



EUROPEAN CENTRAL BANK

EUROSYSTEM



WORKING PAPER SERIES

NO 1475 / SEPTEMBER 2012

INFORMATION FLOWS AND DISAGREEMENT

Cristian Badarinza
and Marco Gross



In 2012 all ECB
publications
feature a motif
taken from
the €50 banknote.



NOTE: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Cristian Badarinza

at Goethe-Universität Frankfurt, Senckenberganlage 31, 60325 Frankfurt am Main, Germany;
e-mail: badarinza@wiwi.uni-frankfurt.de

Marco Gross

at European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany; e-mail: marco.gross@ecb.europa.eu

© European Central Bank, 2012

Address

Kaiserstrasse 29, 60311 Frankfurt am Main, Germany

Postal address

Postfach 16 03 19, 60066 Frankfurt am Main, Germany

Telephone

+49 69 1344 0

Internet

<http://www.ecb.europa.eu>

Fax

+49 69 1344 6000

All rights reserved.

ISSN 1725-2806 (online)

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from <http://www.ecb.europa.eu> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=2148383.

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, <http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html>

Abstract

The aim of this study is to assess the extent to which the degree of heterogeneity of inflation expectations is driven by the flow of information related to current and future price developments. To that end, we follow three routes: i) We propose different measures of information flow that have either a sender or a receiver perspective; ii) We present empirical results for the US and selected EU countries that aim to corroborate the hypothesis that news have the ability to densify expectations, i.e. to reduce forecast heterogeneity; and iii) We augment some otherwise standard models of expectation formation by allowing the individual updating frequency to depend on the observed measure of information flow; since the updating frequency is higher at times of high inflation and decreasing thereafter, this mechanism can contribute to upward biases in inflation expectations over long periods of time.

Keywords: Heterogeneous beliefs, disagreement, public information

JEL classification: D12, D84, E31, E37

Non-technical summary

This paper aims at analyzing the role of one potential determinant, the intensity of information flow, on belief heterogeneity about the future course of price inflation. The underlying assumption is that the economy comprises a large portion of individual agents that face frictions with respect to the arrival or processing of information, either due to delays from sender or limited capability of processing from a receiver perspective.

The hypothesis that we aim to substantiate is that information which agents require to update their expectation does not flow at a constant rate but varies systematically over time, and in the sequel determines the level of disagreement throughout the population. More news shall induce people to agree more.

We employ survey data for the US and a small panel of EU countries to construct a simple measure of disagreement in expectations based upon which we test the hypothesis that information has an impact on the level of agreement, while controlling for the general macroeconomic environment. For obtaining measures of information flow, alternative sources are taken into consideration which allow us to take either a *sender* or *receiver* perspective and compare the resulting measures. For the US, the two measures follow very similar paths over time which is of avail for making the empirical findings from the paper more robust. We find clear evidence in favor of the hypothesis that news are able to densify expectations, i.e. to reduce belief heterogeneity.

Besides the empirical analysis, the purpose is to augment some otherwise conventional models of expectation formation that feature some form of friction with regard to the processing capability of agents, to incorporate the fact that the individuals' update frequency be depending on the intensity of information flow. The sticky information as well as the epidemiological model are better able to replicate the observed patterns in disagreement compared to the original model set-ups where the update frequency was assumed to be constant through time.

The finding that peoples' attention to, or the denseness of information that

they are faced with, increases at times of high inflation and tends to fall during subsequent periods of disinflation offers an additional rationale for why the aggregate expectation may be biased upward over long periods of time. Indeed, we can observe that update frequencies rise at periods of high inflation, fall thereafter, and in the sequel let expectations be sticky and take longer to adjust to falling inflation.

1 Introduction

The fact that households have access to different information sets and use different models when forming expectations has for long been rather neglected by macroeconomic theory. This was perhaps a fortunate development, in particular after the advent of the rational expectations hypothesis in the 1970s, since it permitted building analytically elegant modeling frameworks that were tractable, directly usable for policy analysis and reasonably well equipped to match key properties of the data. However, as soon as detailed survey data on expectations became available, researchers' attention started to concentrate on the determinants and effects of private information and disagreement, as well as on behavioral responses to news or policy announcements, a trend that has gained momentum in recent years.

The theoretical basis for our analysis lies in the nature of the information acquisition process at individual level. Detailed micro-foundations for a model in which information disseminates only slowly through the population have been proposed and popularized by Mankiw and Reis (2002) who develop a framework which is now commonly referred to as the *Sticky Information* (SI) model. Besides the model's various implications for aggregate dynamics and the effect of monetary policy, it implies that agents should have different expectations as a result of the staggered diffusion of information. Only a fraction of agents update their expectations every period, while the remainder of the population continues to form expectations based on outdated information. The SI model setting has found some support, though in a somewhat different conceptual framework, by Carroll (2003): It is an *epidemiological model* (EPI) that parallels the spread of information through the population with that of a disease. The assumption is that individual agents do not form an expectation on their own but rather adopt, i.e. "get infected with", the views of professional forecasters that are conveyed via the media. The difference between the EPI and the SI model is that in the former agents update toward professional forecasts while in the latter towards the

latest rational forecast that agents form themselves.¹ The calibration and simulation exercise that will be presented in the second part of our paper attempts to shed light on these models' properties.

Another strand in the literature, an early one, that the present study is related to, is the one by McCombs and Shaw (1972) on the agenda-setting function of mass media. The central point is that media can have a marked impact on people's awareness of certain topics, where an assumption is that concentration on salient issues leads the population to perceive this issue as more relevant. Importantly, the theory rests also on the assumption that media can shape news in a way that may distort reality. As most of the related work that has appeared since McCombs and Shaw, they explore the theory in a political context and find a positive relation between news intensity and what voters found most relevant in political campaigns.

From an empirical perspective, only few papers have so far attempted to analyze what the determinants of disagreement are. Mankiw *et al.* (2003) present stylized facts and empirical regularities in survey measures of disagreement for the US and analyze as well their relation to macroeconomic variables. They demonstrate that the SI model is capable of explaining observed patterns in the level and the dispersion of survey expectations. One concrete finding in Mankiw *et al.* (2003), however, is that the SI model is not able to replicate the apparent positive relationship between disagreement and level inflation that can be found in the data. We shall later argue that this finding, and the model's inability to reproduce the empirical regularity, may be an artifact of the chosen measure of disagreement. A broader overview on the determinants of disagreement not only for the case of inflation but also real variables such as GDP, investment and employment can be found in Doornik *et al.* (2009).

Maag and Lamla (2009) employ a Bayesian learning model in which media coverage of inflation affects forecast disagreement by influencing both the information sets as well as the predictor choice. In their model, agents update their

¹Yet another theoretical framework has been proposed by Roberts (1997); here, the updating is done with reference to the past realization of inflation.

prior expectations about inflation by absorbing news transmitted by television and newspapers, while acknowledging that media reports contain only noisy signals about future inflation. Agents therefore face a signal extraction problem which is solved through Bayesian updating. Moreover, they allow for heterogeneous forecasting models along the lines of Kandel and Zilberfarb (1999). The approach is innovative particularly with regard to the analytical differentiation between the volume of news and their content: More news shall induce the agent to put less weight on prior beliefs, but it is the specific content which determines heterogeneity and disagreement at aggregate level. Thus, in terms of testable implications the model suggests that i) both a higher volume of media reporting and a lower heterogeneity (information entropy) of the statements about inflation lead to lower forecast disagreement as agents converge more and more to the same information set and ii) that if all media reports contain the identical message, the variance of the noise component collapses to zero, agents end up choosing identical predictors and at aggregate level the cross-sectional dispersion of expectations decreases. The empirical findings that we present in the present paper stand somewhat in contrast to Maag and Lamla (2009) who conduct their analysis for households and professional forecasters in Germany: unlike for Germany, we find that disagreement among citizens of the US and the selected group of EU countries does depend on media coverage; a more intense information flow induces people to agree more.

The contribution of our paper can be seen along three dimensions: First, we present a set of alternative schemes to measure information flow and in particular do we distinguish between a *sender* and a *receiver* perspective of information.² It turns out that the two distinct sources of information, one representing the sender and the other one the receiver side, result in very similar measures of news intensity. This provides us with confidence in using one or the other measure for subsequent empirical analysis and also corroborates the avail of such measures for European countries where the receiver measure as such is not available. Second, we present a set of regression results which aim to address the hypothesis

²Our work also comes close to the approach of Veldkamp (2006) in terms of the quantification of the sender side of the information flow through measuring mass media news intensities.

that information flow has the ability to impact the width of the cross-sectional distribution of expectations. The finding is that an increasingly intense information flow induces agents to agree more about the future. Third, we augment a set of standard micro-founded models of disagreement by allowing the update frequency to be time-varying, in particular by linking it directly to our measure of information flow, to then show that the sticky information as well as the epidemiological model set-up are better able to replicate the observed patterns in disagreement relative to the original model settings where the updating frequency was assumed to be constant through time.

2 Data and methodology

The main data source that the empirical section of the paper refers to is the US Michigan Survey of Consumers. We draw upon the cross-sectional archive of monthly survey waves (covering the period between January 1987 and June 2011), each containing a set of recurrent questions that address different aspects of consumer attitudes and expectations. The monthly cross-sectional samples cover a pool of approximately 500 individual respondents, chosen such as to be representative of the US population excluding Alaska and Hawaii. The surveys are released during the last week of a month. Macroeconomic data for the US and the EU countries were retrieved from the St. Louis Fed's FRED[©] and the ECB database, respectively.

For the European Union countries, we refer to the European Commission's Business and Consumer Survey to obtain survey data which allow us monitoring the perception of consumers with respect to past, current, as well as expected future conditions in Europe. Five sub-surveys, addressing the manufacturing industry, construction, consumers, retail trade and services are conducted on a monthly basis. Only data from the consumer section of the survey is relevant for our analysis and only a subset of questions therein will in fact be taken into consideration.

For obtaining further measures of news intensity (alternative to one based on the Michigan Consumer Survey), we use Google’s *Insights for Search* tool and the professional news service provider *Factiva*.

Quarterly data on inflation expectations by professional forecasters are obtained from the United States Survey of Professional Forecasters which we convert to monthly frequency by means of a linear interpolation scheme. The reference series measures expectations of changes in consumer price inflation (*CPIA*) at the 1-year horizon and covers the period from January 1982 to June 2011. For the analysis that will be presented later in the paper, we use the median of the cross-section of responses that were recorded over time.

2.1 Quantifying disagreement

The Michigan Survey of Consumers contains two questions that address price expectations based upon which we construct measures of central tendency and dispersion of expected inflation.

Question **PX1Q2** (henceforth referred to as **PX1**) reads as follows:

By about what percent do you expect prices to go up/down on the average during the next 12 months?

Respondents are supposed to provide a point estimate in percent or may choose to answer *don’t know*. Based on the cross-sectional distribution of answers to this question, we compute the standard deviation of the individual cross-sectional observations of expected inflation to which we refer to as *quantitative* disagreement.

$$\sigma_t^q = stdev(\{\pi_{i,t+12}^e\}_{i=1,\dots,I_t})$$

with $\pi_{i,t+12}^e$ denoting the 1-year ahead point forecast of individual i at time t and $stdev$ being the standard deviation. Such quantitative measures have been

used e.g. in Carroll (2003) and Mankiw *et al.* (2003).

An alternative measure of disagreement can be derived from question **PX1Q1** which is phrased as follows:

During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now?

Respondents can choose between four answer categories: *Go up* (e^1), *same* (e^2), *go down* (e^3), or *don't know*. To quantify disagreement we here employ the following measure:

$$\sigma_t^c = \sum_{i=1}^2 F_t^i (1 - F_t^i) \quad (1)$$

with F_t^i being the cumulative relative frequencies for the i -th answer category at time t . Note that the third answer category is excluded from the sum since the cumulative frequency is one and therefore does not contain additional information on the distribution of the response shares.³

This measure has been proposed by Lacy (2006) to whom we refer for details.⁴ Importantly, the measure is ordinal in nature, i.e. it does not require one to assume that the distance between categories be equal. Statistics such as a standard deviation shall not be applied here because one would have to have a variable measured at interval scale. The distances between response categories cannot be quantified, nor shall one assume that they are equally far from one another. The maximum refers to a fully *polarized* distribution (as Lacy (2006) refers to this case) in which all responses fall into two response categories to equal shares. The other extreme arises when all responses fall into a single category (full agreement); the measure will reach its minimum at 0 in this case. To this

³The first three answer categories have been rescaled so as to sum to unity at every point in time, i.e. the *don't know* answer share is evenly distributed among the three remaining categories.

⁴The quantification method has also been used for obtaining disagreement measures in the context of central bank communication and transparency by Ehrmann *et al.* (2010).

ordinal measure of disagreement we refer to as *categorical* in the following.

Figure 1 shows the two quantitative measures (standard deviation and interquartile range) along with the categorical (ordinal) measure of disagreement for the US.

The categorical disagreement measure, we argue, is at least as appropriate as the interquartile range for quantifying the extent of heterogeneity in beliefs. The only paper that we are aware of that employs a similar ordinal measure of dispersion is Maag and Lamla (2009); it is the so-called index of qualitative variation that these authors compute.

In Figures 2 and 1, a positive relationship between level inflation and the quantitative measure of disagreement can be observed; this empirical regularity has been documented by various authors, including e.g. Mankiw *et al.* (2003). The strong positive correlation is interpreted perhaps righteously as reflecting the fact that inflation uncertainty rises with higher levels of inflation. We think, however, that as far as the macroeconomic implications on consumer behavior are concerned, drawing conclusions on the basis of this relationship may be misleading. Consumers that answer a telephone interview question about what they expect inflation will be in the near future may not be very precise if asked to provide a precise level expectation, while their answer would supposedly be more reliable and less subject to noise when asked to position themselves in ordinal categories. For example, at some selected point in the sample, about 2% of agents stated they expect prices to go down, with - at the same time - the average answer to the quantitative question slightly exceeding 5%. Thus, if one had to decide which of the two answers to trust more, the ordinal one should supposedly be preferred. Moreover, in terms of the economic interpretation, the same extent of dispersion, say, at 5 percentage points, around an average inflation expectation of, say, 10%, versus that same dispersion around a 1% level may reflect quite different realities and generate quite different consumer behavior, since the latter would imply that some portion of agents would in fact consider deflation likely to occur.

For the European Union countries, Questions 5 and 6 of the questionnaire are relevant for our analysis (we refer to them as Questions A and B). They read as follows.

Question A: *How do you think that consumer prices have developed over the past 12 months? They have...*

- ++ *risen a lot*
- + *risen moderately*
- = *risen slightly*
- *stayed about the same*
- *fallen*
- N *don't know*

Question B: *By comparison with the past 12 months, how do you expect that consumer prices will develop in the next 12 months? They will...*

- ++ *increase more rapidly*
- + *increase at the same rate*
- = *increase at a slower rate*
- *stay about the same*
- *fall*
- N *don't know*

Analogously to the quantification method applied to the Michigan Survey data, to measure the degree of heterogeneity in beliefs in the EU countries (concerning current as well as future conditions) we use again the Lacy (2006) measure:

$$\sigma_t^{c,p} = \sum_{i=1}^4 F_t^{p,i} (1 - F_t^{p,i}) \quad \text{and} \quad \sigma_t^{c,e} = \sum_{i=1}^4 F_t^{e,i} (1 - F_t^{e,i}) \quad (2)$$

where $F_t^{p,i}$ and $F_t^{e,i}$ are the cumulative relative frequencies for the i -th category at time t from question A and B. Figure 8 shows the resulting measures based

on current perceptions and expectations across the group of EU countries.

2.2 Quantifying information flow

In order to measure the flow of information related to current and expected future price developments we follow three routes.

First, we refer once more to the Michigan Survey, that is, to questions **NEWS1** and **NEWS2**. They read as follows:

During the last few months, have you heard of any favorable or unfavorable changes in business conditions? What did you hear?

Respondents can choose two items out of 80, six of which we consider being related to prices, price inflation respectively.

- ...
- n^{31} *Lower/stable prices, less inflation*
- n^{32} *Higher prices, inflation is good*
- n^{37} *Other references to prices/credit*
- n^{71} *Prices falling, deflation*
- n^{72} *Prices high, inflation*
- n^{77} *Other price/credit references*
- ...
- n/a *don't know*

Based on the answers to this question, we compute a monthly share of agents who were considering changing price conditions relevant. This first measure of information flow can be thought of as having a **receiver** perspective and will in the following be referred to as *survey-based news*. We denote the value of this measure in period t as f_t^s , for later reference.

Second, we use Google's *Insights for Search* tool which can be used to analyze search patterns for self-specified geographical areas. The keyword that we instruct the search engine to analyze is *inflation* to then obtain a weekly time

series in the form of an index, with the maximum over the sample period being normalized to 100. The weekly frequency is converted to monthly by taking period averages. Of this second measure of information flow one can think as also having a **receiver** perspective, though it is different in nature compared to the survey-derived measure because agents put own effort into the search via Google.

Third, we employ the inflation news intensity measure from Badarinza and Gross (2009). The professional news service provider *Factiva* allows to retrieve the number of articles that contained the term *inflation* in their headlines or lead paragraphs which are then divided by the number of news contained in the parent directory which counts all economic news. Thus, we obtain a ratio of inflation-related news that appear in print and online media with a monthly frequency back until January 2000 both for the US and the selected set of EU countries. This source of information flow can be thought of as having a **sender** perspective. It will in the following be referred to as *public news*. The value of this measure in period t will be denoted as f_t^p , for later reference.

Figures 3 and 4 illustrate how the two receiver and the one sender information flow series compare. In particular, the survey-based receiver-side measure and the Factiva sender measure of information flow appear to follow very similar paths over time, except during the beginning of the 1990s. During the latter part of the sample period, starting from the late 1990s, the correlation between the two is rather close to one. As can be seen in Figure 4, the Google news intensity measure deviates slightly from the path drawn by the other two measures; its correlation with the survey-based and the public news measures equals .4 and .5, respectively. Despite originating from very different sources, it turns out that the three distinct measures of information flow, representing either a sender or a receiver side, result in very similar measures of news intensity. This also corroborates the avail of such measures for the EU countries, where the receiver measure as such is not available due to lack of appropriate survey data.

3 Determinants of disagreement

3.1 The US perspective

The starting point of the present analysis and the complement to the more detailed results obtained for a wider set of EU countries in Badarinza and Gross (2009) is the very pronounced negative relationship that Figure 5 suggests for disagreement and the measure of news intensity. This empirical regularity is in line with the predictions from benchmark theoretical models of expectation formation that draw on either informational frictions or some general form of Bayesian learning.

However, the univariate relation is likely to be subject to an omitted variable bias: News co-move positively with the level of the inflation rate (the common sample correlations are .59 and .63 respectively for public and survey-based news) and thus a fall in disagreement may not be causally driven by a more intense information flow but be explained by the fact that agents either put more effort into updating their information sets (the Google search frequency correlates with inflation by a factor of .65) or that news media agencies have a reporting bias towards high (or rising) inflation. Also, care should be taken when drawing conclusions as to the survey-based news measure since it is reflective of information/news that agents have purposely chosen to be exposed to. This decision of information consumption and, moreover, the decision to choose from a long list of 80 items one of the 6 relating to price inflation, may thus not be exogenous to the individual expectation formation mechanism.

In order to properly control for these effects, we include in all regressions the inflation level, the square of the inflation level and the square of inflation in first differences as a proxy for short-term inflation volatility. While there is no clear theoretical implication for the sign of the coefficient related to the inflation level, short-term volatility shall unambiguously increase forecast heterogeneity according to all standard models of expectation formation.

We let π_t denote the inflation rate and estimate a set of models that are all nested in the following, general specification:

$$\sigma_t^i = \beta_0 + \beta_1 \sigma_{t-1}^i + \beta_2 f_t^j + \beta_3 \pi_t + \beta_4 \pi_t^2 + \beta_5 (\pi_t - \pi_{t-1})^2 + \varepsilon_t^{i,j}, \quad \varepsilon_t^{i,j} \sim N(0, \xi^{i,j})$$

with $i = \{c, q\}$ indicating categorical (c) and quantitative (q) disagreement and $j = \{s, p\}$ referring to the survey-based (s) and the public news (p) measures of information flow.

Table 1 shows a first set of regression results. Full-sample period estimates suggest that the effect of inflation-related survey-based news is negative and significant at the 1% level respectively based on both quantitative and categorical disagreement.

For the sub-period from 2000 until the end of the sample, this effect remains negative and significant only in the case of the categorical disagreement measure. As concerns magnitudes, the effect is measured to be more pronounced compared to the full-sample estimate. Figure 6 confirms that the coefficient estimates from our benchmark models are rather stable over time.

As Mankiw *et al.* (2003), we document a positive relationship between quantitative disagreement and the inflation level, with the corresponding p -value being smaller than 1%. For the sub-sample from 2000 onwards, however, the effect becomes negative (and significant and the 1% level).

Unlike for the effect of the level of the inflation rate, the effect of short-term inflation volatility appears robust throughout the sample period. The magnitudes of the estimated normalized coefficients are comparable across different model setups: the estimates suggest a significant positive effect, irrespective of the choice of the disagreement measure and for the full- and sub-sample periods. As mentioned above, we interpret this as strong evidence in favor of the hypothesis that higher fundamental uncertainty generates higher disagreement, as agents either form expectations based on outdated information sets, use different forecasting models, or have different product groups in mind when thinking about the likely evolution of the price level in the near future.

The magnitude of the effects that variation in the explanatory variables induce for disagreement can be directly compared because the coefficients have been normalized: a 1 standard deviation (STD) upward move in received news about inflation result in a .47 STD fall in categorical disagreement for the sub-sample from 2000-2011. The effect of a 1 STD rise in inflation volatility induces c.p. a .38 STD increase in categorical disagreement.

The marginal contribution of the news variable to the fraction of variation in categorical disagreement explained by the model (differences in respective R -square measures) equals +16 and +20 percentage points for the 1978-2011 and 2000-2011 sample periods, respectively.

In Table 2 we report results for regression models that include, respectively, quantitative and categorical disagreement as the dependent variable, while we now take the perspective of the sender of information and let the public news measure enter as an explanatory variable.

The results suggest that the variation in information flow has a negative effect on quantitative disagreement, albeit not being statistically significant at conventional confidence levels. Since quantitative disagreement appears to be quite persistent (with an estimated AR(1) coefficient equal 0.8), with its corresponding p -value and the increase in the model's R -square from 26% to 62% when including the lag, the effect of news cannot be separately identified.

We arrive at similar conclusions when looking at the categorical disagreement measure where model estimates suggest that the information flow exerts a negative effect on disagreement levels irrespective of whether or not a lag was included in the model. The coefficient on inflation has a negative sign, supporting the interpretation that as the inflation level moves in and out of a certain ordinal category, agents' expectations are also driven more or less in the same direction, without this actually saying much about true behaviorally relevant changes in the heterogeneity of beliefs.

The estimation results overall confirm the view that changes in information flow have the potential to densify expectations, i.e. to generate more agreement

among consumers about the likely future path of inflation. To further seek to robustify the results, Table 3 shows model estimates based on first differences of both the dependent and independent model variables. The findings discussed above remain robust.

Figure 7 shows a set of scatter plots that aim to assess how robust results are to the observations frequency.⁵ We use annual averages of all variables and compute residuals from regressions of the disagreement level and the information flow measures, employing as well the complete set of inflation control variables. We scatter the residuals from these first-stage regressions against news intensity and draw the corresponding regression lines. The picture that emerges confirms the results based on the models based on monthly frequency. Quantitative disagreement increases unconditionally along with rising news intensity, an effect which gets reversed once one controls for aggregate conditions. Second, when looking at categorical disagreement, the strong unconditional negative effects remain visible based on the annual frequency and are robust to controlling for aggregate conditions.

3.2 The European perspective

Figure 8 shows how the two disagreement measures evolve over time in a panel of seven EU countries. In particular in Germany, a well-known fact is that around the introduction of the Euro currency, the public perception of the inflation rate surged without official statistics actually confirming such perception. Brachinger (2006) explains this as being the result of weights that consumers have in mind when forming their perception having moved asymmetrically: "In the year 2001 something structurally fundamental happened regarding the change in the prices of consumer goods: above-average price increases were applied precisely to those goods characterized by an above-average purchasing frequency. The event thus features a sudden positive correlation between price change and frequency of purchase, something that had not happened previously and that shows an astounding

⁵We are grateful to Chris Carroll for noticing this point and suggesting to conduct this kind of robustness check.

persistence to the end of 2005 (p. 29).” Consistent with this view, we document a significant increase in agreement with respect to perceived inflation, indicating that the misperception that Brachinger points to was pervasive across the population. However, unlike the level perception of inflation rates, the cross-sectional dispersion in perceptions quickly rebounded, as the inflation fears were proven unfounded.

As concerns inflation expectations, a similar effect can be documented for about one year in advance, which in our view serves to describe the overall phenomenon quite clearly: having to replace the Deutsche Mark with the Euro in common trades was expected to be a painful event and when it actually happened, the population misperceived the limited signals it obtained from a subset of products as confirming the previously held expectations. A similar pattern could be observed in the Netherlands and Italy; it seemed to be absent in Spain and France.

In general, we observe a considerable degree of synchronization between the heterogeneity in perceptions and expectations. The only exception to this is Sweden, where agents persistently agree about the present inflation rate, but disagree markedly with respect to 1-year ahead expectations. We attribute this to inflation in Sweden having remained rather low and stable, at least before the financial crisis.

In Figure 9, we plot inflation rates for the set of EU countries along with the information flow measures for respective countries (public news measure). While inflation rates were contained and stable across the EU countries up until 2007, the first inflationary pressures started to materialize thereafter. Behind the scene, the low interest rate policy of the European Central Bank and the Federal Reserve, as well as the sharp increases of commodity prices caused by a booming global demand were forcefully driving the inflation rates up in all the seven countries considered here. The period was also characterized by the media’s extensive reporting about inflation, with our public news measure attaining local maxima in late 2007 in almost all countries. As the inflation rate rapidly plummeted afterwards in the course of the financial crisis, also the inflation-related information

flow fell sharply, a phenomenon that we regard as crucial for the discussion of the structural models of expectation formation in Section 4 of the paper.

Turning to the relationship between disagreement and news, we observe strong negative correlations throughout the sample period, which can be attributed to a large extent to the effect of the inflation term itself. In order to isolate the effect of news properly, we first regress the respective measures of disagreement on a set of explanatory variables including the inflation rate, squared inflation and the squared first-difference of inflation - our preferred measure of volatility or inflation uncertainty. Table 4 shows the country-individual model estimates. We plot the resulting residuals from the first-stage regressions along with the respective measures of public information flow in Figure 10. The negative relationship between news and disagreement is preserved for most countries after controlling for inflation and inflation volatility.

In order to further substantiate that point, we estimate a set of regression models similar to the ones for the US now in a dynamic panel format:

$$\sigma_{t,k}^i = \beta_0 + \beta_1 \sigma_{t-1,k}^i + \beta_2 f_{t,k}^p + \beta_3 \pi_{t,k} + \beta_4 \pi_{t,k}^2 + \beta_5 (\pi_{t,k} - \pi_{t-1,k})^2 + \bar{\beta}_k + \varepsilon_{t,k}^i$$

$$\varepsilon_{t,k}^i \sim N(0, \xi^i)$$

with $i = \{e, q\}$ indicating disagreement in expectations (e) and perceptions (p), $k = 1, \dots, 7$ representing the country and $\bar{\beta}_k$ being a country-specific fixed-effects term.

Results are shown along with the individual country equation estimates in Table 4. A more intense information flow induces disagreement to fall, an effect that is statistically significant at a 1% confidence level and holding both with respect to disagreement in perceptions and expectations. Notably, the effect of inflation uncertainty is also preserved: when inflation volatility (uncertainty) increases, agents tend to disagree more.

For Italy, the Netherlands, Sweden and the UK the main effects are robust and similar in magnitude to what we have observed for the US, while in Germany, Spain and France the respective coefficients are statistically insignificant.

We attribute this to the rather short time-series, as well as to the rather calm macroeconomic environment prevailing over a large portion of the sample period. Also, it remains unclear whether the public news measure really captures the relevant inflation-related information flow. There is no micro survey (Michigan survey-type) data available that would help judge whether public news indeed reached the German or Spanish consumers.

4 Models of information diffusion

4.1 Theoretical framework

This section aims to provide a comparative overview of what different models of expectation formation imply regarding the level of inflation expectations and the level of disagreement. The economy is assumed to consist of an infinity of agents which form expectations about monthly values of aggregate variables subject to their individual (possibly cross-sectionally heterogeneous) information sets. In each period, there is thus a continuum of individual forecasts.

We assume that a set of macroeconomic aggregates evolve according to a VAR(12) law of motion:

$$x_t = \Phi_1 x_{t-1} + \dots + \Phi_{12} x_{t-12} + \varepsilon_t \text{ with } \varepsilon \sim N(0, \sigma), \quad (3)$$

with x_t being a vector of observable (de-meanned) variables:

$$x_t \equiv \begin{bmatrix} \pi_t \\ r_t \\ y_t \end{bmatrix},$$

with π_t being the inflation rate, r_t the Federal Funds rate and y_t the economy-wide industrial production.

By appropriately stacking variables and lags, we represent the law of motion in VAR(1) format:

$$X_t = AX_{t-1} + B\epsilon_t, \text{ where } X_t \equiv \begin{bmatrix} x_t \\ x_{t-1} \\ \vdots \\ x_{t-11} \end{bmatrix} \text{ and } \epsilon_t \equiv \begin{bmatrix} \epsilon_t \\ \epsilon_{t-1} \\ \vdots \\ \epsilon_{t-11} \end{bmatrix} \quad (4)$$

At this point, the question arises as to how agents process the information in order to come up with point and categorical forecasts of the inflation rate twelve months ahead.

We consider four different expectation formation schemes:

- **Rational expectations** The optimal time t forecast, conditional on the full set of observable variables and their history, corresponds to applying the expectation operator onto the law of motion and conditioning on time t information. More precisely, we have the matrix A_t being estimated at each point in time, consistent with the idea that households learn over time the nature of the true data generating process. If all agents share the same time t information, their forecasts are all identical and the aggregate average expectation is:

$$E_t X_{t+12} = A_t^{12} X_t \quad (5)$$

- **Sticky information.** This version of the model has at its heart the information diffusion mechanism for which Reis (2006) derives explicit micro-foundations. Agents are assumed to update their information sets infrequently, with only a fraction δ_t having perfect knowledge of the whole ob-

servable vector x_t and a fraction $(1 - \delta_t)$ building expectations conditional on their $t - 1$ individual information set (i.e. not necessarily updated at $t - 1$, but at some point in the past). On aggregate, the evolution of the average expectation is given by:

$$\begin{aligned}
E_t^{SI} X_{t+12} &= \delta_t E_t X_{t+12} \\
&\quad + (1 - \delta_t) \delta_{t-1} E_{t-1} X_{t+12} \\
&\quad + (1 - \delta_t) (1 - \delta_{t-1}) \delta_{t-2} E_{t-2} X_{t+12} \\
&\quad \dots \\
&= [\delta_t \quad (1 - \delta_t) \delta_{t-1} \quad (1 - \delta_t) (1 - \delta_{t-1}) \delta_{t-2} \dots] \begin{bmatrix} A_t^{12} X_t \\ A_{t-1}^{13} X_{t-1} \\ A_{t-2}^{14} X_{t-2} \\ \vdots \end{bmatrix}
\end{aligned}$$

- **Sticky expectations.** In this version of the model, agents again update infrequently, with only a fraction δ_t forming expectations rationally based on time t information; we assume, however, that individuals amounting to a fraction of $(1 - \delta_t)$ deviate from rationality in the sense that they stick to their period $t - 1$ forecast, such that - e.g. as an outcome of rational inattention - they either do not update the information set or build a new forecast based on outdated information. One can think of this framework as implying that agents, when they update the information set form a certain expectation about one-year ahead inflation to which they then stick up until the infinite future, unless at some point they update again. There is no explicit micro evidence for this type of behaviour in the literature and indeed we consider it a rather extreme case of deviation from rationality (especially the fact that agents are assumed to neglect the existence of a term structure of inflation expectations). We think, however, that the sticky expectation scheme may be an interesting alternative description of inertial behavior and thus want to confront also this model scheme with

data. At the aggregate level, expectations then evolve as follows:

$$E_t^{SE} X_{t+12} = \delta_t A_t^{12} X_t + (1 - \delta_t) E_{t-1}^{SE} X_{t+11} \quad (6)$$

- **Epidemiological diffusion.** Finally, we consider the model by Carroll (2003) which once more implies a rather significant departure from the rational expectation hypothesis: no single consumer actually observes the set x_t of macroeconomic aggregates, but a fraction δ_t of them have access to public media through various channels of information transmission so as to receive a forecast that has been sent by professional forecasters; the remaining fraction $(1 - \delta)$ is inert in the sense that their time t forecast is derived based on the published professional forecast at some point in the past. If we denote by E^{prof} the published professional forecast, the average aggregate expectations evolve according to the following process:

$$E_t^{EPI} X_{t+12} = \delta_t E_t^{prof} X_{t+12} + (1 - \delta_t) E_{t-1}^{EPI} X_{t+11} \quad (7)$$

As can be seen above, we intend to let the share of the population that updates its beliefs to be *time-varying*. In order to do so, we employ the measures of news intensity proposed in the previous section.

The question then arises as to how one can quantify, identify respectively, the δ , that is, how some news intensity measure can be mapped into δ_t . Since there is no previous reference in the empirical literature on how this mapping could be accomplished, as a first step we take the survey-based news measure at face value and assume it be equal to δ_t . The survey-based news variable measures the fraction of respondents to the Michigan Survey which state that they had received news about inflation during the relevant period. We consider the measure a reasonable first approximation of δ_t .

The evolution of the raw survey-based news measure in Figure 3 suggests how δ may have been varying over the sample period: During the late 1970s and early 1980s, the updating share of the population has risen to about 30% per month; during the 1990s it was more or less constant at about 6% and

during the latter part of the 2000s decade it has approached about 50%. We also consider an alternative to this strict identification of δ via the survey-based news measure because it may well be subject to measurement problems. Specifically, in the simulation rounds labeled with "adjusted", we rescale the news measure by multiplying with a factor equal to the inverse of the maximum observed value, such that the new maximum of the information flow (updating) is normalized to one.

4.2 Expected inflation

As Mankiw *et al.* (2003) have shown, the sticky information model is successful in matching the observed time path of consumer inflation expectations. They find an optimal (constant) stickiness parameter of $\delta = 10\%$. Our experiments confirm these findings for a sample which covers also the entire post-2001 decade.

We find that it is not only in terms of the dynamics that the sticky information model performs better compared to a rational expectations benchmark, but also with respect to the level of inflation expectations. The upper panel of Table 5 indicates that the Root-Mean-Squared-Error (RMSE) of predicting the data with the model decreases by around 2 percentage points, with the sticky information model performing better than its counterparts.

Allowing for time variation in the updating share by relating it to the information flow enhances the models' ability to explain variation in the data significantly. First, during the years of high inflation, accounting for the fact that more than 30% of agents update their information, improves the model's ability to explain the rather quick adjustments observed in the data. Second, during the low-inflation years ensuing after the mid-1980s, accounting for information flow seems to be important for understanding why, after increases in the level of inflation, to which agents seem to adjust quite swiftly, there follows a period of inertia during which expectations remain at high levels (see Figure 11). Especially in the wake of high news intensity, expectations seem to be sticky which can be well explained by the fact that the share of agents who update is falling

sharply.

4.3 Disagreement

We now turn to analyzing the model-implied measures of disagreement. Disagreement is measured as the standard deviation of the cross-sectional distribution of forecasts, such that for some period t we have different cohorts, each of them weighted by the corresponding updating weight. Again, as shown by Mankiw *et al.* (2003), the sticky information model performs well in matching the observed dynamics in quantitative disagreement (i.e. the cross-sectional standard deviation of point forecasts), with an in-sample correlation coefficient of around 40%.

However, as Mankiw *et al.* (2003) also document, the model heavily underestimates the scale of disagreement (see top panel of Figure 12). This arises because the only source of belief heterogeneity in the model is the time variation of inflation while in reality there are a plethora of other sources of disagreement, either of a behavioral nature or resulting from the imperfect information that agents have access to, as well as from differences in model and forecasting technologies. Indeed, when we allow the forecasting technique to differ in each period, i.e. we let the matrix A be time-varying, the amount of disagreement rises by around 2 percentage points (see top panel of Figure 13), matching thereby quite closely the disagreement levels observed towards the end of sample, as well as to a somewhat lower degree the ones observed during the 1980s. Accounting for news improves the model fit slightly, but still we see no reason to believe that the heterogeneity of point forecasts was driven mainly by the infrequent updating of information sets and would rather adhere to a view pointing to observation or forecast noise.

A different picture arises when analyzing categorical disagreement instead (see bottom panel of Figure 12 and bottom panel of Table 5). First, all three models appear to perform better at accounting for time-variability in categorical than in quantitative disagreement throughout the sample period. Second, the effect of letting δ be time-varying is very pronounced and it improves the models' fit

substantially. Third, it is the epidemiological model which seems to be preferred by the data, especially during the rather calm period of the 1990s and during the more recent financial crisis period.

We thus consider a modified version of the epidemiological hypothesis relevant which would rest on the assumption that agents are not necessarily aware of the point forecasts of professionals, and instead develop a general feel for the macroeconomic environment, with fluctuations in this feel generating disagreement. Figure 14 further points to such aspect of the models. It is apparent that especially during the crisis years at the end of the 2000s, when also the inflation-related news intensity is more volatile than in previous periods, the disagreement dynamics are much better matched when accounting for the time variation in agents' updating frequencies. These very pronounced effects are to be seen as a model-consistent counterpart to the more stylized statistical results reported in the previous sections.

5 Conclusions

Understanding the determinants of the cross-sectional dispersion in expectations is important for both monetary and fiscal policy makers since the effectiveness of inflation targeting on one side and the impact of expansionary policies on the other critically depend on the way agents expect the economy to evolve in the future. Also, there are strong theoretical reasons to believe that there is a connection between the heterogeneity of beliefs and issues concerning asset pricing, stock market trading and financial stability. The empirical results reported in this paper contribute to this literature by substantiating the hypothesis that agents' beliefs about future inflation are converging in the presence of a more intense flow of public information. Our findings are in line with theoretical predictions by Carroll (2003) and Mankiw and Reis (2002) as well as with the literature on the social value of public information, as exemplified by Morris and Shin (2002). They also serve to provide an answer to the apparent inconsistency between the theory and the data as pointed to in Mankiw *et al.* (2003).

The two key elements of our analysis are the measurement of the inflation-related information flow and the quantification of cross-sectional disagreement in inflation expectations. We approach these questions from different directions, provide a comparative overview of standard methods and also propose a set of alternative measures. More research is due, however, in order to be able to draw firm conclusions about the very delicate interaction between the available information sets, the formation of beliefs and the degree to which these determine policy-relevant actions. For example, one unresolved question is whether the cross-sectional standard deviation or inter-quartile range of quantitative survey answers really capture a relevant sort of disagreement or simply reflect imperfect information processing on behalf of the agents and thus do not materialize in heterogeneous market behavior. Also, the distinction between the volume of public information and its content has recently been shown to be of importance for determining beliefs and behavior, with a gamut of psychological, social and cultural factors possibly also influencing individual decision making and market outcomes.

References

- BADARINZA, C. and GROSS, M. (2009). Inflation perceptions and expectations in the euro area - the role of news. *ECB Working Paper No. 1088*.
- BRACHINGER, H. W. (2006). *Euro or Teuro?: The euro-induced perceived inflation in Germany*. DQE Working Papers 5, Department of Quantitative Economics, University of Freiburg/Fribourg Switzerland.
- CARROLL, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics*, **118** (1), 269–298.
- DOVERN, J., FRITSCHÉ, U. and SLACALEK, J. (2009). *Disagreement among forecasters in G7 countries*. Working Paper Series 1082, European Central Bank.
- EHRMANN, M., EIJJFINGER, S. and FRATZSCHER, M. (2010). *The role of central bank transparency for guiding private sector forecasts*. Working Paper Series 1146, European Central Bank.
- EIFE, T. A. and COOMBS, W. T. (2007). *Coping with people’s inflation perceptions during a currency changeover*. Working Papers 0458, University of Heidelberg, Department of Economics.
- KANDEL, E. and ZILBERFARB, B.-Z. (1999). Differential interpretation of information in inflation forecasts. *The Review of Economics and Statistics*, **81** (2), 217–226.
- LACY, M. (2006). An explained variation measure for ordinal response models with comparisons to other ordinal R^2 measures. *Sociological Methods and Research*, **32**, 469–520.
- MAAG, T. and LAMLA, M. J. (2009). *The role of media for inflation forecast disagreement of households and professionals*. KOF Working papers 09-223, KOF Swiss Economic Institute, ETH Zurich.

- MANKIW, N. and REIS, R. (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics*, **117** (4), 1295–1328.
- MANKIW, N. G., REIS, R. and WOLFERS, J. (2003). Disagreement about inflation expectations. In *NBER Macroeconomics Annual 2003, Volume 18*, NBER Chapters, National Bureau of Economic Research, Inc, pp. 209–270.
- MCCOMBS, E. and SHAW, D. (1972). The agenda-setting function of mass media. *Public Opinion Quarterly*, **36** (2), 176–187.
- MORRIS, S. and SHIN, H. (2002). Social value of public information. *American Economic Review*, **92** (5), 1521–1534.
- REIS, R. (2006). Inattentive consumers. *Journal of Monetary Economics*, **53** (8), 1761–1800.
- ROBERTS, J. M. (1997). Is inflation sticky? *Journal of Monetary Economics*, **39** (2), 173–196.
- VELDKAMP, L. (2006). Media frenzies in markets for financial information. *American Economic Review*, **96** (3), 577–601.

Appendix

Table 1
Disagreement and survey news in the US

Categorical disagreement										
	Full sample				2000-2011					
Lagged				0.870	(0.00)		0.761	(0.00)		
Survey news			-0.471	(0.00)	-0.136	(0.08)	-0.647	(0.00)	-0.469	(0.01)
Inflation	-1.129	(0.00)	-0.781	(0.01)	-0.142	(0.40)	-0.618	(0.01)	-0.354	(0.05)
Inflation ²	0.760	(0.00)	0.642	(0.01)	0.158	(0.35)	0.415	(0.09)	0.419	(0.11)
(Δ Inflation) ²	0.112	(0.00)	0.189	(0.00)	0.258	(0.00)	0.206	(0.00)	0.384	(0.00)
obs.	401		401		401		138		138	
R^2	0.24		0.40		0.84		0.65		0.85	
Quantitative disagreement										
	Full sample				2000-2011					
Lagged				0.803	(0.00)		0.689	(0.00)		
Survey news			-0.180	(0.00)	-0.045	(0.31)	0.169	(0.45)	0.145	(0.38)
Inflation	0.831	(0.00)	0.964	(0.00)	0.544	(0.00)	-0.848	(0.00)	-0.301	(0.08)
Inflation ²	-0.082	(0.71)	-0.127	(0.59)	-0.131	(0.36)	0.793	(0.00)	0.303	(0.23)
(Δ Inflation) ²	-0.001	(0.96)	0.028	(0.19)	0.043	(0.09)	0.159	(0.00)	0.115	(0.11)
obs.	401		401		401		138		138	
R^2	0.57		0.60		0.86		0.27		0.62	

Note: The table reports coefficient estimates that have been normalized by multiplying coefficients with the standard deviation of the regressor and dividing by the standard deviation of the level of the dependent variable. Whenever the regression contains the lagged dependent variable on the right-hand side, normalization is done with the short-run standard deviation of the dependent variable. P -values derived from heteroskedasticity and autocorrelation robust standard errors (Newey-West) are reported in parentheses.

Table 2
Disagreement and public news in the US

	Categorical disagreement		Quantitative disagreement	
Lagged		0.838 (0.00)		0.698 (0.00)
Public news	-3.027 (0.00)	-0.430 (0.62)	-0.713 (0.46)	-1.305 (0.16)
Inflation	-0.412 (0.03)	-0.173 (0.17)	-0.855 (0.00)	-0.273 (0.07)
Inflation ²	-0.004 (0.99)	0.051 (0.81)	0.991 (0.00)	0.503 (0.03)
(Δ Inflation) ²	0.138 (0.00)	0.324 (0.00)	0.192 (0.00)	0.148 (0.01)
obs.	138	138	138	138
R^2	0.60	0.83	0.26	0.62

Note: Coefficient estimates reported in the table have been normalized (see footnote to Table 1). P -values derived from heteroskedasticity and autocorrelation robust standard errors (Newey-West) are reported in parentheses.

Table 3
Regressions in first-differences

	Quantitative disagreement		Categorical disagreement		
					Public news
	Full sample	2000-2011	Full sample	2000-2011	2000-2011
Δ News	0.154 (0.00)	0.246 (0.03)	-0.251 (0.00)	-0.472 (0.00)	-0.111 (0.18)
Δ Inflation	0.051 (0.18)	0.047 (0.54)	-0.159 (0.00)	-0.163 (0.01)	-0.322 (0.00)
(Δ Inflation) ²	0.054 (0.04)	0.091 (0.17)	0.194 (0.00)	0.262 (0.00)	0.252 (0.00)
obs.	401	138	401	137	137
R^2	0.03	0.08	0.16	0.41	0.23

Note: Coefficient estimates reported in the table have been normalized (see footnote to Table 1). P -values derived from heteroskedasticity and autocorrelation robust standard errors (Newey-West) are reported in parentheses.

Table 4

Dynamic panel model estimates for selected EU countries

Disagreement in expectations										
	Lagged		Public news		Inflation		Inflation ²		$(\Delta\text{Infl.})^2$	
Germany	0.921	(0.00)	0.032	(0.39)	-0.071	(0.39)	-0.055	(0.42)	-0.050	(0.28)
Spain	0.915	(0.00)	0.014	(0.45)	0.235	(0.11)	-0.181	(0.18)	0.189	(0.01)
France	0.911	(0.00)	0.068	(0.24)	-0.111	(0.27)	0.070	(0.36)	0.076	(0.16)
Italy	0.819	(0.00)	-0.164	(0.03)	0.400	(0.08)	-0.235	(0.20)	0.119	(0.07)
Netherlands	0.782	(0.00)	-0.283	(0.00)	-0.506	(0.07)	0.334	(0.14)	-0.102	(0.10)
Sweden	0.892	(0.00)	-0.176	(0.02)	-0.598	(0.04)	0.536	(0.07)	0.030	(0.36)
UK	0.692	(0.00)	-0.213	(0.01)	1.019	(0.01)	-0.863	(0.01)	0.026	(0.38)
Full sample	0.889	(0.00)	-0.121	(0.00)	0.050	(0.30)	-0.043	(0.32)	0.060	(0.03)
Disagreement in perceptions										
	Lagged		Public news		Inflation		Inflation ²		$(\Delta\text{Infl.})^2$	
Germany	0.928	(0.00)	-0.030	(0.39)	0.213	(0.21)	-0.358	(0.10)	0.097	(0.12)
Spain	0.937	(0.00)	0.160	(0.05)	-0.108	(0.32)	-0.134	(0.26)	0.156	(0.03)
France	0.852	(0.00)	0.059	(0.27)	-0.089	(0.32)	-0.237	(0.16)	0.098	(0.12)
Italy	0.855	(0.00)	-0.231	(0.01)	0.015	(0.48)	-0.302	(0.14)	0.029	(0.36)
Netherlands	0.931	(0.00)	-0.261	(0.00)	0.076	(0.41)	-0.255	(0.21)	0.015	(0.43)
Sweden	0.815	(0.00)	0.119	(0.08)	-0.026	(0.47)	0.168	(0.31)	0.082	(0.16)
UK	0.892	(0.00)	-0.189	(0.03)	0.650	(0.08)	-0.583	(0.10)	0.134	(0.06)
Full sample	0.929	(0.00)	-0.103	(0.00)	-0.020	(0.42)	-0.086	(0.18)	0.085	(0.00)

Note: Coefficients reported in this table have been normalized. Lines to which individual country names are assigned show individual regression results. Fixed effects dynamic panel regression estimates are presented in the rows referred to as *Full sample*. *P*-values derived from heteroskedasticity and autocorrelation robust standard errors (Newey-West) are reported in parentheses.


Table 5
Assessment of model properties

	Rational forecast	Constant updating		Survey-implied updating		Survey-implied updating (adjusted)				
		SI	SE	EPI	SI	SE	EPI	SI	SE	EPI
Inflation expectations										
Correlation coef.	0.887	0.889	0.872	0.244	0.918	0.886	0.230	0.913	0.908	0.372
RMSE	0.061	0.042	0.054	0.053	0.039	0.055	0.053	0.046	0.052	0.050
Regression R^2	0.315	0.790	0.761	0.059	0.842	0.785	0.053	0.833	0.825	0.139
Quantitative disagreement										
Correlation coef.		0.365	0.336	0.512	0.312	0.281	0.519	0.089	0.096	0.485
RMSE		0.240	0.231	0.245	0.241	0.231	0.245	0.249	0.242	0.250
Regression R^2		0.133	0.113	0.262	0.097	0.079	0.269	0.008	0.009	0.235
Categorical disagreement										
Correlation coef.		0.158	0.330	0.315	0.214	0.375	0.335	0.413	0.486	0.349
RMSE		0.034	0.043	0.028	0.032	0.044	0.028	0.026	0.034	0.022
Regression R^2		0.025	0.109	0.099	0.046	0.141	0.112	0.170	0.237	0.121

Note: Summary statistics reported in this table measure the distance between model-implied and survey-based measures of inflation expectations and quantitative/qualitative disagreement. The RMSE (Root-Mean-Squared-Error) and R^2 estimates are based on is based and the residuals from a regression of model-implied upon survey-based variables. The three variants of the model are sticky information (SI), sticky expectations (SE) and epidemiological diffusion (EPI). The updating frequency used in the versions of the models reported in the right-hand column is based on the survey-implied information flow measure, normalized as described in Section 4.1.

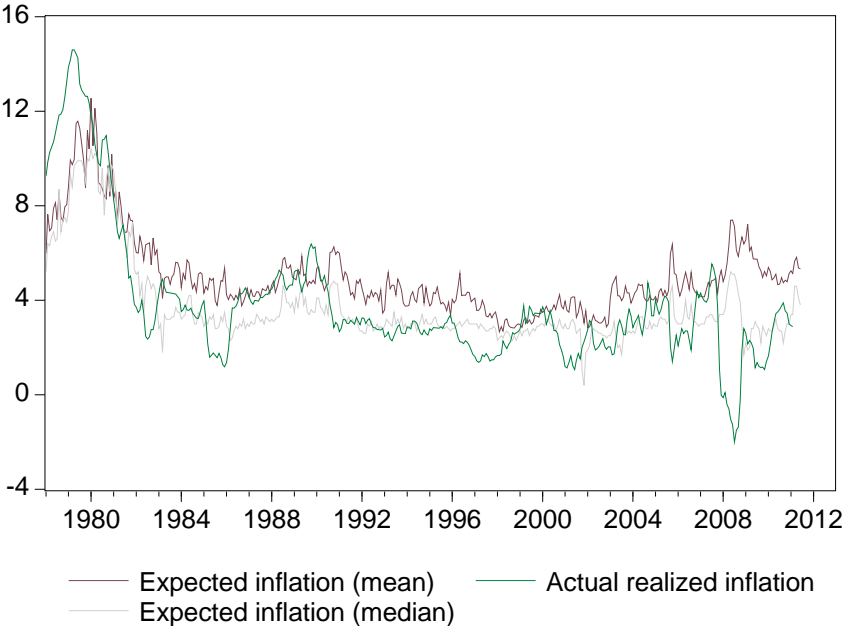
Figure 1

Disagreement in expectations in the US



Note: Details as to the construction of the disagreement measures shown in this Figure can be found in Section 2.1.


Figure 2
Inflation expectations in the US



Note: See Section 2.1 for details.

Figure 3


Public versus survey-based news intensity measures for the US



Note: Details as to the construction of the news measures shown in this Figure can be found in Section 2.2.

Figure 4



Sender versus receiver-based news intensity measures for the US



Note: The Google news series is the year-on-year rate of change based on otherwise raw search frequencies. All variables are normalized by subtracting their mean and dividing by respective 2005-2011 sample standard deviations.

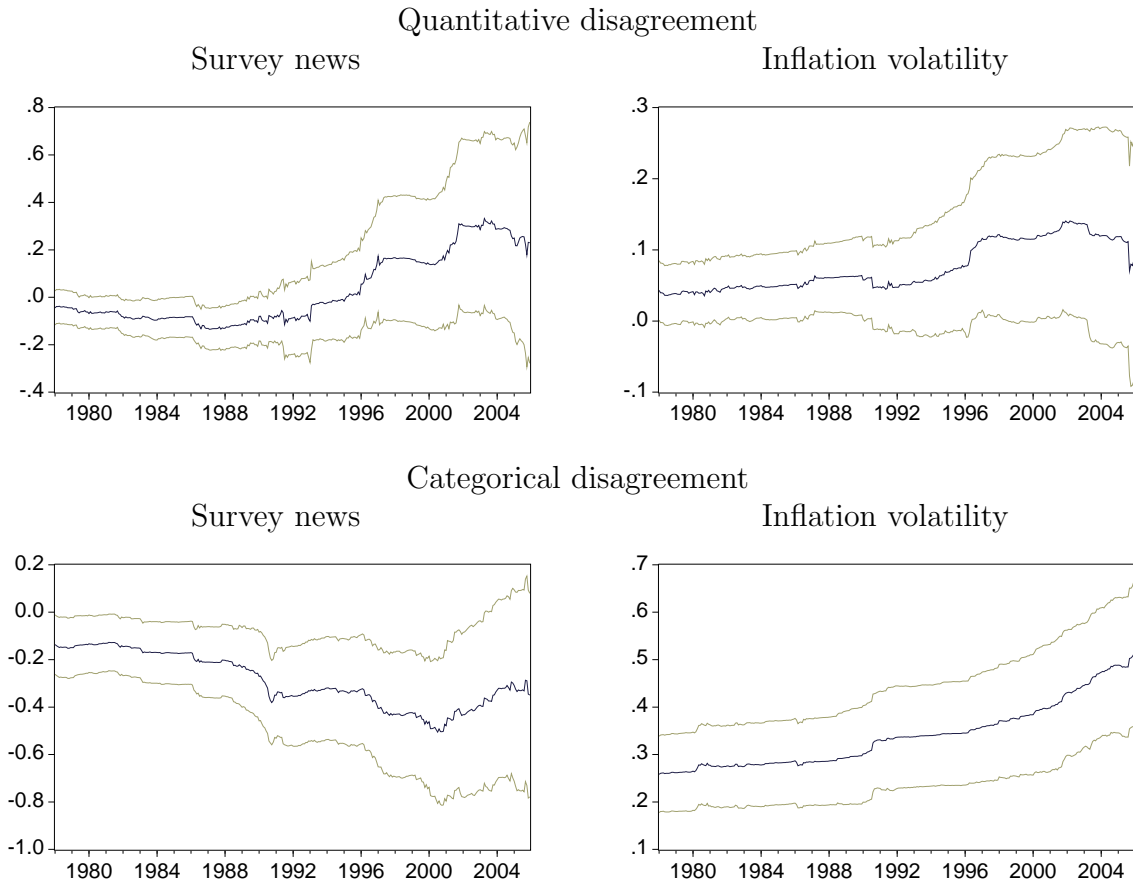
Figure 5

News intensity versus disagreement in the US



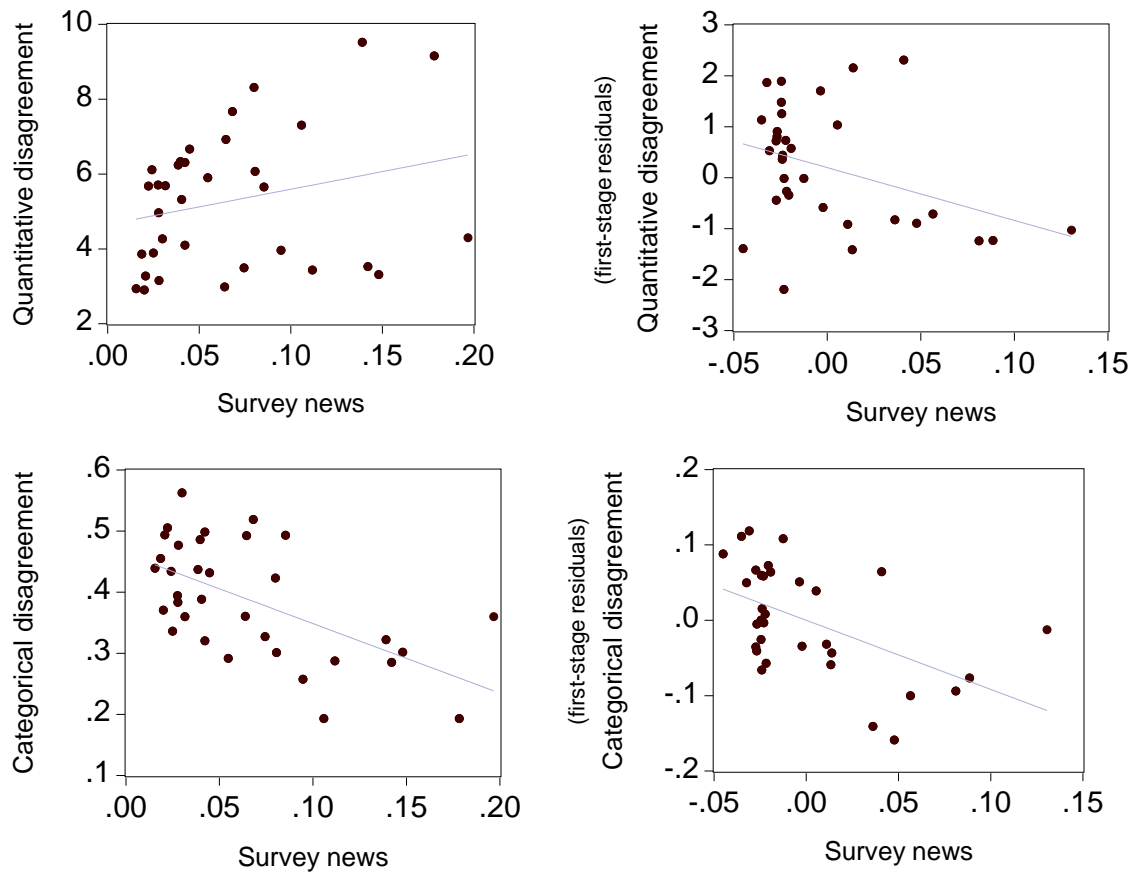
Note: All variables are normalized by subtracting their mean and dividing by respective sample standard deviations.

Figure 6
Parameter stability



Note: The blue line shows the evolution of normalized coefficients from contracting-window regressions of categorical disagreement on survey-based news, inflation controls and an AR(1) term. The first point on the line thus corresponds to a full sample regression, as in the benchmark case reported in Table 1. Gray lines are 90% confidence bands computed from HAC (Newey-West) robust standard errors.

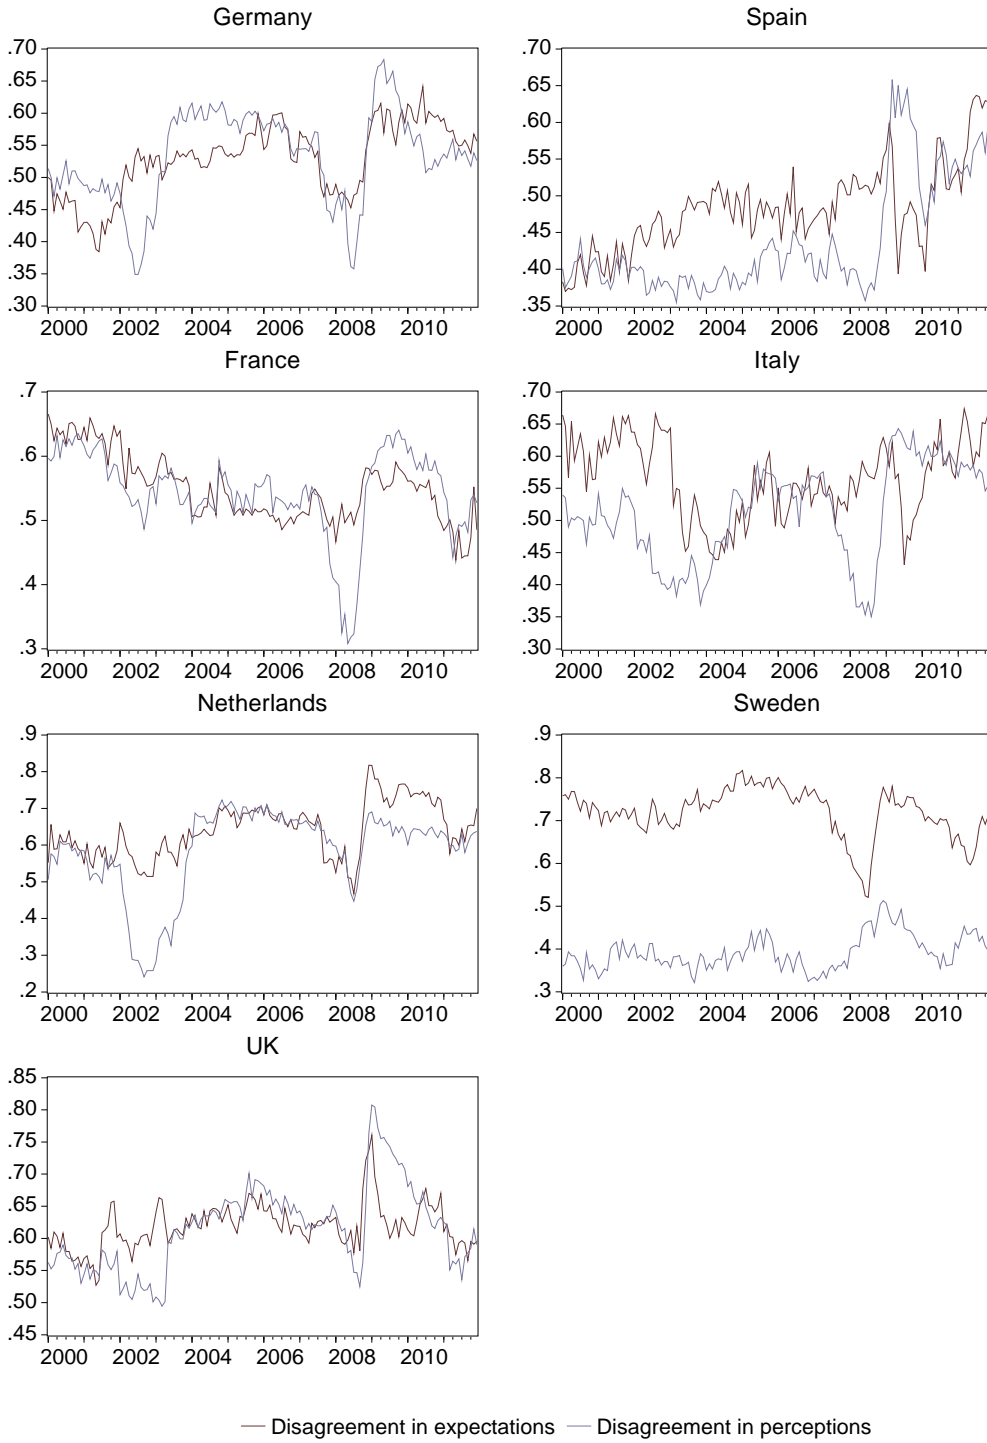
Figure 7
Robustness to observations frequency



Note: The underlying data have an annual frequency. The scatter plots on the right-hand side involve the residuals from first-stage regressions on the complete set of inflation controls. See Section 3.1 for details.

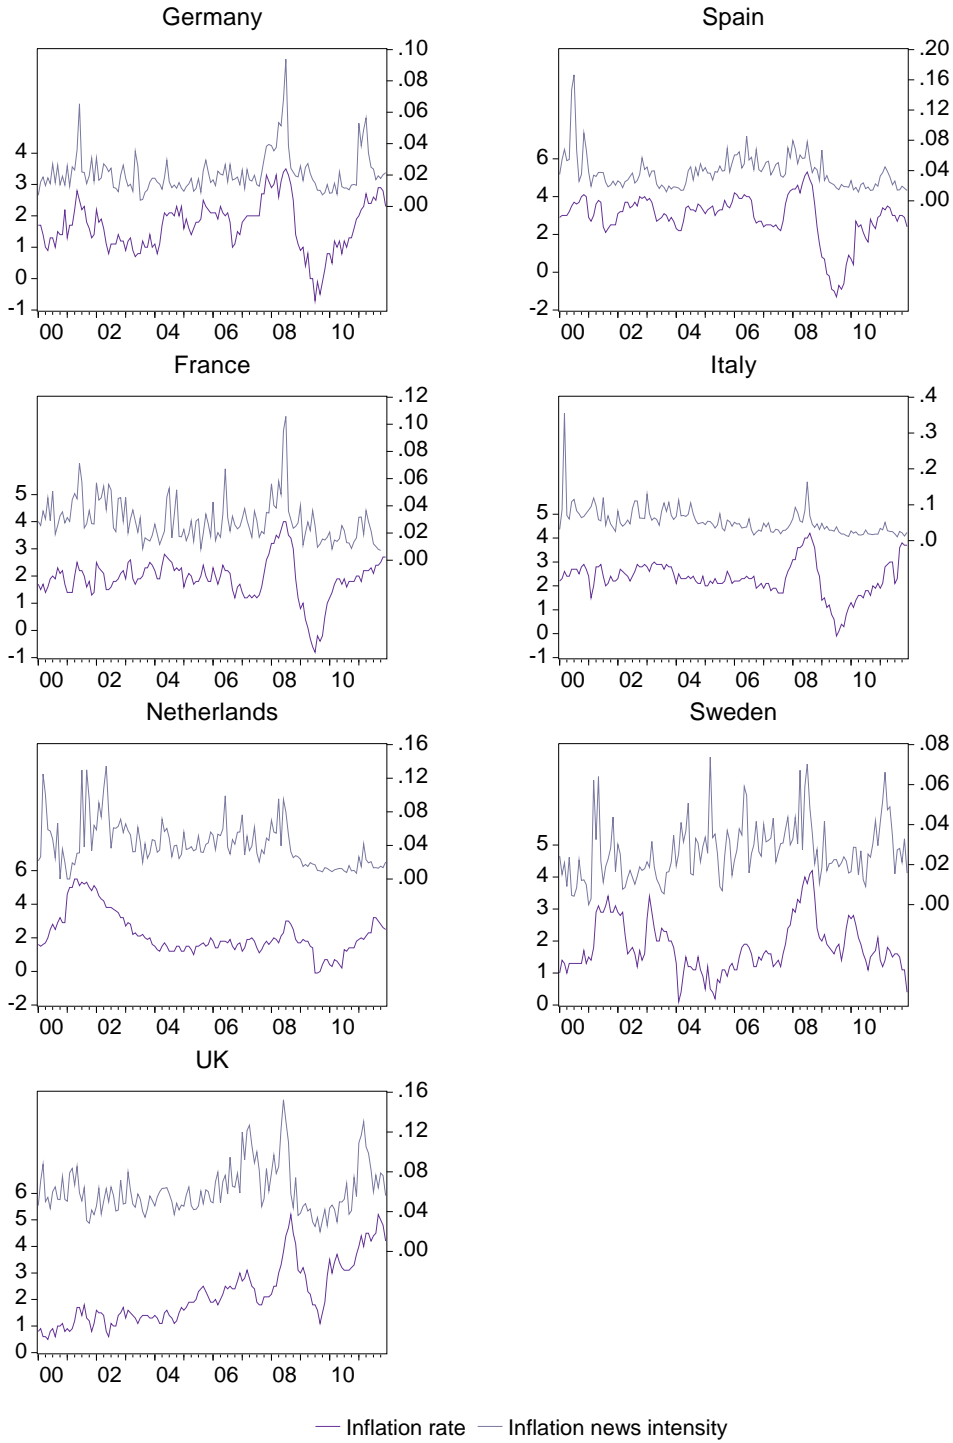
Figure 8

Disagreement in inflation perceptions and expectations for a panel of EU countries



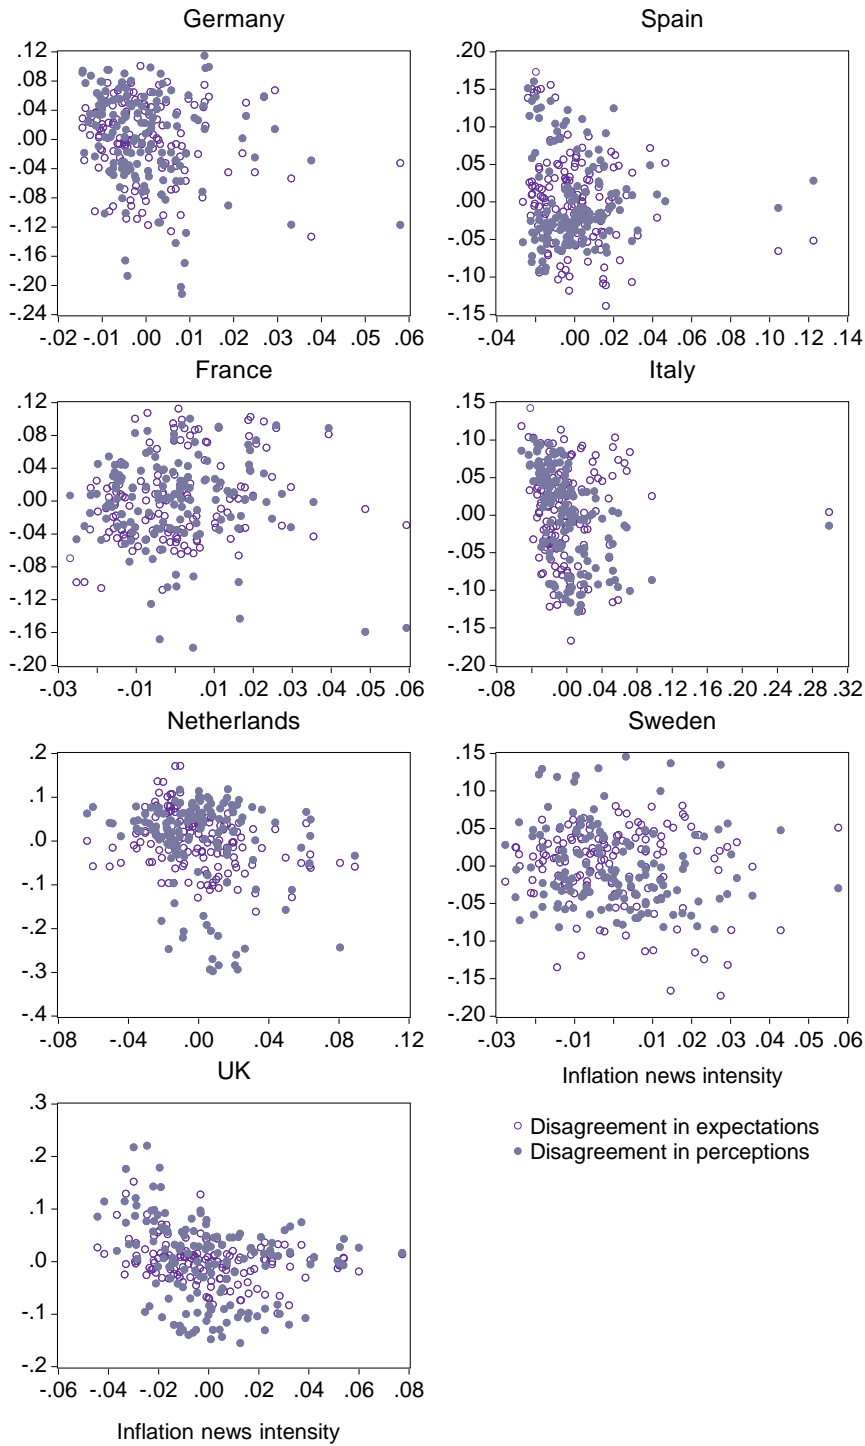
Note: See Section 3.2 for details.

Figure 9
News intensity in selected EU countries




Note: See Section 3.2 for details.

Figure 10
Effects of public news media in a panel of EU countries



Note: The scatter plots involve the residuals from first-stage regressions on the complete set of inflation controls (scattered against news intensity). See Section 3.2 for details.


Figure 11
Model-implied inflation expectation levels



Note: The figure shows model-implied inflation expectations involving time-varying updating frequencies that are based on the survey-based news measure and full-sample coefficient estimates. With respect to inflation expectation levels, the figure is representative of all other considered model variants, including the alternative learning specifications and the adjustments of the information flow series. The blue shaded bars show the raw survey-based news measure, appropriately scaled to facilitate readability of the graph. The series referred to as *Data* is the median 1-year ahead inflation expectation in the Michigan Survey of Consumers.


Figure 12

Model-implied disagreement levels (benchmark case)



Note: The survey-implied categorical disagreement level and the model-implied standard deviation of forecasts are normalized in the bottom panel by subtracting their mean and dividing by respective sample standard deviations. The model version reported in this figure assumes time-varying survey-based non-adjusted updating frequencies and constant full-sample coefficient estimates.


Figure 13
 Model-implied disagreement levels (rational learning)



Note: The survey-implied categorical disagreement level and the model-implied standard deviation of forecasts are normalized in the bottom panel by subtracting their mean and dividing by respective sample standard deviations. The model version reported in this figure assumes time-varying survey-based non-adjusted updating frequencies and time-varying coefficient estimates, according to an expanding-window sample selection scheme.

Figure 14

Effects of accounting for time variation in updating frequencies



Note: The survey-implied categorical disagreement level and the model-implied standard deviation of forecasts are normalized by subtracting their mean and dividing by respective sample standard deviations. This normalization explains the slight scaling differences between this figure and Figures 12 and 13. The model version reported in this figure assumes time-varying survey-based adjusted updating frequencies and full-sample coefficient estimates.