasts to the midpoint of the corresponding bucket of survey responses.

2004 Q1 as an example), the information lag structure is not enough to generate the dispersion seen in the survey, although the model can at least fill any bucket from 0-5%.

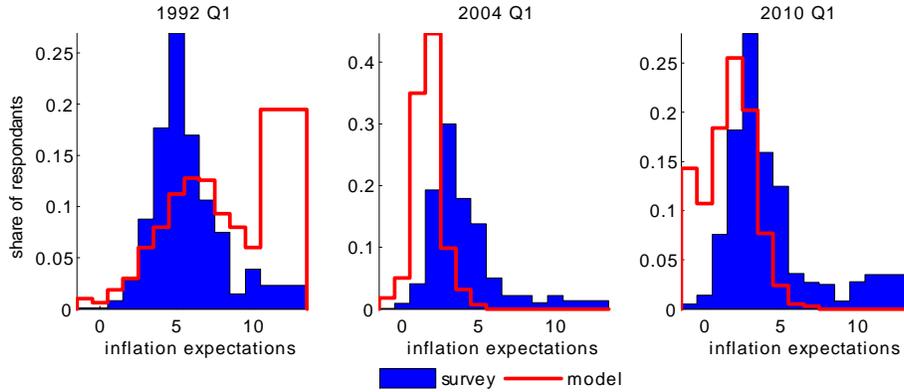


Figure 7: Survey and model distribution for sticky information model

What does the full information model look like? Figure 8 shows that this model tracks the mean of inflation expectations about right and generates significant dispersion. On the other hand, the mean of inflation is much more erratic than the data in the late 1980s and the early 1990s.

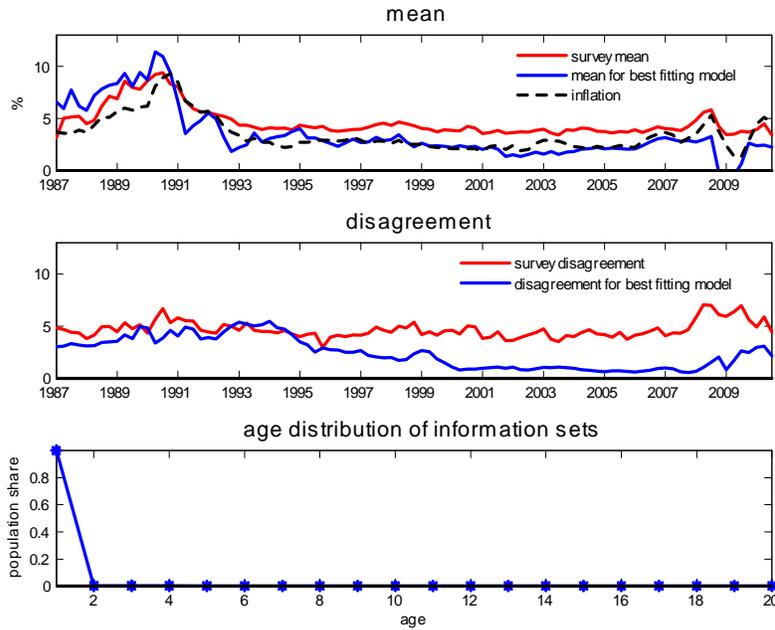


Figure 8: Mean, disagreement, age distribution for full information BVAR model

Turning a closer look at the cross sectional dispersion for this model, Figure 9 shows that parameter uncertainty alone can account for almost all of the dispersion in survey inflation expectations in 1992 Q1, and in 2010

Q1. Furthermore, the full information model fits the shape of the survey distribution surprisingly well at these two points in time. Comparing Figures 7 and 9 visually suggests that it might be hard to distinguish between the two models based on fitting the entire survey distribution at these three selected points in time. Of course, these are just selected dates and a complete model comparison would compute distance measures between the survey distribution and the distribution predicted by the models at each point in time. We turn to this exercise in the next subsection.

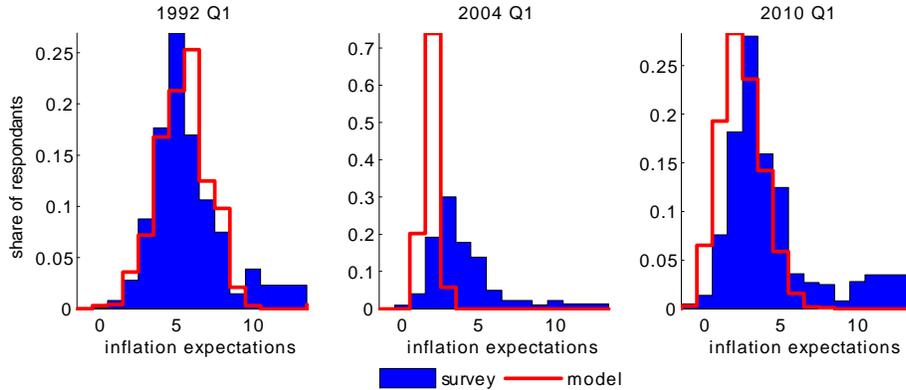


Figure 9: Survey and model distribution for full information model

### 4.3 Assessing model fit

The previous subsection has shown the need for a proper measure of model fit, to discriminate between different theories of dispersion in inflation expectations. Robertson et al. (2005) and Cogley et al (2005) use the Kullback- Leibler information criterion (KLIC<sup>3</sup>) to measure how close the two densities are. This criterion is a non-symmetric measure of the difference between the two probability distributions and is defined as:

$$Klic(p, p^*) = \int \log \left[ \frac{p^*(Y_{j,t...t+4}|Y_{t-1})}{p(Y_{j,t...t+4}|Y_{t-1})} \right] p^*(Y_{j,t...t+4}|Y_{t-1}) dY_{j,t...t+4} \quad (1)$$

where  $p$  is the estimated empirical density function and  $p^*$  is the estimated Barclay Basix's density function. We discuss the details of the non-parametric density estimation in the Appendix.

As mentioned at the beginning of the section, we perform two experiments using the Barclay Basix survey. The first one, in which we assume that everybody is fully informed, is the one that we would like to compare the others against. We call this the full information exercise. The second one assumes that not everybody is fully informed and as a result the mean of the estimated inflation rate reflects forecast made with outdated information. We call this the sticky information exercise. Within this we allow the share with which the forecasts produced by different information vintages are weighted to vary. We obtain candidate share of forecasts made with different information vintages by using a beta distribution and chose the shares that minimise the Klic

<sup>3</sup>In probability theory and information theory this is also called the information divergence, information gain or relative entropy.

distance measure first during the entire sample (static sticky information model) and then to be more general chose those that minimise Klic at each point in time (time varying sticky information model). For comparison we also present results obtained by assuming that all information vintages lead to forecasts that are uniformly weighted. Although this is a case nested within our second exercise it is interesting to see how this fairs.

For the second exercise we use a beta distribution that is governed by two shape parameters that we vary a fine grid which allows for quite general distributions, such as uniform, hump shaped, inverted bell shapes and others . As done until now, we assume that no information set is outdated by more than 20 quarters. Therefore we then pick 20 evenly spaced points in the support of the beta distribution, evaluate the pdf at these points and impose the normalization that these values add up to one. This gives us candidate population shares for the information sets that we use to evaluate the Klic and chose those shares that minimise the Klic criterion i) over the entire sample and ii) at each point in time. Following Branch (2007), we conclude that the model with the smallest Klic is the model most consistent with the data.

In contrast to our findings in the previous section, the full information model performs much worse than any of the thee versions of sticky information models. The full information model was closest to the other three between 1993 and 1999, started to perform worse during the great stability and was the furthest away from the survey data in 2009. In 11 we plot the estimated empirical densities against the estimated survey's density as an example of how different or alike these are. We chose 4 periods: the first, 1990 Q3, coincides with the highest inflation since 1990, the second, 2000 Q1, is a sample of the great stability, with the third and fourth two dates during the financial crises. Realised inflation increased significantly in 2008 before falling back sharply in 2009 Q1 and the full information model predicts that inflation will continue to fall in Q3 2009 . However expectations did not suffer from a sudden and significant fall, as the actual inflation did and so the full information model has a hard time explaining the survey(see bottom right hand corner of figure 11). All models have a hard time in matching the survey during the second part of 2009 and beginning of 2010 although, perhaps not surprising the time varying weights model performs best. We do not want to overplay the importance of this result as this model is likely to be overparametrised, but for what it is worth a model in which the share of people with different information vintages differs from year to year fits the Basix survey of inflation expectations best.

If one believes in the sticky information models then one would expect that, during periods of economic tranquility, the share of those who updated in the latest quarter would decline as this would represent a cost set against a very small benefit given the stability of the economy. This is exactly what we find when we estimate the time varying weights. In the top pannel of 12 one can see that the share of those who updated in the last quarter has significantly decreased during the great stability from about 60% to under 10% while the share of those who update less frequently increased. Put differently, before 1995 and during the higher and more volatile inflation rates in the 1990s, households were updating their information sets just under once a year while during the great stability these updated their information sets every other year. The main implication of this finding is that the distribution of agents is not always geometric and thus sometimes the highest proportion of households update, on average, only infrequently. This is in contrast with Mankiw and Reis (2002) model that predicts

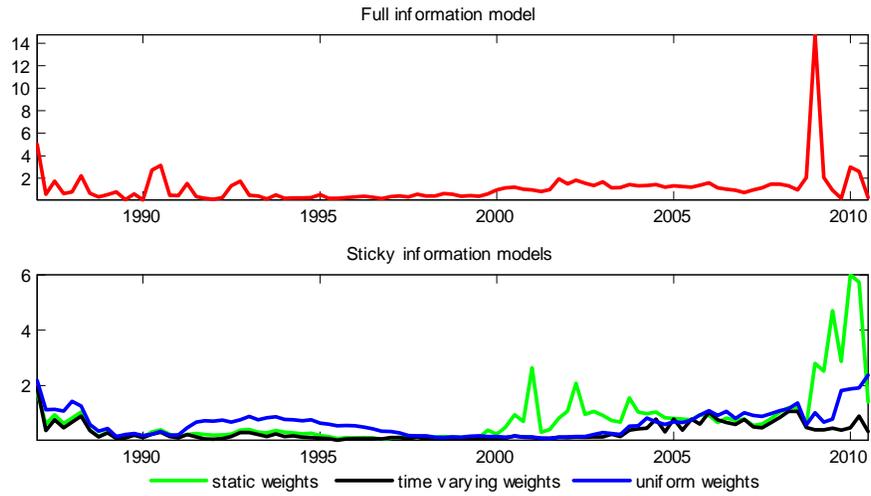


Figure 10: Kullback- Leibler distance measure

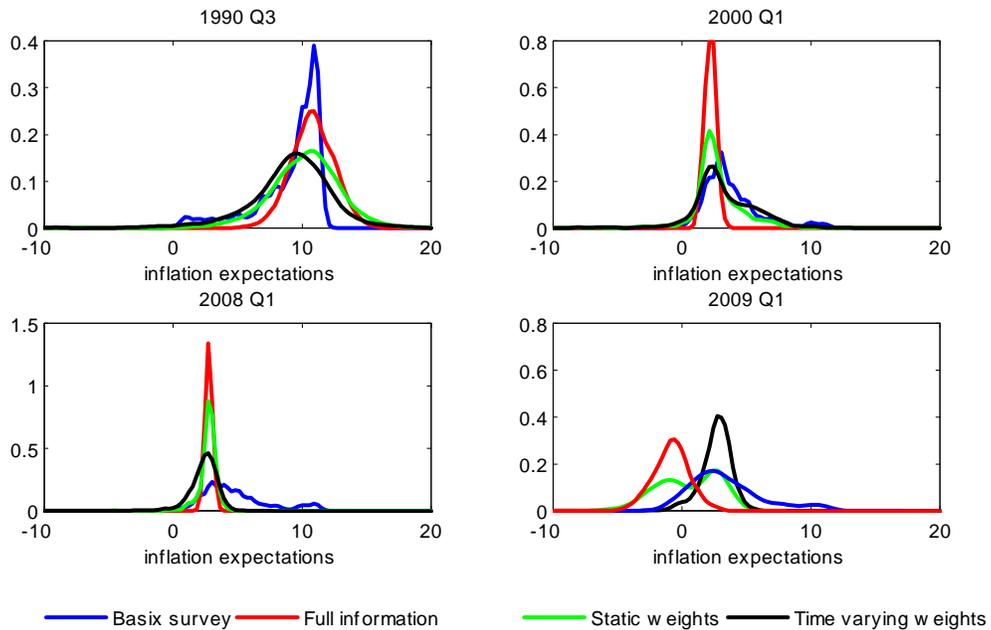


Figure 11: Estimates of density of inflation expectations for various sticky information models.

that the highest share of agents update each period.

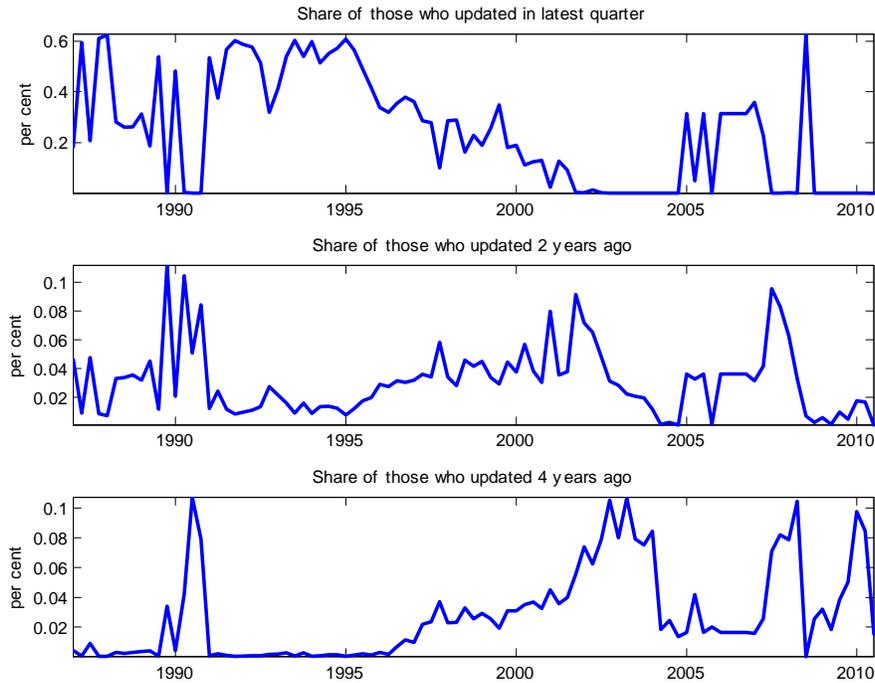


Figure 12: Share of those who update at different points in time

To conclude this section, the model that fits best the Barclays Basix survey best is a time varying sticky information model. The full information model performed worst indication that a forecast as simple as the one produced by this model does not generate enough dispersion and does not fit the mean as well as the sticky information models. In this section we use the Kullback-Leibler probability distance measure to assess the fit of the model. However the time varying sticky information is likely to be overfitted and so further work in finding a way to penalise overparametrisation is required.

#### 4.4 Assessing the sticky information model: combining macro and micro data

Our analysis so far has compared means and dispersion of inflation expectations predicted by a sticky information model with those from the surveys. In this subsection, we link this macro view of the sticky information model with a micro-founded approach. The Bank NOP survey contains rich information about individual specific characteristics, such as age, education, gender, whether or not a respondent has a mortgage etc. One would expect that individual specific characteristics influence the cost of updating information sets and therefore the frequency with which this updating occurs for any specific individual. It is therefore a natural question to ask if we can relate the personal characteristics of individuals forming inflation expectations in the survey to those who make VAR forecasts using particular vintages of information.

The theory would receive more support if those respondents that one would expect to have very low costs

of updating their information sets reported inflation expectations that correlate well with the VAR forecasts of very recent information set vintages. Conversely, based on the theory one would expect that those households with very high costs of updating information sets would - more often than not - make forecasts that are based on rather outdated information sets. Their responses should thus - more often than not - correlate with the VAR forecasts of older vintages. Of course, even infrequent updaters will on occasion use the most recent information set. This exercise is aimed at finding out what are the characteristics of those respondents that have expectations consistent with to VAR forecasts based on certain vintages of information.

It is important to note that this conceptual framework is a slight departure from the baseline sticky information model of Mankiw and Reis (2002). In that model all agents are the same except for the timing when they update their information set. In this subsection, we still assume that agents are heterogeneous in the timing that they update, but they also could differ in their personal characteristics. We will therefore simply use the VAR forecasts of each vintage of the information set and make no assumption about the distribution of these forecasts in the survey.

In particular, we match the NOP survey data for those years when we have individual specific information with the empirical predictions generated by the constant coefficient VAR for RPIX.<sup>4</sup> We define three broad groups in relation to their frequency of updating their information set:

- ‘frequent updaters’: individuals who last updated their information set between the previous quarter and the 4 quarters before that, i.e. approximately in the previous year;
- ‘infrequent updaters’: individuals who last updated their information set between 6 and 15 quarters before, i.e. over a year before but no more than 4 years before ;
- ‘rare updaters’: those who last updated between 16 and 20 quarters earlier, i.e. between 4 and 5 years before.

In the survey, we generate one dummy variable for each of these groups, which takes value one if the survey respondent gave an answer within one percentage point (i.e.  $\pm 0.5$ ) of the value predicted by the constant coefficient VAR with different frequencies of updating. We also generate a dummy variable for the individuals who are not able to form an inflation expectation and therefore answer ‘Don’t know’ in the survey.

On average across time, around half of the survey respondents have a match in the model-based distribution. Of the remaining half, around 40% gave an expectation that does not have a match in the model-based distribution, and another 10% could not formulate an inflation expectation and answered “Don’t know”. We then identify the survey responses which have a unique match in the model-based distribution, such that a survey respondent could only be one of the frequent/infrequent/rare updater/‘don’t know’ type.

We run a probit regression with the four dummies ‘frequent updaters’, ‘infrequent updaters’, ‘rare updaters’ and ‘don’t know’ on the personal characteristics included in the NOP survey: age, gender, education level, income level, whether in work and whether the respondent owns a house outright, has a mortgage, rents or else. Figure 10 in the Appendix reports the regression results. One striking difference stems from a comparison

---

<sup>4</sup>This individual specific information is available annually in 2001 and 2002 and quarterly from 2003 onwards.

of those who have formed an inflation expectation (column 1-3) versus those who have not (column 4). Those who have formed an expectation, regardless of how often they update their information set, tend to be male, have a degree, own a house and/or have a mortgage and to be on average incomes. Those who do not form an inflation expectation tend to be female, not to have a degree or even middle education, not to own a house or to have a mortgage and tend to be on lower incomes (i.e. the baseline category of less than £9,500).

The main difference among those who have formed an inflation expectation, instead, is that the likelihood of having a degree increases with the frequency of updating the information set, with ‘frequent updaters’ being 8 percent more likely to have a degree than ‘rare updaters’. In other words, the ‘frequent updaters’ are most likely to have a degree than the rest of the population.

Housing tenure also tend to have some significant correlation with the frequency of updating information. The ‘frequent updaters’ are the group most likely to have a mortgage, but not to own their house outright or rent. This may be associated with lower costs of tracking inflation because the individual would already be tracking the economic conjuncture more broadly in order to assess the best way to finance his debt. The ‘infrequent updaters’ are also likely to have a mortgage, but could also own their house outright. And the ‘rare updaters’ could have either of those or might as well be renting. The other control variables included appear to be characterize the individuals less precisely. Age and income do not have a clear pattern of significance, nor does whether a person is in work or not.

Overall, the results from the probit regression lend some support to the sticky information hypotheses. One would expect that those who hold a degree have smaller costs of updating their information sets and would therefore do so more frequently. It appears reasonable that individuals with a degree have easier and more regular access to newspapers and other media that contain information about current inflation. Current inflation in turn is often used as a naive predictor for future short term inflation. The probit regression result that associates people who hold a degree with frequent updaters thus provides some support for the theory of sticky information.

## 5 Conclusion

Our paper asks whether the sticky information framework of Mankiw and Reis (2002) is consistent with the survey data of U.K. household inflation expectations. We find moderate support for the theory in the data. The survey data shares some key characteristics predicted by the theory, such as a strong increase in cross-sectional dispersion in periods where the macro-data is more volatile over the recent past. Our estimates of the speed of information updating are in line with those reported for other countries and suggest that U.K. households update their informations sets about once a year on average.

In our last exercise we use non-parametric methods to assess how well each model fits. We performed two exercises using the Barclay Basix survey. First, by assuming that everybody is fully informed . We call this the full information exercise. Second, by assuming that not everybody is fully informed and as a result the

mean of the estimated inflation rate reflects forecasts made with outdated information. We call this the sticky information exercise. Within, this we allow the share with which the forecasts produced by different information vintages are weighted to vary over time. We find that the full information model performs much worse than any of the versions of sticky information models and that, over all, the sticky information model with time varying weights captures the time variation of the mean and dispersion of the Barclay Basix survey of inflation expectations reasonably well. We also find that the distribution of agents is not always geometric and thus sometimes the highest proportion of households update, on average, only infrequently. This is in contrast with Mankiw and Reis (2002) model that predicts that the highest share of agents update each period.

And we have found mild additional support for the theory by looking at the micro characteristics that make respondents more likely to be frequent updaters of information sets. Personal characteristics such as education could plausibly reduce the costs of updating information sets and therefore increase the frequency of information update, as one would expect from the theory.

Overall, our results suggest that analysing models of information stickiness is a valuable initial step towards understanding the rich dynamics of inflation expectations. However future research that will consider model uncertainty as well as different assumption about the basket of goods households could be considering when making expectations will further enhance our understanding of the process behind inflation expectation formation.

## 6 Appendix:

### 6.1 Results from the probit model

Depvar:	'frequent updaters'	'infrequent updaters'	'rare updaters'	'don't know'
	(1)	(2)	(3)	(4)
Male	<b>0.038289***</b> (0.0082)	<b>0.024276***</b> (0.0068)	<b>0.017681**</b> (0.0074)	<b>-0.029236***</b> (0.0027)
Edu: Degree	<b>0.085341***</b> (0.0139)	<b>0.023379**</b> (0.0112)	-0.017234 (0.0115)	<b>-0.028980***</b> (0.0039)
Edu: Secondary	<b>0.044008***</b> (0.0108)	0.014493 (0.0089)	-0.002338 (0.0096)	<b>-0.019379***</b> (0.0035)
Age: 25-34	<b>0.036027*</b> (0.0199)	-0.017781 (0.0149)	0.014330 (0.0175)	-0.007384 (0.0057)
Age: 35-44	<b>0.050454**</b> (0.0201)	-0.014633 (0.0151)	0.018684 (0.0176)	<b>-0.015975***</b> (0.0055)
Age: 45-54	<b>0.051714**</b> (0.0210)	0.027974 (0.0171)	0.021162 (0.0184)	<b>-0.031399***</b> (0.0053)
Age: 55-64	<b>0.049813**</b> (0.0223)	0.022180 (0.0180)	0.013775 (0.0193)	<b>-0.030599***</b> (0.0054)
Age: 65+	0.017360 (0.0221)	-0.017654 (0.0169)	-0.005186 (0.0191)	0.001259 (0.0068)
Own house outright	0.018407 (0.0142)	<b>0.025326**</b> (0.0122)	<b>0.033273**</b> (0.0133)	<b>-0.022740***</b> (0.0042)
Has a mortgage	<b>0.025938**</b> (0.0124)	<b>0.023919**</b> (0.0106)	<b>0.020861*</b> (0.0115)	<b>-0.025032***</b> (0.0040)
Rents	<b>-0.061831***</b> (0.0128)	0.011728 (0.0117)	<b>0.038794***</b> (0.0130)	0.002907 (0.0043)
Working	-0.014484 (0.0109)	-0.002757 (0.0090)	0.003159 (0.0099)	0.000948 (0.0036)
Income: £9.5k-£17.5k	<b>0.032427***</b> (0.0124)	0.016852 (0.0103)	<b>0.038770***</b> (0.0112)	<b>-0.026809***</b> (0.0033)
Income: £17.5k-£25k	<b>0.039696**</b> (0.0154)	<b>0.045301***</b> (0.0135)	<b>0.031025**</b> (0.0140)	<b>-0.031425***</b> (0.0039)
Income: >£25k	<b>0.087035***</b> (0.0144)	<b>0.037034***</b> (0.0120)	<b>0.025025*</b> (0.0128)	<b>-0.051018***</b> (0.0041)
Pseudo-R2	0.0249	0.0104	0.0034	0.0366
Obs.	12,832	12,832	12,832	45,655

The dependent variables in columns 1 to 3 are dummies with value equal to one for each NOP survey respondent who has one (and only one) match in the empirical distribution generated by the constant coefficient VAR for RPIX.

'Frequent updaters' have a unique match in the empirical distribution generated by updating the information set between the previous quarter and the 4 quarters before that; 'infrequent updaters' and 'rare updaters' have a unique match respectively with those who last updated between 6 and 15 quarters before and 16 to 20 quarters before.

Coefficients are marginal effects from probit estimation, with standard errors in parentheses.

The asterisks indicate significance at the 10% (\*) , 5% (\*\*) or 1% (\*\*\*) level.

Baseline categories are: female, low education (GCSE in UK), age 15-24, 'other' housing tenure, not working, income <£9,500.

Figure 13: Estimation results for the Probit model

## 6.2 Description of the time varying parameter model

We then move on to assume that households put variable weight on each new piece of information depending, for example, on the past forecast errors. Therefore under this scenario, households are assumed to run a time varying VAR with stochastic volatility of the following in 2. Although the description of the model below may look complicate and unrealistic as a model of households' expectations formation, this could be however a reasonable way of thinking about expectations insofar this model could be a reflection of the fact that expectations are formed adaptively at each point in time. Branch (2007) uses a similar system (albeit without stochastic volatility) and fit it to the Michigan survey of inflation expectations. The system we estimate can be described as following:

$$Y_t = X_t' \theta_t + \varepsilon_t \quad (2)$$

where  $Y_t$  is a vector of observed endogenous variables and the matrix  $X_t$  includes one lag of  $Y_t$  and constants,  $\theta_t$  collects the time varying parameters and  $\varepsilon_t$  are assumed to be zero-mean normally distributed, with time varying covariance matrix  $R$  such that  $\varepsilon_t = R^{1/2} \zeta_t$ . The normalized innovations  $\zeta_t$  are assumed standard normal and we will describe  $R$  in more detail below.

The VAR parameters follow a driftless a random walk with a reflecting barrier<sup>5</sup> that keeps them from entering regions of the parameter space where the VAR is explosive and evolve according to:

$$p(\theta_t | \theta_{t-1}, Q) = I(\theta_t) f(\theta_t | \theta_{t-1}, Q) \quad (3)$$

$I(\theta_t)$  is an indicator function rejecting unstable draws<sup>6</sup> and  $f(\theta_t | \theta_{t-1}, Q)$  is given by:

$$\theta_t = \theta_{t-1} + \eta_t \quad (4)$$

where  $\eta_t$  is  $N(0, Q)$ .

Turning to the VAR's reduced form innovations, following the literature on multivariate stochastic volatility models (see for example Jacquier et al. (1994, 1999)), we specify the drifting variance as:

$$var(\varepsilon_t) = R_t = A_t^{-1} H_t (A_t^{-1})' \quad (5)$$

---

<sup>5</sup>This has become standard when estimating TVP models. See for example, among others, Cogley and Sargent (2002), Cogley and Sargent (2005), Primiceri (2005).

<sup>6</sup>The function  $I(\theta_t) = 0$  when the roots of the associated VAR lag polynomial lie inside the unit circle and  $I(\theta_t) = 1$  otherwise. This is a stability condition for VAR, representing an *a priori* belief about the implausibility of explosive representations for the variables in system.

$H_t$  and  $A_t$  are time varying matrices which are defined as:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 & 0 \\ 0 & 0 & 0 & h_{4,t} & 0 \\ 0 & 0 & 0 & 0 & h_{5,t} \end{bmatrix} \quad A_t \equiv \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 & 0 \\ \alpha_{51,t} & \alpha_{52,t} & \alpha_{53,t} & \alpha_{54,t} & 1 \end{bmatrix}$$

with  $h_{i,t}$  following geometric random walks

$$\ln h_{i,t} = \ln h_{i,t-1} + v_{i,t} \quad (6)$$

We define  $h_t \equiv (h_{1,t}, h_{2,t}, h_{3,t}, h_{4,t}, h_{5,t})'$ . Following Primiceri (2005), we postulate the non-zero and non-one elements of the matrix  $A_t$ - which we collect in the vector  $\alpha_t \equiv (\alpha_{21,t}, \alpha_{31,t}, \dots, \alpha_{43,t})'$ - to evolve as driftless random walks,

$$\alpha_t = \alpha_{t-1} + \tau_{i,t} \quad (7)$$

Following Primiceri (2005) we assume the vector  $(u'_t, \eta'_t, \tau'_t, v'_t)'$  to be distributed as:

$$\begin{bmatrix} u_t \\ \eta_t \\ \tau_t \\ v_t \end{bmatrix} \sim N(0, V), \text{ with } V = \begin{bmatrix} I_4 & & & \\ & Q & & \\ & & S & \\ & & & Z \end{bmatrix} \text{ and } Z = \begin{bmatrix} \sigma_1^2 & & & & \\ & \sigma_2^2 & & & \\ & & \sigma_3^2 & & \\ & & & \sigma_4^2 & \\ & & & & \sigma_5^2 \end{bmatrix} \quad (8)$$

where  $u_t$  is such that  $\varepsilon_t \equiv A_t^{-1} H_t^{1/2} u_t$ . The reason Primiceri (2005) assumes a block-diagonal structure for  $V$  is twofold. On one hand the model is already heavily parametrised and on the other hand Primiceri (2005) stresses that allowing for a 'completely generic correlation structure among different sources of uncertainty would preclude any structural interpretation of the innovations'. Following Primiceri (2005) again, we simplify the problem by assuming a block-diagonal structure for  $S$  as follows:

$$S \equiv Var(\tau_t) = Var(\tau_t) = \begin{bmatrix} S_1 & 0_{1x2} & 0_{1x3} & 0_{1x4} \\ 0_{2x1} & S_2 & 0_{2x3} & 0_{2x4} \\ 0_{3x1} & 0_{3x2} & S_3 & 0_{3x4} \\ 0_{4x1} & 0_{4x2} & 0_{4x3} & S_4 \end{bmatrix} \quad (9)$$

with  $S_1 \equiv var(\tau_{21,t})$ ,  $S_2 \equiv var([\tau_{31,t}, \tau_{32,t}]')$ ,  $S_3 \equiv var([\tau_{41,t}, \tau_{42,t}, \tau_{43,t}]')$  and  $S_4 = var([\tau_{51,t}, \tau_{52,t}, \tau_{53,t}, \tau_{54,t}]')$  which implies that the non-zero and non-one elements of  $A_t$  belonging to different rows evolve independently<sup>7</sup>.

<sup>7</sup>Primiceri (2005) explains in Appendix A.2 that this assumption simplifies the inference and therefore allows us to do Gibbs

We estimate 2-3 using Bayesian methods. Appendix A. 1 describes the choice of priors and the Markov-Chain Monte Carlo algorithm we used to simulate the posterior distribution of the hyperparameters and the states conditional on the data.

## 7 References

Branch, W. (2004), 'The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations. *Economic Journal*, 114(497), pp. 592-621.

Branch, W. (2006), 'A simple recursive forecasting model'. *Economics Letters*, 91, pp1580- 166

Carroll, C. (2003a), 'Macroeconomic expectations of households and professional forecasters', *Quarterly Journal of Economics*, 118(1), pp. 269-98.

Carroll, C. (2003b). The epidemiology of macroeconomic expectations. In *The Economy as an Evolving Complex System, III*, Larry Blume and Steven Durlauf (eds.). Oxford, England: Oxford University Press.

Branch, W.A, Evans, G. W., (2006). 'A simple recursive forecasting model'. *Economics Letters*. 158-166.

Branch, W.A. (2007). 'Sticky information and model uncertainty in survey data on inflation expectations'. *Journal of Economic Dynamics and Control*, 31 (2007), pp 245–276.

Cogley, T., Morozov, S., Sargent T, (2005). 'Bayesian fan charts for U.K. inflation: forecasting and sources of uncertainty in an evolving monetary system'. *Journal of Economics Dynamics and Control*, 29 (2005), pp. 1893-1925.

Geweke, J. (1989). 'Bayesian inference in econometric models using Monte Carlo integration'. *Econometrica* 57, 1317-1339.

King, R.G. (2003). 'Disagreement about Inflation Expectations: Comment'. *NBER Macroeconomics Annual*, Vol. 18 (2003), pp. 248-257.

Mankiw, N. G., Reis, R., Wolfers, J., 'Disagreement about Inflation Expectations'. *NBER Macroeconomics Annual*, Vol. 18 (2003), pp. 209-248.

Mankiw, N. G., Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4):1295-1328.

Morris, S., Shin, H. S. (2006), 'Inertia of Forward-Looking Expectations'. *The American Economic Review*, Vol. 96, No. 2 (May, 2006), pp. 152-157.

Pfajfar, D., Santoro, E., (2006), 'Heterogeneity and learning in inflation expectation formation: an empirical assessment', *Universita degli studi di Trento, Discussion Paper* 7.

Pfajfar, D., Santoro, E., (2008), 'Asymmetries in Inflation Expectation Formation Across Demographic Groups', *Cambridge Working Papers in Economics*, No. 0824.

Robertson, J. C, Tallman, E. W., Whitman, C., (2005). 'Forecasting using relative entropy'. *Journal of Money Credit and Banking*, Vol. 37, pp. 383-402.

---

sampling on the non-zero and the non-one elements of  $A_t$  equation by equation.

Souleles, N. (2001). Consumer sentiment: Its rationality and usefulness in forecasting expenditure-Evidence from the Michigan micro data. *Journal of Money, Credit and Banking* forthcoming

Williams, J. C., 'Disagreement about Inflation Expectations: Comment', *NBER Macroeconomics Annual*, Vol. 18 (2003), pp. 257-268.