

Cole, Stephen J.; Milani, Fabio

Working Paper

Heterogeneity in Individual Expectations, Sentiment, and Constant-Gain Learning

CESifo Working Paper, No. 8343

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Cole, Stephen J.; Milani, Fabio (2020) : Heterogeneity in Individual Expectations, Sentiment, and Constant-Gain Learning, CESifo Working Paper, No. 8343, Center for Economic Studies and Ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/219161>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Heterogeneity in Individual Expectations, Sentiment, and Constant-Gain Learning

Stephen J. Cole, Fabio Milani

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Heterogeneity in Individual Expectations, Sentiment, and Constant-Gain Learning

Abstract

The adaptive learning approach has been fruitfully employed to model the formation of aggregate expectations at the macroeconomic level, as an alternative to rational expectations. This paper uses adaptive learning to understand, instead, the formation of expectations at the micro-level, by focusing on individual expectations and, in particular, trying to account for their heterogeneity. We exploit survey data on output and inflation expectations by individual professional forecasters. We link micro and macro by endowing forecasters with the same information set that they would have as economic agents in a benchmark New Keynesian model. Forecasters are, however, allowed to differ in the constant gain values that they use to update their beliefs. We estimate the best-fitting constant gain for each forecaster. We also extract individual measures of sentiment, defined as the degrees of excess optimism and pessimism that cannot be justified by the near-rational learning model, given the state of the economy and the updated beliefs. Our results highlight the heterogeneity in the gain coefficients adopted by forecasters, which is particularly pronounced at the beginning of the sample. The median values are consistent with those typically estimated using aggregate data, and display some moderate time variation: they occasionally jump to higher values in the 1970-80s, and stabilize in the 1990s and 2000s. Individual sentiment is persistent and heterogeneous. Differences in sentiment, however, don't simply cancel out in the aggregate: the majority of forecasters exhibit excess optimism, or excess pessimism, at the same time.

JEL-Codes: C520, D840, E320, E500, E600, E700, E710.

Keywords: individual survey forecasts, heterogeneous expectations, constant-gain learning, New Keynesian Model, sentiment shocks, waves of optimism and pessimism, evolving beliefs.

Stephen J. Cole
Department of Economics
Marquette University
P.O. Box 1881
USA – Milwaukee, WI 53201
stephen.cole@marquette.edu
<http://www.stephencoleeconomics.com>

Fabio Milani
Department of Economics
University of California, Irvine
3151 Social Science Plaza
USA – Irvine, CA 92697-5100
fmilani@uci.edu
<http://www.socsci.uci.edu/~fmilani>

1 Introduction

Since the introduction of the rational expectations hypothesis in the 1970s, a number of macroeconomists have raised questions about its empirical validity, and have offered potential replacements. The approach that has emerged as the main alternative to rational expectations is probably provided by the literature on adaptive learning (Evans and Honkapohja, 2001, Sargent, 1993, 1999). Under rational expectations, economic agents are endowed with substantial knowledge about the economy: they know the structure of the model, the values of the parameters representing preferences, technology, and policy, and the processes for the exogenous disturbances. Learning relaxes these strict informational assumptions to introduce some limitations to agents' understanding. For instance, agents within the model are no longer assumed to know the magnitudes of all economic relationships; instead, they have to learn about them based on past experiences and historical data.

Various papers have already provided evidence that adaptive learning matters at the macroeconomic level. Learning drives persistence in macroeconomic variables (Orphanides and Williams, 2005, Milani, 2007, 2017), it provides an endogenous source of the observed shifts in postwar volatility (Branch and Evans, 2007, Milani, 2014, Dave and Malik, 2017), and it amplifies business cycle fluctuations (Eusepi and Preston, 2011, Milani, 2011, 2017). Models with learning typically outperform models with rational expectations in their ability to fit macroeconomic time series (Milani, 2007, Slobodyan and Wouters, 2012).

The study of aggregate survey data on expectations has provided additional empirical evidence in favor of learning. Branch and Evans (2006) show that constant-gain learning fits median survey expectations about inflation and output better than learning with decreasing and Kalman-filter gains. Markiewicz and Pick (2014) similarly find that models based on constant gain learning provide a better fit of professional forecaster's expectations concerning a wider range of macroeconomic and financial variables. Malmendier and Nagel (2016) analyze inflation forecasts from the University of Michigan Survey of Consumers; they also show that, at the aggregate level, mean expectations are closely replicated by a constant-gain learning updating rule with a gain coefficient of similar magnitude to that estimated in macroeconomics models (e.g., Milani, 2007).

The main scope of our paper is to contribute to the literature that studies learning and the formation of *individual-level* expectations. In particular, we focus on the heterogeneity of expectations across individual forecasters. Other studies have revealed a significant degree of heterogeneity.

Mankiw, Reis, and Wolfers (2004) document the extent and properties of disagreement related to future inflation expectations. Coibion and Gorodnichenko (2012) show how disagreement in survey data responds to economic shocks. Andrade et al. (2016) use Blue Chip forecasts at short and long horizons to document new stylized facts about disagreement; these facts can be rationalized using models with imperfect information. Branch (2004) examines the role of model uncertainty in producing the heterogeneity observed in the data. Baranowski (2014) studies the type of adaptive learning rule that best fits individual forecasts of inflation, unemployment, and real GDP growth, using various combinations of regressors. In this paper, we aim to impose more structure on the heterogeneity, by examining individual-level heterogeneity through the lenses of a benchmark macroeconomic model with learning. We do so by treating individual forecasters as if they were agents in a model, and we provide them with a similar information set.

We use data on individual forecaster expectations from the Survey of Professional Forecasters (SPF). We use the same forecasts that would enter in a benchmark New Keynesian model: one-quarter-ahead forecasts for output (growth) and inflation. To minimize composition effects coming from the entry and exit of forecasters, we retain in the sample only observations for forecasters that remain in the survey for at least ten (and, as robustness, twenty) periods.

We then investigate how these individual expectations are formed. As background, we assume that the underlying economy is summarized by a canonical New Keynesian model. In a first step, we estimate the aggregate model by matching expectations in the model to the mean from our panel of forecasts, and assuming that aggregate expectations are formed under constant-gain learning. From the estimated aggregate model, we obtain the filtered structural disturbances, which are typically part of the information set for agents in the model, and that we assume to be part of the individual forecasters' information set as well.

In the second step, we then turn to the analysis of individual expectations. We assume that those expectations are formed from a perceived linear model of the economy (PLM). Our assumption is that forecasters are given the same model and the same information set that they would have as economic agents in a benchmark New Keynesian model. Therefore, they use a PLM that is equal to the Minimum State Variable (MSV) solution of the corresponding macro model under rational expectations, and they are assumed to observe the disturbances (for robustness, we will also relax this assumption, endowing agents with knowledge of the lagged endogenous variables, but not of disturbances). For each forecaster, we minimize the distance between their observed forecasts and

the expectations formed from the learning PLM. As a result, we obtain the best-fitting constant gain for each individual forecaster. The gain governs their speed of learning for the sample during which they are in the survey: it can be interpreted as their perceived probability that the variable they are forecasting will be subject to a structural break, as well as their memory of past observations.

For each forecaster, expectations are well explained by the learning model. But we denote the difference between observed expectations and the portion that is explained by the learning model as ‘excess optimism and pessimism’, or ‘sentiment’. Those optimism and pessimism terms may be serially correlated and represent an individual-level version of the aggregate sentiment analyzed in Milani (2011, 2017), which show that sentiment shocks are responsible for about half of business cycle fluctuations.

The heterogeneity in individual expectations in our environment, therefore, stems from two sources: different constant gains and different degrees of optimism/pessimism. We can evaluate how gains change across the cross-section of forecasters at each point in time, and how they change over time. Also, we can track sentiment at the individual level and evaluate whether optimism and pessimism shocks average out across forecasters or are, instead, cross-sectionally correlated.

Main Results. We document a substantial heterogeneity in the learning approach of individual forecasters. Their gain coefficients are heterogeneous: in many periods, forecasters who are largely unresponsive to new information coexist with forecasters who employ gains around 0.1 or higher. The gains vary over time: they are often higher in the 1970s and 1980s, with averages that rise to values of 0.03-0.05, and they decline in the second part of the sample, stabilizing around 0.015, and with a lower dispersion. The micro evidence is, therefore, consistent with switches in the gain as identified in Milani (2014), who also proposed time variation in the gain as a potential driver of stochastic volatility in output and inflation.

Beliefs about macroeconomic relationships estimated at the individual level also reveal substantial heterogeneity across forecasters and changes over time. On average, perceptions about the persistence of inflation increase over the sample, before reverting back later on. Forecasters also significantly revise their beliefs about the effectiveness of monetary policy: the perceived sensitivity of output to interest rates fall to values between -1.3 and -2 for most of the 1970s, and it moves upward with Volcker’s disinflation. At the end of the sample, the perceived sensitivity has been reduced to a coefficient of -0.5. Individual beliefs can also affect the dynamics of the aggregate model: impulse responses and the role played by different shocks can be very different depending

on which beliefs prevail in the population.

Moreover, we provide estimates of sentiment series at the individual level. Excesses of optimism and pessimism by single forecasters do not cancel out in the aggregate, rather they typically move in herd. The evolution of mean sentiment mirrors the series that is estimated at the aggregate level. Sentiment is persistent, and it has a volatility that is comparable to that of other structural disturbances.

Related literature. The paper aims to add to the literature on adaptive learning in macroeconomics, which has historically been more focused on the formation of expectations at the aggregate level. Evans and Honkapohja (2001) and Sargent (1993, 1999) review the foundations of the adaptive learning approach. Adaptive learning has important implications for which monetary policy strategies are desirable (as described, for example, by Orphanides and Williams, 2005, Preston, 2006, Gaspar, Smets, and Vestin, 2006, Eusepi and Preston, 2010), fiscal policy (Evans, Honkapohja, and Mitra, 2009), and for the effects of ‘forward guidance’ (Cole, 2020a,b). In addition, previous research has examined how learning behavior can explain fluctuations in the macroeconomy. Milani (2007, 2014) shows that learning is successful in capturing the persistence and volatility of macroeconomic data. Eusepi and Preston (2011) find that learning helps explain the propagation of shocks over the business cycle. Prior literature also provides evidence that learning accurately captures the formation of aggregate survey expectations (Orphanides and Williams, 2005, Branch and Evans, 2006, Markiewicz and Pick, 2014, Bräuning and van der Cruijssen, 2019).

Fewer papers delve, instead, into the the formation of *individual expectations*, which is the main objective of this paper. In particular, we provide evidence on the importance of heterogeneity at the microeconomic level by exploiting individual survey expectations. Therefore, our work is more closely connected to Branch (2004), Pfajfar and Santoro (2010), and Malmendier and Nagel (2016). Branch (2004) describes how the forecasting models used by agents are not necessarily constant, but they can shift over time. Different shares of agents may switch between models in forming expectations, generating heterogeneity. Pfajfar and Santoro (2010) study heterogeneity by examining the time series of different percentiles from the cross-sectional distribution of inflation forecasts. Malmendier and Nagel (2016) use an adaptive learning model to argue that consumers discount the past differently based on their age. We provide a link between micro and macro approaches, instead, by treating forecasters in the same way as we would treat them as agents in a benchmark New Keynesian model, with the same perceived model and the same information set.

We document substantial levels of heterogeneity, while maintaining a parsimonious model for the source of such heterogeneity: in our case, it derives from the dispersion of constant gain values across forecasters, and from different degrees of excess optimism and pessimism.

A number of papers provide theoretical foundations and interpretation for learning gain coefficients. Evans and Honkapohja (2001) examine and discuss the gain coefficient in terms of convergence of the learning model to its rational expectations counterpart. Barucci (1999) and Honkapohja and Mitra (2006) describe that gain coefficients can also be interpreted as the degree of memory forecasters attach to past observations. Berardi and Galimberti (2017) document appropriate approaches for calibrating and interpreting gain coefficients. Berardi (2019) offers a Bayesian framework for interpreting the gain coefficient as the probability of estimated parameters changing every period. Our paper adds to this literature by shedding light on realistic values for gains at the individual level, and it reveals both heterogeneity in the cross-section of forecasters and time-variation over the sample.

Our results regarding the importance of sentiment in individual expectations also provide an important rationale at the microeconomic level for the type of aggregate sentiment that has been recently introduced in a variety of macroeconomic models. Milani (2011, 2017) utilizes adaptive learning and aggregate survey expectations from the SPF and finds that sentiment shocks explain a significant portion of business cycle fluctuations.¹ Angeletos, Collard, and Dellas (2018) also describe the importance of sentiment (or “confidence”) shocks for explaining the business cycle. When firms are faced with a signal extraction problem for their goods, Benhabib et al. (2015) show that sentiment can lead to equilibria away from a standard rational expectations solution. Moreover, our results add to the previous research by providing compelling evidence of sentiment shocks as an additional source of heterogeneity at the microeconomic level.

In the next Section, we present the individual expectations data from the SPF. Section 3 lays out the New Keynesian model, which is assumed as a benchmark model for the underlying economy. The model is estimated using Bayesian methods and aggregate expectations matched to the mean survey expectations. Section 4 investigates the expectation formation at the individual forecaster level to extract best-fitting gains and micro-level optimism and pessimism. The results are discussed in Section 5, and their sensitivity is assessed in Section 6. Section 7 presents the conclusions.

¹The papers by Charalampidis and Milani (2020), Chatterjee and Milani (2019), and Cole and Milani (2019) also highlight the empirical importance of sentiment in different frameworks.

2 Individual Survey Expectations Data

We use individual expectations data from the SPF, hosted by the Federal Reserve Bank of Philadelphia. We focus on forecasts about future real GDP growth and the future inflation rate calculated from the GDP Implicit Price Deflator. The series of reference are ‘*RGDP*’ and ‘*PGDP*’, and specifically we use ‘*RGDP2*’, ‘*RGDP3*’, ‘*PGDP2*’, and ‘*PGDP3*’: they refer to expectations formed by forecasters at time t , while being able to observe the published values of the same variables up to $t - 1$, about the value of the variables at the end of the current period t and of the next period $t + 1$. Expectations about $t + 1$, hence, have the same horizon as those that would enter in a benchmark New Keynesian model.

Mansky (2011) has highlighted the potential composition effects that can arise due to the entry and exit of forecasters. Therefore, we keep in the panel only those forecasters that remain in the survey and submit forecasts for at least ten periods (in the robustness section, we consider a sample with those that remain twenty periods), and for which we have both output and inflation forecasts. After constructing this data set, our sample includes 204 individual forecasters that participate in the survey at different points, and for a number of periods above the threshold, between the last quarter of 1968 and the third quarter of 2016.

Given our interest in inferring the learning process and any excess optimism/pessimism in real-time, it is crucial that we try to capture the actual information set that was available to forecasters at the time the forecasts were produced. We do so by exploiting the real-time data series that the SPF provides in correspondence of each forecast (obtained through the Real-Time Data Set for Macroeconomists, also hosted by the Federal Reserve Bank of Philadelphia). Therefore, we use the corresponding real-time data for output and inflation (the series with acronym *ROUTPUTQvQd* and *PQvQd*) as our observables for the realized variables. For each series, we use the first-vintage observation.

Figure 1 shows the evolution of individual expectations over the sample, along with the implied mean and the actual realized series.² Figure 2 plots the disagreement across forecasters, measured as the cross-sectional standard deviation.

We will assume that individual expectations are formed in a way that is consistent with a typical Perceived Law of Motion (PLM) from a canonical macroeconomic model. We turn to the

²We find that difference between the mean and median are trivial (with a correlation above 0.99), once we have cleaned the sample to include only ‘long’-term participants to the survey.

presentation of the model and the expectation formation assumptions next.

3 New Keynesian Model

We assume that our individual forecasters are endowed with the same information set that economic agents would have in a benchmark New Keynesian model. Therefore, in forming their macroeconomic forecasts, they use information from past realizations of the endogenous variables (output gap, inflation, and interest rates), as well as information about structural disturbances to demand and supply (autoregressive natural-rate and cost-push shocks). The set of variables that they use in their forecasting models, therefore, corresponds to the same variables that appear in the MSV solution of the model under rational expectations.

We assume that the underlying aggregate economy is characterized by a canonical New Keynesian model (e.g., Woodford, 2003), extended to include endogenous sources of persistence as habit formation and inflation indexation. The model is summarized by the following equations:

$$\tilde{y}_t = \widehat{E}_t \tilde{y}_{t+1} - \psi(i_t - \widehat{E}_t \pi_{t+1} - r_t^n) \quad (1)$$

$$\tilde{\pi}_t = \beta \widehat{E}_t \tilde{\pi}_{t+1} + \kappa (\omega y_t + \psi^{-1} \tilde{y}_t) + u_t \quad (2)$$

$$i_t = \rho_t i_{t-1} + (1 - \rho_t) [\chi_t^\pi \pi_t + \chi_t^y y_t] + \sigma_{\varepsilon,t} \varepsilon_t \quad (3)$$

where $\tilde{y}_t \equiv y_t - \eta y_{t-1}$, $\tilde{\pi}_t \equiv \pi_t - \gamma \pi_{t-1}$, $\psi \equiv \sigma(1 - \eta)$, $\kappa \equiv (1 - \alpha\beta)(1 - \alpha)/\alpha$, and where y_t denotes the output gap, π_t denotes inflation, and i_t denotes the short-term nominal interest rate, which serves as the monetary policy instrument. The coefficient σ denotes the elasticity of intertemporal substitution, ω the inverse of the Frisch elasticity of labor supply, β the household's discount factor, α the Calvo price stickiness parameter, and ρ_t , χ_t^π , and χ_t^y are Taylor rule coefficients that denote the inertia of interest rate decisions, and the monetary policy reaction to inflation and the output gap. The degree of (external) habit formation in consumption is measured by η and the extent of indexation to past inflation in price setting by γ .

The model includes three exogenous disturbances: the demand (real natural rate) disturbance r_t^n , the supply (cost-push or price markup) disturbance u_t , which are assumed to evolve as AR(1) processes with autoregressive coefficients ρ_t^j and standard deviations σ_t^j , with $j = r, u$, and the

monetary policy shock, which, following the convention in the literature, is assumed to be *i.i.d.*:

$$r_t^n = \rho_t^r r_{t-1}^n + \sigma_t^r \varepsilon_t^r \quad (4)$$

$$u_t = \rho_t^u u_{t-1} + \sigma_t^u \varepsilon_t^u. \quad (5)$$

To improve the fit of the model to postwar data, we allow some of the coefficients to vary over time (and they are denoted with a t subscript). The Taylor rule coefficients, including the volatility of monetary policy shocks, are allowed to assume different values in the pre-1979 sample, in the non-borrowed-reserve targeting experiment years between 1979 and 1982, and in the post-1982 period:

$$\left\{ \begin{array}{ll} \left[\rho_{pre79}, \chi_{pre79}^\pi, \chi_{pre79}^y, \sigma_{\varepsilon,pre79} \right] & t \leq 1979 : 3 \\ \left[\rho_{79-82}, \chi_{79-82}^\pi, \chi_{79-82}^y, \sigma_{\varepsilon,79-82} \right] & 1979 : 4 \leq t \leq 1982 : 4 \\ \left[\rho_{post82}, \chi_{post82}^\pi, \chi_{post82}^y, \sigma_{\varepsilon,post82} \right] & t \geq 1983 : 1 \end{array} \right.$$

Moreover, substantial evidence points toward a break in the volatility of the macroeconomic shocks around 1984 (e.g., McConnell and Perez-Quiros, 2000). Hence, we allow the remaining disturbance parameters (both the autoregressive coefficients and standard deviations) to potentially switch between the pre-1984 and post-1984 samples:

$$\left\{ \begin{array}{ll} \left[\rho_{pre84}^j, \sigma_{pre84}^j \right] & t < 1984 : 1 \\ \left[\rho_{post84}^j, \sigma_{post84}^j \right] & t \geq 1984 : 1 \end{array} \right.$$

for $j = r, u$. Expectations in the model are denoted by \widehat{E}_t and they are measured by the mean of expectations from our sample of individual forecasters: $\widehat{E}_t = \int_j \widehat{E}_{t,j} dj$.

3.1 Near-Rational Expectations

Following the literature on adaptive learning in macroeconomics (Evans and Honkapohja, 2001, 2013, Sargent, 1999), we assume that agents in our model follow the ‘cognitive consistency’ principle: they don’t enjoy a knowledge advantage compared with the modeler, and they try to infer relationships among variables by analyzing historical data, as econometricians would. To produce forecasts about future variables (e.g., output or inflation), they employ a linear perceived model, estimated using standard techniques (e.g., OLS or WLS). As new information arrives every period, they update forecasts accordingly, thus continuously learning about the economy.

Therefore, aggregate expectations in the model are assumed to be formed as in Milani (2011, 2017), i.e., from the following Perceived Law of Motion (PLM):

$$Y_t = a_t + b_t Y_{t-1} + c_t \epsilon_t + \nu_t \quad (6)$$

where $Y_t = [y_t, \pi_t, i_t]'$, $\epsilon_t = [r_t^n, u_t]'$, and a_t , b_t , and c_t , are vectors and matrices of coefficients of appropriate dimensions. The term ν_t denotes an econometric error term. As common in the adaptive learning literature, economic agents are assumed to use a correctly-specified model to generate their forecasts: the model corresponds to the Minimum State Variable (MSV) solution of the system under rational expectations. Agents, hence, use the correct set of endogenous variables in their perceived model, for which they observe data up to $t - 1$, and they are assumed to observe the contemporaneous disturbances. The model contrasts with rational expectations, since agents are assumed to lack knowledge about the reduced-form coefficients in the PLM: therefore, they do not know the magnitude of the relationships among variables. For example, they do not know how sensitive output and inflation are to interest rate changes or to demand and supply shocks, or the persistence of output and inflation, or the slope of the Phillips curve. This approach is still typically interpreted as a minimal deviation from rational expectations.

Given their imperfect knowledge, agents attempt to learn the magnitudes of the relationships over time, based on the realizations of macroeconomic data that they observe. They update their beliefs at each t according to the constant-gain learning formula:

$$\hat{\phi}_t = \hat{\phi}_{t-1} + \bar{\mathbf{g}} R_t^{-1} X_t (Y_t - \hat{\phi}_{t-1}' X_t)' \quad (7)$$

$$R_t = R_{t-1} + \bar{\mathbf{g}} (X_t X_t' - R_{t-1}) \quad (8)$$

where $X_t \equiv [1, Y_{t-1}, \epsilon_t]'$, and $\hat{\phi}_t = [a_t, b_t, c_t]'$. The key coefficient of interest is $\bar{\mathbf{g}}$, the constant-gain coefficient. The gain governs the speed at which agents learn and adjust their beliefs to new information. The gain can also be interpreted as the degree of memory that agents have, given that they discount past information more heavily than recent observations (at the rate $(1 - \bar{\mathbf{g}})^j$ for observations falling j periods in the past). Given the PLM and the updated beliefs $\hat{\phi}_t$, the aggregate expectations entering in equations (1)-(3) are formed as

$$\hat{E}_t Y_{t+1} = \left(I + \hat{b}_t \right) \hat{a}_t + \hat{b}_t^2 Y_{t-1} + (\hat{c}_t \rho + \hat{b}_t \hat{c}_t) \epsilon_t + ds_t. \quad (9)$$

The expectation formation mechanism includes two components: one endogenous and due to learning about the economy and responding to observed conditions, and the second (ds_t , where d is

simply a selection matrix), exogenous. The latter represents the components of expectations that cannot be justified by the near-rational learning model. These terms, denoted by s_t define, as in Milani (2011, 2017), “sentiment”, or waves of excess optimism and pessimism, in the model.

Sentiments about output and inflation are assumed to evolve as

$$s_t = \rho_t^s s_{t-1} + \Sigma_t^s \zeta_t,$$

where $s_t = [s_t^y, s_t^\pi]'$, $\rho_t^s = [\rho_t^r, 0; 0, \rho_t^\pi]$, and $\Sigma_t^s = [\sigma_{r,t}, 0; 0, \sigma_{\pi,t}]$, with autoregressive coefficients and standard deviations allowed to switch before and after 1984, as for the other disturbances.

3.2 Bayesian Estimation of New Keynesian Model

The previous model, with aggregate expectations formed as in (9), can be expressed in state-space form as:

$$OBS_t = H_0 + H\mathbf{Y}_t \tag{10}$$

$$\mathbf{Y}_t = A_t + F_t\mathbf{Y}_{t-1} + G\tilde{\epsilon}_t \tag{11}$$

where OBS_t collects the observable variables to be matched in the estimation, \mathbf{Y}_t collects the endogenous variables, the expectations, and the exogenous disturbances, and $\tilde{\epsilon}_t$ collects the exogenous innovations.

We use real-time data (first-vintage) in the estimation of the DSGE model and in trying to match the individual forecasters’ expectations. The real-time series are obtained from the Real-Time data set for Macroeconomists, hosted on the Federal Reserve Bank of Philadelphia’s website. As explained there, quarterly vintages correspond to the real time data available to forecasters in February, May, August, November of each year. For this reason, when downloading the Federal Funds Rate, we also use the corresponding values in the same months.

Inflation is obtained as log first difference of the GDP Implicit Price Deflator; we use the log first difference of Real GDP for output growth. Expectations are given by the mean across our set of forecasters of one-period-ahead output growth and inflation expectations, using the series described in the previous section. In the estimation, we assume a piecewise-linear trend for output, following the evidence in Perron and Wada (2009). They find that when changes in the slope of the trend function are accounted for, there is no longer evidence of stochastic trends, and the resulting cyclical component aligns well with NBER recession dates. We allow for changing slopes between

the 1954:III-1973:III, 1973:IV-1994:IV, 1995:I-2007:I, and 2007:II-2016:III subperiods, based on the growth facts presented in Jones (2016).

A potential issue in the estimation regards the existence of a binding zero-lower bound (ZLB) starting in 2009, which could introduce a nonlinearity into our model. We solve this issue, as others have done, by using data on the ‘shadow’ short rate. Unlike the short-term US nominal interest rate, the shadow short rate is allowed to have negative values to capture a more accommodating stance of monetary policy at the ZLB (for example, due to unconventional monetary policy interventions). We utilize the shadow short rate made available by Krippner (2013) in place of the FFR from 2009:Q1 to 2016.

The last choice before estimating the New Keynesian model with adaptive learning concerns initial beliefs $\hat{\phi}_{t=0}$ and $R_{t=0}$. To obtain values for these coefficients, we utilize a presample estimation. We start from uninformative initial beliefs at the beginning of the pre-sample period, i.e. 1954:Q3, which is the first quarter of availability of Federal Funds rate data. We next run the estimation from 1954:Q3 to 1968:Q3 without expectations data to give us initial beliefs for our main estimation, which includes the expectations series starting in 1968:Q4. The likelihood is then computed for the 1968:Q4-2016:Q3 sample.

The model is estimated using Bayesian methods as in Milani (2007, 2011). Table 1 shows the chosen prior distributions, along with the posterior estimates for our vector of structural, disturbance, and learning, parameters. With survey expectations and learning, we estimate lower degrees of habit formation and inflation indexation ($\eta = 0.366$, $\gamma = 0.088$) than in corresponding models under rational expectations. The estimated response of monetary policy to inflation, as well as the degree of interest rate inertia, are higher during the 1979-1982 experiment, than in other periods. Cost-push shocks are close to *iid*, while natural rate shocks display significant persistence. Sentiment shocks, both related to output and inflation, are also persistent and they have comparable volatility to that of fundamental disturbances. In line with the Great Moderation literature, the standard deviations of most shocks fall in the second part of the sample. Finally, we provide an estimate of the best-fitting constant gain coefficient in a macroeconomic model with expectations matched to aggregate survey expectations. The posterior mean for the gain equals 0.015.

4 Individual Expectations

4.1 Perceived Model and Constant-Gain Learning

The previous section estimated the New Keynesian model at the aggregate level, which allowed us to obtain mean estimates of the model parameters and the filtered structural disturbances. We now turn to examining expectations at the individual level. We impose structure on our forecasters' expectation formation process by assuming that they form expectations from a near-rational model that allows for learning. Agents have a correctly-specified PLM, which has the same endogenous variables as the solution under rational expectations, and the same aggregate disturbances.

Hence, the PLM is the same as (6) for each individual forecaster j :

$$Y_t = a_t^j + b_t^j Y_{t-1} + c_t^j \epsilon_t + \nu_t^j \quad (12)$$

where $Y_t = [y_t, \pi_t, i_t]'$, $\epsilon_t = [\varepsilon_t^r, \varepsilon_t^u]'$, and a_t , b_t , and c_t are vectors and matrices of coefficients; ν_t is the usual regression error term. The estimation of the New Keynesian macroeconomic model allows us to include the filtered structural disturbances in the individual agents' information sets.³

While forecasters are assumed to base their expectations on a correctly-specified model, they may differ in their beliefs. Each forecaster j updates beliefs through constant-gain learning:

$$\widehat{\phi}_t^j = \widehat{\phi}_{t-1}^j + \overline{\mathbf{g}}^j (R_t^j)^{-1} X_t (Y_t - \widehat{\phi}_{t-1}^j X_t)' \quad (13)$$

$$R_t^j = R_{t-1}^j + \overline{\mathbf{g}}^j (X_t X_t' - R_{t-1}^j) \quad (14)$$

where $X_t \equiv [1, Y_{t-1}, \epsilon_t]'$, and $\widehat{\phi}_t^j = [a_t^j, b_t^j, c_t^j]'$. Expectations for each individual forecaster are, therefore, assumed to be formed as

$$\widehat{E}_t^j Y_{t+1} = \left(I + \widehat{b}_t^j \right) \widehat{a}_t^j + \left(\widehat{b}_t^j \right)^2 Y_{t-1} + \left(\widehat{c}_t^j \rho + \widehat{b}_t^j \widehat{c}_t^j \right) \epsilon_t + \widehat{d} s_t^j, \quad (15)$$

where ϵ_t collects the structural AR(1) disturbances and s_t denotes sentiment, or unjustified optimism and pessimism (that is, unjustified based on the state of the economy and the updated beliefs). Therefore, equation (15) provides a way to study and extract excess optimism and pessimism at the micro level.

Initial beliefs for each forecaster also need to be specified. Instead of fixing them at ad hoc values, we proceed in the following way. In the SPF dataset, forecasters may enter at different

³In the robustness Section 6.1, we will address the sensitivity of the results to the exclusion of information about disturbances: in that case, agents only use a VAR(1) plus constant in the observable variables as their PLM.

time periods, and thus, each individual may happen to utilize a different length of histories when producing an initial forecast. Therefore, we jointly estimate the initial beliefs, using the relevant pre-sample data for each forecaster, which would have been obtained by the forecaster using the same constant gain learning approach as

$$\widehat{\phi}_\tau^j = \left[\sum_{i=1}^{\tau} (1 - \bar{\mathbf{g}}^j)^{(i-1)} X_{\tau-i} X'_{\tau-i} \right]^{-1} \left[\sum_{i=1}^{\tau} (1 - \bar{\mathbf{g}}^j)^{(i-1)} X_{\tau-i} Y'_{\tau-i+1} \right] \quad (16)$$

where $X_t \equiv [1, Y_{t-1}, \epsilon_t]'$ denotes the regressors in forecaster j 's PLM and $\bar{\mathbf{g}}^j$ is the constant gain parameter *specific* to individual j . We assume $R_\tau = c * I$, a scaled identity matrix, with $c = 0.1$ to allow for a larger degree of uncertainty characterizing the initial beliefs. As described above, the estimation of the macroeconomic model allows us to obtain filtered series of the structural disturbances.

The following example will help to clarify how forecasters form initial beliefs. If an individual enters the survey in 1979:Q4 and stays for twenty periods, we would assume that she can observe the presample data from 1954:Q3 to 1979:Q3, discount them (we obtain the best-fitting gain in the estimation, which governs this discounting) and then update beliefs and form expectations for the subsequent years. In this way, the forecasters' learning speed and initial beliefs are both estimated, but in a parsimonious way.

Overall, forecasters are assumed to have both similarities and differences. Each individual j produces expectations using an adaptive learning model with the same set of variables, based on the correctly-specified solution under rational expectations. They all have access to the same set of presample data. However, forecasters expectations can vary based on the endogenous learning component influenced by $\bar{\mathbf{g}}$, and the exogenous sentiment shocks s_t displayed in (15).

4.2 Estimation

For each individual forecaster, we compare observations on their one-period ahead forecasts for output growth and inflation to their counterparts implied by the adaptive learning model. We find the best-fitting gain coefficient by minimizing the loss function implied by the mean squared errors for these series

$$\underset{\bar{\mathbf{g}}^j}{\operatorname{argmin}} \left[E_t^{j,obs} Z_{t+1} - \widehat{E}_t^j Z_{t+1} \right]' \left[E_t^{j,obs} Z_{t+1} - \widehat{E}_t^j Z_{t+1} \right] \quad (17)$$

where $E_t^{j,obs} Z_{t+1}$ denotes the observed survey forecast from the SPF for forecaster j , with $Z = [\Delta y, \pi]'$ and $\widehat{E}_t^j Z_{t+1}$ denotes the implied expectations obtained from the learning model. Thus, the best-fitting gain is the one that minimizes the mean squared forecast error for each forecaster. The unexplained component of each expectations series, that is the part that is not explained by the learning model with the best-fitting gain, is denoted as sentiment: forecasters are either more or less pessimistic than their near-rational learning model implies.

The heterogeneity across forecasters, therefore, can stem from two sources: different learning speeds, as measured by different constant gains (motivated by different agent’s perceptions about incoming structural change or by different memories), and different sentiment.

5 Empirical Results

We now examine the two potential sources of heterogeneity in individual forecasters’ expectations, as well as the implications these differences have for the economy.

5.1 Best-Fitting Constant Gain Parameter

The source of heterogeneity in individual respondents’ forecasts can stem, first, from their constant gain parameters \bar{g}^j . As explained above, this parameter can have two interpretations. First, it governs the speed at which agents adjust their beliefs to new information about the economy, possibly because they are concerned about future structural breaks of unknown form. It can also be interpreted as the degree of memory agents have about past data (Malmendier and Nagel, 2016, provide a behavioral explanation for differences in discounting, showing that older agents assign more weight to observations in the more distant past). To understand how the best-fitting constant gain parameter can be a source of heterogeneity in SPF respondents, we examine the \bar{g}^j that minimizes equation (17) for each forecaster j .

Figure 3 displays the value of the estimated best-fitting constant gain for each forecaster for each period she submits a response. The vertical axis denotes the value of the constant gain parameter from 0.0001 to 0.2. The horizontal axis indicates the date. In each time period, we represent the distribution of the best-fitting gain of forecasters with a boxplot. The target signifies the median, the edges of the box the 25th and 75th percentiles, the whiskers the extreme values not considered outliers, and the ‘+’ symbol outliers.

The results show that heterogeneity in individual expectations can stem from different values of

the best-fitting constant gain across SPF forecasters. Before the beginnings of the Great Moderation (i.e., prior to the mid to late 1980s), there exists larger dispersion in the value of \bar{g} . For instance, the interquartile range is much larger during this period relative to after the Great Moderation. The median gains during these decades often fluctuate and occasionally rise to values around 0.05. Several forecasters place a large degree of weight on new information about incoming structural change in the economy and they substantially discount the past. As the sample period moves into the 1990s and concludes in 2016, the distribution of constant gains tightens and steadily coalesces around smaller values of \bar{g} . The upper limits of the outliers and extreme values tend to take on smaller numbers; the interquartile range is also generally decreasing over this time period. The median value of the constant gain also centers around lower numbers (0.01-0.02) of \bar{g} , with a value of 0.0146 in the final period of our sample (i.e., 2016:Q3). Forecasters do not perceive a high prospect of structural change in the economy implying lower values of the best-fitting \bar{g} .

5.1.1 Beliefs

The previous subsection showed that the heterogeneity in forecasts can stem from different best-fitting constant gain values across forecasters. These differences then affect agents' beliefs, that is, the elements of $\hat{\phi}_t = [a_t, b_t, c_t]'$. Thus, a natural question regards the implications for the economy when there exists heterogeneity in forecasters' estimates of a, b , and c . To attempt to answer this question, we analyze selected beliefs of SPF respondents.

Figures 4 – 6 display the results. In each figure, we represent the distribution of the belief coefficient each period with a boxplot as before. Figure 4 shows the slope of the Phillips Curve parameter, that is, $b_{2,1}$. The inflation persistence parameter (i.e, $b_{2,2}$) is displayed in Figure 5. Figure 6 presents the policy parameter $b_{1,3}$, which governs the sensitivity of the output gap to changes in the interest rate. In addition, the belief coefficients correspond to the best-fitting gain of each forecaster.

Three important takeaways emerge after examining Figures 4 – 6. First, looking at the evolution of the slope of the Phillips Curve $b_{2,1}$, there exists notable shifts in the values of the coefficients during recessions: the curve appears to steepen during recessions. Second, the median value of the perceived inflation persistence parameter (i.e., $b_{2,2}$) increases in the middle of the sample and then declines again at the end. Finally, forecasters perceive the influence of policy on output gap (i.e., $b_{1,3}$) to be less effective as time elapses. Figure 6 shows that the median value of $b_{1,3}$ stays negative

the entire period, but moves in an upward trajectory towards zero. The median coefficients are close to -2 in the 1970s, but they are revised closer to -0.5 after the Great Recession. This result is consistent with the VAR evidence from (Boivin and Giannoni, 2006) on the reduced effectiveness of monetary policy over the sample.

5.1.2 Evolving Beliefs, Heterogeneity, and Responses to Shocks

The response of the economy to shocks will be substantially different depending on the state of private-sector beliefs. In Figure 7, we show the responses of output, inflation, and interest rates, to the structural shocks (natural rate, cost-push, and monetary policy), which would exist if the aggregate beliefs in the New Keynesian model were assumed to be equal to those held by each forecaster j . We show the responses in the early part of the sample (1971:Q1) and in the late part, before the Great Recession (2006:Q1).

As the figure shows, the heterogeneity is more pronounced in 1971, as the ranges of responses are usually wider. For example, a portion of forecasters believes that positive cost-push shocks have only mild (and possibly even positive) effects on the economy, while others believe that cost-push shocks have extremely large recessionary effects on output. The majority of forecasters lay in between these two extremes. In a self-referential system, individual forecasters' beliefs can be partially self-fulfilling: if agents perceive supply, or any other, shocks to be particularly effective, they will indeed play a larger role than in an economy in which their effects are perceived as trivial in the formation of expectations. The responses for 2006 indicate that significant dispersion still exists regarding the magnitude of the effects, but without major disagreement on the sign and overall shape of the responses.

Figure 8 shows, instead, the contribution of each shock to the forecast error variance of output, when, as before, the aggregate beliefs are assumed to be fixed, in turn, to exactly match the beliefs of each single forecaster j . It is apparent from the figure that the state of beliefs plays a central role for the transmission and importance of shocks: the natural rate shock can explain anywhere from close to zero to almost all output fluctuations. Depending on forecasters' perceptions, monetary policy shocks can end up explaining between few percentage points and 60% of business cycle movements.

5.2 Sentiment

Individual forecasters have been assumed to construct forecasts using the same adaptive learning model in equation (12), but with potentially different constant gains across forecasters. However, a natural question regards other origins that may contribute to the heterogeneity in expectations. In this section, we examine a potential second source of differences: sentiment shocks. These shocks are defined as the difference between observed expectations and model-implied expectations for one-period ahead output and inflation:

$$s_t^{j,y} = E_t^{j,obs} y_{t+1} - \hat{E}_t^j y_{t+1} \quad (18)$$

$$s_t^{j,\pi} = E_t^{j,obs} \pi_{t+1} - \hat{E}_t^j \pi_{t+1}. \quad (19)$$

As stated in Milani (2011, 2017), these shocks can be defined as waves of excess optimism and pessimism by agents about the economy in a particular time period.

The sentiment shocks of each individual forecasters over time corresponding to the best-fitting gain are displayed in Figures 9 and 10. As in the previous section, we represent the distribution each period with a boxplot. Figure 9 corresponds to the sentiment shock for expectations of one-period ahead output. Figure 10 shows the sentiment shock for expectations of one-period ahead inflation.

Figures 9 and 10 show that an additional source of heterogeneity in forecasts can stem from different waves of optimism and pessimism for each individual forecaster. These disparities are apparent when examining pre- and post-Great Moderation periods in the U.S. In the former, the distribution of sentiment shocks is very wide reflecting the higher volatility of macroeconomic variables during this time period. After mid-1980s, forecasters' sentiment shocks are still different, but not as volatile and much tighter. They tend to cluster together during the Great Moderation era. In addition, the output sentiment shock tends to align with downturns in the U.S. economy. For instance, during the 2001 and 2007 - 2009 recessions in the U.S., the median value for the output shock turns negative. This result is not surprising as agents are becoming more pessimistic about the economy and matches well with Milani (2011, 2017) who shows these sentiment shocks can explain greater than 40% of business cycle fluctuations in the U.S.

An important question relates to the process of sentiment for individual forecasters. To answer this question, we fit the sentiment shocks to two types of models: an AR(1) plus constant and a VAR(1) plus constant. For the first, we specify each type of sentiment shock as an AR(1) plus

constant model:

$$s_t^{j,y} = a^{j,y} + b^{j,y} s_{t-1}^{j,y} + \varepsilon_t^{j,y} \quad (20)$$

$$s_t^{j,\pi} = a^{j,\pi} + b^{j,\pi} s_{t-1}^{j,\pi} + \varepsilon_t^{j,\pi} \quad (21)$$

where $\varepsilon_t^{j,y}$ and $\varepsilon_t^{j,\pi}$ are the usual (scalar) white noise error term assumed not correlated with each other. In addition, the VAR(1) plus constant model allows for the possibility of correlation across output and inflation sentiment shocks and is defined as

$$s_t^j = \Phi_0^j + \Phi_1^j s_{t-1}^j + \varepsilon_t^{j,\Phi} \quad (22)$$

where $s_t^j = [s_t^{j,y}, s_t^{j,\pi}]'$, $\varepsilon_t^{j,\Phi}$ is a 2×1 vector of usual white noise error terms and the coefficient matrices are given by

$$\Phi_0 = \begin{bmatrix} \phi_{0,1}^j \\ \phi_{0,2}^j \end{bmatrix} \quad (23)$$

$$\Phi_1 = \begin{bmatrix} \phi_{1,1}^j & \phi_{1,2}^j \\ \phi_{2,1}^j & \phi_{2,2}^j \end{bmatrix} \quad (24)$$

For each forecaster and model, we estimate the coefficients using OLS (line-by-line OLS in the VAR(1) plus constant case). The results are presented in Figures 11 – 13. Figures 11 and 12 display histograms of the coefficients across forecasters in AR(1) plus constant specifications. Figure 13 denotes the histograms of the coefficients across SPF respondents in the VAR(1) plus constant model.

Three important takeaways emerge after examining Figures 11 – 13. First, the sentiment shocks seem to be (slightly) biased across both types of models. The histogram of the estimated constants in the AR(1) plus constant and the VAR(1) plus constant models are skewed right towards positive values. In addition, the sentiment shocks of SPF forecasters seem to exhibit persistence. Across both AR(1) plus constant and VAR(1) plus constant specifications, the histogram of the estimated values of the autoregressive terms are centered over positive values. Finally, the shocks do not seem to be highly (if at all) correlated across output and inflation sentiment shocks. This result is seen with the VAR(1) plus constant model, which allows for the possibility of this correlation with estimation of the off-diagonal coefficients (i.e., $\phi_{1,2}$ and $\phi_{2,1}$). In Figure 13, the histogram of these estimated coefficients is slightly skewed left. However, it is centered over zero suggesting minimal to no correlation across sentiment shocks.

Overall, the heterogeneity across forecasters can stem from two sources: (1) different learning speeds, as measured by different constant gains (motivated by different agent’s perceptions about incoming structural change or by different memories); and (2) different degrees of excess optimism and pessimism. The distribution of $\bar{\mathbf{g}}^j$ across forecasters tends to be more dispersed pre-Great Moderation and tightens up around lower values towards the end of the sample. In addition, the individual sentiment shocks track the U.S. business cycle fairly well. These shocks also seem to be biased upwards and exhibit persistence across time. In many periods, particularly at the beginning of the sample, forecasters who display excess optimism for either output or inflation coexist with others who display excess pessimism. But this is mostly due to forecasters in the tails of the distribution. Overall, sentiments don’t cancel out: the 25-75% interquantile ranges show that many forecasters move in herd, tending to be overly optimistic or pessimistic at the same time.

6 Robustness

6.1 Alternative Forecasting Model

Our benchmark results showed that the heterogeneity in forecasters can be explained by different constant gains and sentiment shocks under the assumption that survey forecasters use an adaptive learning rule to construct expectations. Their adaptive learning PLM included a constant, lagged endogenous variables, and knowledge of structural disturbances, that is, equation (12). However, a natural question regards the structure of this forecasting model. Specifically, how would the heterogeneity across forecasters change if their forecasting model assumed less knowledge about the economy? In other words, how would the results change if the PLM of SPF respondents was based on a misspecified, but realistic, model?

We attempt to answer the previous question by comparing the benchmark results to a PLM that consists of a VAR(1) plus constant. In other words, equation (12) is modified to be the following:

$$Y_t = a^j + b^j Y_{t-1} + \epsilon_t^j. \quad (25)$$

We repeat the benchmark exercise of Section 5.1 and display the outcomes of the best-fitting gain in Figure 14. As before, we represent the distribution of individual forecasts with a boxplot.

With a misspecified forecasting model, the best-fitting $\bar{\mathbf{g}}^j$ across SPF respondents tend to cluster over smaller values as compared to the benchmark case in Figure 3. The extent of time variation is also more modest.

6.2 SPF Respondents Submit Forecasts for at Least 20 Periods

A potential issue with using a survey dataset such as the SPF is the entry and exit of respondents as described by Mansky (2011). To remedy this issue, our baseline exercise included only those respondents that submit forecasts for at least ten periods. We chose this number partly so that we had a sufficient number of data points per respondent. However, given that this observation requirement is somewhat arbitrary, we analyze the benchmark results for those respondents that submit forecasts for at least twenty periods.

The results of this sensitivity analysis show that the baseline outcomes are largely unchanged. Figure 15 reports the best-fitting constant gains for each forecasters with a boxplot representing the distribution of forecasters as in Section 5.1. As in the baseline case, there exists larger dispersion in the value of $\bar{\mathbf{g}}$ before the Great Moderation relative to after this time period. Thus, the results of this section indicate that our benchmark requirement of ten observations for a respondent to remain in our dataset allows sufficient information about the distribution of best-fitting constant gains.

6.3 Alternative Initialization of R_τ

Our benchmark results of Section 5 were obtained under the assumption that the initial R matrix (i.e., R_τ) of each individual forecaster was equal to $c * I$, a scaled identity matrix, with $c = 0.1$. This specification allowed for a larger degree of uncertainty characterizing the initial beliefs of individual forecasters. However, it is natural to investigate the results under an alternative parameterization of R_τ for each forecaster.

We proceed by estimating the initial beliefs of R for each forecaster as an additional extension exercise. Similar to the assumption for the initial $\hat{\phi}_t$, each forecaster is assumed to use their relevant pre-sample data using the same constant-gain learning approach to estimate R_τ :

$$R_\tau = \bar{\mathbf{g}}^j \sum_{i=1}^{\tau} (1 - \bar{\mathbf{g}}^j)^{(i-1)} X_{\tau-i} X'_{\tau-i}. \quad (26)$$

To be consistent with Section 5, each forecaster is assumed to scale this initial matrix R by $c = 0.1$, hence making our prior knowledge about $\hat{\phi}_t$ more diffuse. Moreover, we perform the benchmark best-fitting constant gain exercise of Section 5.1 and display the results in Figure 16.

The distributions of the best-fitting constant gains still highlight significant heterogeneity over the whole sample. The median values have similar magnitudes as before, but without the spikes

at 0.05. They also converge to values around 0.02 at the end of the period, as in our benchmark analysis in Figure 3 and aggregate results in Table 1.

7 Conclusions

We analyze the formation of expectations at the individual forecaster level. We treat forecasters in the same way as we would treat agents in a benchmark New Keynesian model. They are assumed to dispose of a perceived model of the economy that resembles the MSV solution under rational expectations, and they have the same information sets: therefore, they have a correctly-specified model, and they observe the same endogenous variables and disturbances that they would in the RE solution.

Their expectations can, however, be heterogeneous since different forecasters are allowed to have different gain coefficients. The best-fitting gain coefficient for each individual forecaster in the sample is estimated by minimizing the mean squared errors between the actual forecast and the forecast implied by the corresponding PLM. Moreover, each forecaster may be subject to sentiment, i.e., waves of excess optimism and pessimism, identified as in Milani (2011, 2017).

Our results reveal gain coefficients at the micro level that are, on average, of similar values to those estimated on aggregate data for macro models. The gains are, however, heterogeneous, with a dispersion that is higher in the 1970s and 1980s and much smaller by the end of the sample. The median gains are occasionally higher in the 1970s and 1980s. As a consequence, beliefs about key economic magnitudes are also heterogeneous and vary over time: for example, perceptions are consistent with a declining effectiveness of monetary policy over time.

Finally, we provide evidence at the micro level of the kind of sentiment shocks that have been shown in the recent literature to be important determinants of business cycles. Individual excesses of optimism and pessimism do not cancel out in the aggregate, but they are instead consistent with aggregate contagion or herd behavior.

In future research, it will be important to start from the evidence of heterogeneity at the micro level and investigate more thoroughly the implications of heterogeneous beliefs and sentiment for the macroeconomy.

References

- [1] Andrade, P., Crump, R.K., Eusepi, S., and Moench, E. (2016). “Fundamental disagreement,” *Journal of Monetary Economics*, 83(C), 106-128.
- [2] Angeletos, G. M., Collard, F., and Dellas, H. (2018). “Quantifying confidence,” *Econometrica*, 86(5), 1689-1726.
- [3] Baranowski, R. (2014). “Individual heterogeneity and adaptive learning in survey data,” mimeo, Coe College.
- [4] Barucci, E., (1999). “Heterogeneous beliefs and learning in forward looking economic models,” *Journal of Evolutionary Economics*, 9(4), 453-464.
- [5] Benhabib, J., Wang, P., and Wen, Y. (2015). “Sentiments and aggregate demand fluctuations,” *Econometrica*, 83(2), 549-585.
- [6] Berardi, M. (2019). “A probabilistic interpretation of the constant gain algorithm,” mimeo, University of Manchester, U.K.
- [7] Berardi, M., and Galimberti, J.K. (2017). “Empirical calibration of adaptive learning,” *Journal of Economic Behavior and Organization*, 144(C), 219-237.
- [8] Boivin, J., and Giannoni, M.P. (2006). “Has Monetary Policy Become More Effective?,” *Review of Economics and Statistics*, 88(3), 445-462.
- [9] Branch, W.A. (2004), “The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations,” *The Economic Journal*, 114(497): 592-621.
- [10] Branch, W.A. and Evans, G.W., (2006). “A simple recursive forecasting model,” *Economics Letters*, 91(2), 158-166.
- [11] Branch, W. A., and Evans, G. W. (2007). “Model uncertainty and endogenous volatility,” *Review of Economic Dynamics*, 10(2), 207-237.
- [12] Bräuning, C., and van der Crujisen, C. (2019). “Learning dynamics in the formation of European inflation expectations,” *Journal of Forecasting*, 38(2), 122-135.

- [13] Charalampidis, N., and Milani, F. (2020). “Evolving beliefs and animal spirits in the euro area,” mimeo, University of California Irvine.
- [14] Chatterjee, P., and Milani, F. (2019). “Perceived uncertainty shocks, excess optimism-pessimism, and learning in the business cycle,” mimeo, University of California Irvine.
- [15] Coibion, O., and Gorodnichenko, Y. (2012). “What can survey forecasts tell us about information rigidities?,” *Journal of Political Economy*, 120(1), 116-159.
- [16] Cole, S.J., (2020a). “Learning and the effectiveness of central bank forward guidance,” forthcoming, *Journal of Money, Credit, and Banking*.
- [17] Cole, S.J., (2020b). “The limits of central bank forward guidance under learning,” forthcoming, *International Journal of Central Banking*.
- [18] Cole, S., and F. Milani (2019). “The misspecification of expectations in new keynesian models: a DSGE-VAR approach,” *Macroeconomic Dynamics*, 23(3), 974-1007.
- [19] Dave, C., and Malik, S. (2017), “A tale of fat tails,” *European Economic Review*, 100(C), 293-317.
- [20] Eusepi, S., and Preston, B. (2010). “Central bank communication and expectations stabilization,” *American Economic Journal: Macroeconomics*, 2(3), 235-71.
- [21] Eusepi, S., and Preston, B. (2011). “Expectations, learning, and business cycle fluctuations,” *American Economic Review*, 101(6), 2844-72.
- [22] Evans, G. W. and Honkapohja, S. (2001). *Learning and expectations in macroeconomics*. Princeton University Press, Princeton, NJ.
- [23] Evans, G. W., Honkapohja, S., and Mitra, K. (2009). “Anticipated fiscal policy and adaptive learning,” *Journal of Monetary Economics*, 56(7), 930-953.
- [24] Evans, G. W., and Honkapohja, S. (2013). “Learning as a rational foundation for macroeconomics and finance,” *Rethinking expectations: The way forward for macroeconomics*, 68.
- [25] Gaspar, V., Smets, F., and Vestin, D. (2006). “Adaptive learning, persistence, and optimal monetary policy,” *Journal of the European Economic Association*, 4(2-3), 376-385.

- [26] Honkapohja, S., and Mitra, K. (2006). “Learning stability in economies with heterogeneous agents,” *Review of Economic Dynamics*, 9(2), 284-309.
- [27] Jones, C. (2016). “The Facts of Economic Growth,” Chapter 1, *Handbook of Macroeconomics*, 2016, vol. 2, 3-69, Elsevier.
- [28] Krippner, L. (2013). “Measuring the stance of monetary policy in zero lower bound environments,” *Economics Letters*, 118(1), 135-138.
- [29] Malmendier, U., and Nagel, S. (2016). “Learning from inflation experiences,” *The Quarterly Journal of Economics*, 131(1), 53-87.
- [30] Mankiw, N.G., Reis, R., and Wolfers, J. (2004). “Disagreement about inflation expectations,” NBER Chapters, in: *NBER Macroeconomics Annual*, Volume 18, pages 209-270, National Bureau of Economic Research, Inc.
- [31] Mansky, C.F. (2011). “Interpreting and Combining Heterogeneous Survey Forecasts,” in M. Clements and D.Hendry (editors), *Oxford Handbook on Economic Forecasting*, Chapter 16, 457-472, Oxford: Oxford University Press.
- [32] Markiewicz, A., and Pick, A. (2014). “Adaptive learning and survey data,” *Journal of Economic Behavior & Organization*, 107(PB), 685-707.
- [33] McConnell, M. M., and Perez-Quiros, G. (2000). “Output fluctuations in the United States: What has changed since the early 1980’s?,” *American Economic Review*, 90(5), 1464-1476.
- [34] Milani, F. (2007). “Expectations, learning and macroeconomic persistence,” *Journal of Monetary Economics*, 54(7), 2065-2082.
- [35] Milani, F. (2011). “Expectation shocks and learning as drivers of the business cycle,” *Economic Journal*, 121(552), 379-401.
- [36] Milani, F. (2014). “Learning and time-varying macroeconomic volatility,” *Journal of Economic Dynamics and Control*, 47, 94-114.
- [37] Milani, F. (2017). “Sentiment and the US business cycle,” *Journal of Economic Dynamics and Control*, 82, 289-311.

- [38] Orphanides, A., and Williams, J.C. (2005). “Imperfect knowledge, inflation expectations and monetary policy.” In *Inflation Targeting*, edited by Ben S. Bernanke and Michael Woodford. University of Chicago Press.
- [39] Perron, P., and Wada, T. (2009). “Let’s take a break: Trends and cycles in US real GDP”, *Journal of Monetary Economics*, 56(6), 749-765.
- [40] Pfajfar, D., and Santoro, E. (2010). “Heterogeneity, learning and information stickiness in inflation expectations,” *Journal of Economic Behavior & Organization*, 75(3), 426-444.
- [41] Preston, B. (2006). “Adaptive learning, forecast-based instrument rules and monetary policy,” *Journal of Monetary Economics*, 53(3), 507-535.
- [42] Sargent, T.J., (1993). *Bounded Rationality in Macroeconomics*. Oxford University Press: Oxford, U.K.
- [43] Sargent, T.J., (1999). *The Conquest of American Inflation*. Princeton: Princeton University Press.
- [44] Slobodyan, S., and Wouters, R. (2012). “Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models”, *American Economic Journal: Macroeconomics*, 4(2), 65-101.
- [45] Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

	Prior Distribution	Posterior Mean	95% Posterior Interval
IES	$\sigma \sim \Gamma(1, 0.75)$	0.694	[0.454,1.040]
Calvo	$\xi_p \sim B(0.6, 0.05)$	0.889	[0.864,0.911]
Habits	$\eta \sim B(0.5, 0.2)$	0.366	[0.229,0.531]
Indexation	$\gamma \sim B(0.5, 0.2)$	0.088	[0.016,0.212]
IRS	$\rho_{pre79} \sim B(0.7, 0.2)$	0.807	[0.679,0.916]
	$\rho_{79-82} \sim B(0.7, 0.2)$	0.929	[0.887,0.969]
	$\rho_{post82} \sim B(0.7, 0.2)$	0.613	[0.411, 0.961]
Resp. Infl.	$(\chi_{\pi,pre79} - 1) \sim \Gamma(0.5, 0.25)$	1.350	[1.099,1.741]
	$(\chi_{\pi,79-82} - 1) \sim \Gamma(0.5, 0.25)$	1.644	[1.146,2.425]
	$(\chi_{\pi,post82} - 1) \sim \Gamma(0.5, 0.25)$	1.457	[1.111,1.904]
Resp. Output	$\chi_{y,pre79} \sim \Gamma(0.25, 0.15)$	0.237	[0.118,0.372]
	$\chi_{y,79-82} \sim \Gamma(0.25, 0.15)$	0.200	[0.076,0.348]
	$\chi_{y,post82} \sim \Gamma(0.25, 0.15)$	0.209	[0.110,0.289]
AR Nat. Rate	$\rho_{r,pre84} \sim B(0.5, 0.2)$	0.727	[0.546,0.890]
	$\rho_{r,post84} \sim B(0.5, 0.2)$	0.841	[0.730,0.923]
AR Cost-push	$\rho_{u,pre84} \sim B(0.5, 0.2)$	0.066	[0.011,0.177]
	$\rho_{u,post84} \sim B(0.5, 0.2)$	0.016	[0.003,0.040]
AR Out. Sent.	$\rho_{s_y,pre84} \sim B(0.5, 0.2)$	0.559	[0.336,0.761]
	$\rho_{s_y,post84} \sim B(0.5, 0.2)$	0.694	[0.576,0.819]
AR Infl Sent.	$\rho_{s_\pi,pre84} \sim B(0.5, 0.2)$	0.854	[0.761,0.953]
	$\rho_{s_\pi,post84} \sim B(0.5, 0.2)$	0.561	[0.358,0.719]
Std. Nat. Rate	$\sigma_{r,pre84} \sim IG(0.3, 1)$	1.011	[0.697,1.417]
	$\sigma_{r,post84} \sim IG(0.3, 1)$	0.547	[0.383,0.754]
Std. Cost-push	$\sigma_{u,pre84} \sim IG(0.3, 1)$	0.447	[0.366,0.550]
	$\sigma_{u,post84} \sim IG(0.3, 1)$	0.241	[0.214,0.271]
Std. MP	$\sigma_{\varepsilon,pre79} \sim IG(0.3, 1)$	0.266	[0.215,0.324]
	$\sigma_{\varepsilon,79-82} \sim IG(0.3, 1)$	0.254	[0.226,0.287]
	$\sigma_{\varepsilon,post82} \sim IG(0.3, 1)$	0.280	[0.091,0.840]
Std. Out. Sent.	$\sigma_{s_y,pre84} \sim IG(0.3, 1)$	0.984	[0.791,1.204]
	$\sigma_{s_y,post84} \sim IG(0.3, 1)$	0.494	[0.419,0.573]
Std. Infl Sent.	$\sigma_{s_\pi,pre84} \sim IG(0.3, 1)$	0.331	[0.274,0.409]
	$\sigma_{s_\pi,post84} \sim IG(0.3, 1)$	0.183	[0.161,0.208]
Constant Gain	$\bar{g} \sim B(0.025, 0.01)$	0.015	[0.010,0.021]

Table 1: Prior and Posterior Distributions for estimated New Keynesian model coefficients.

Note: Γ refers to Gamma distribution, B to Beta, and IG to Inverse Gamma. The numbers in parenthesis refer to the chosen means and standard deviations for each distribution.

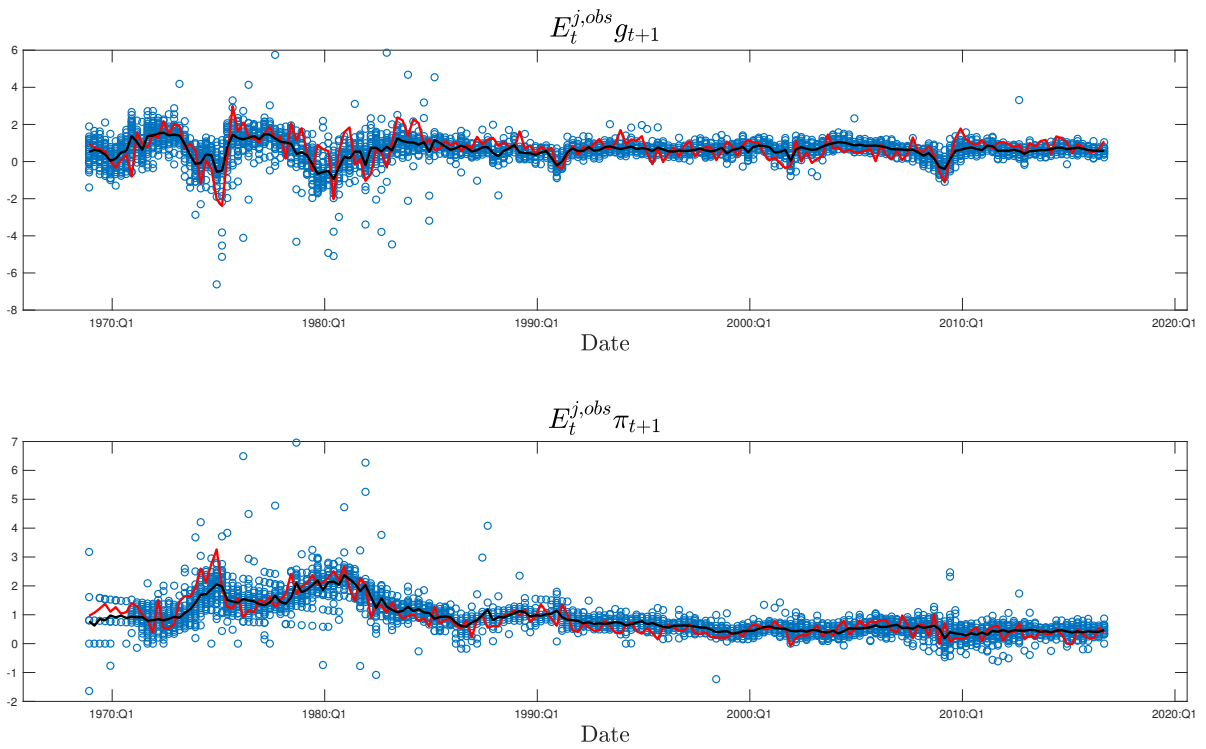


Figure 1: Evolution of Expected Output Growth and Inflation Over the Sample.

Note: ‘ g ’ denotes the growth rate of output. Blue Circles: Individual SPF Forecasters. Black Line: Implied Mean of SPF Forecasters. Red Line: Actual Realized Series.

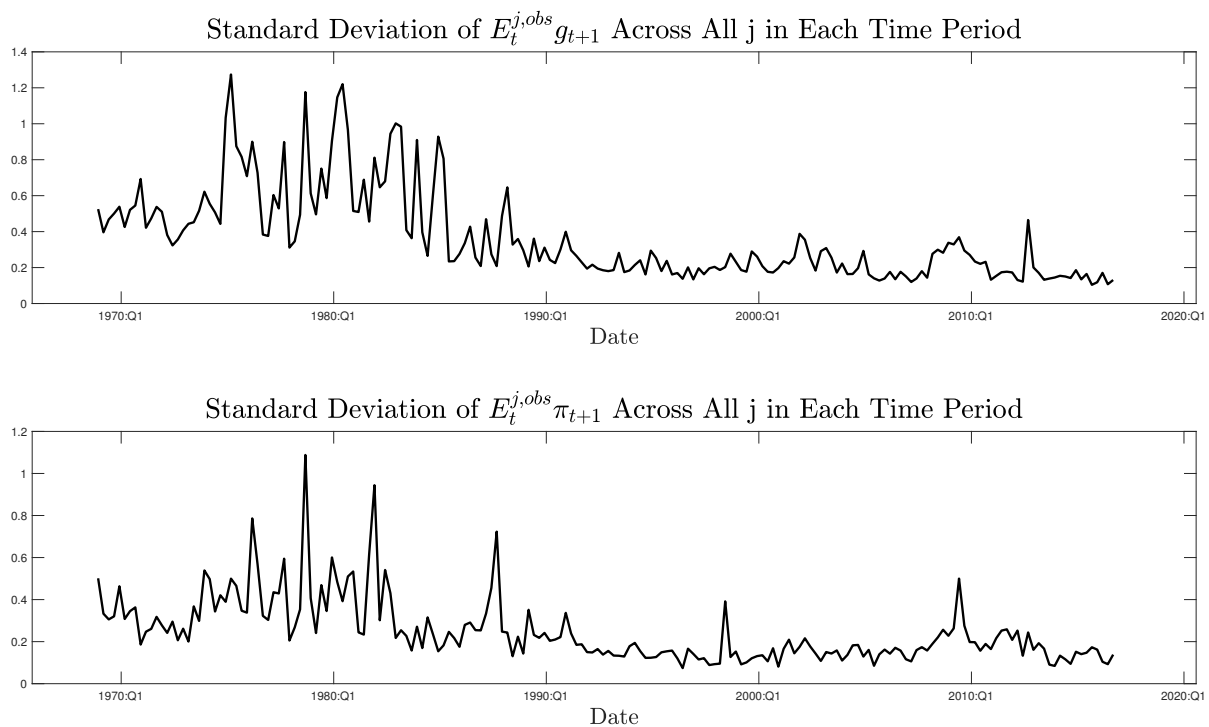


Figure 2: Disagreement Across SPF Forecasters Each Time Period

Note: ‘ g ’ denotes the growth rate of output. The disagreement across forecasters each time period is measured as the cross-sectional standard deviation.

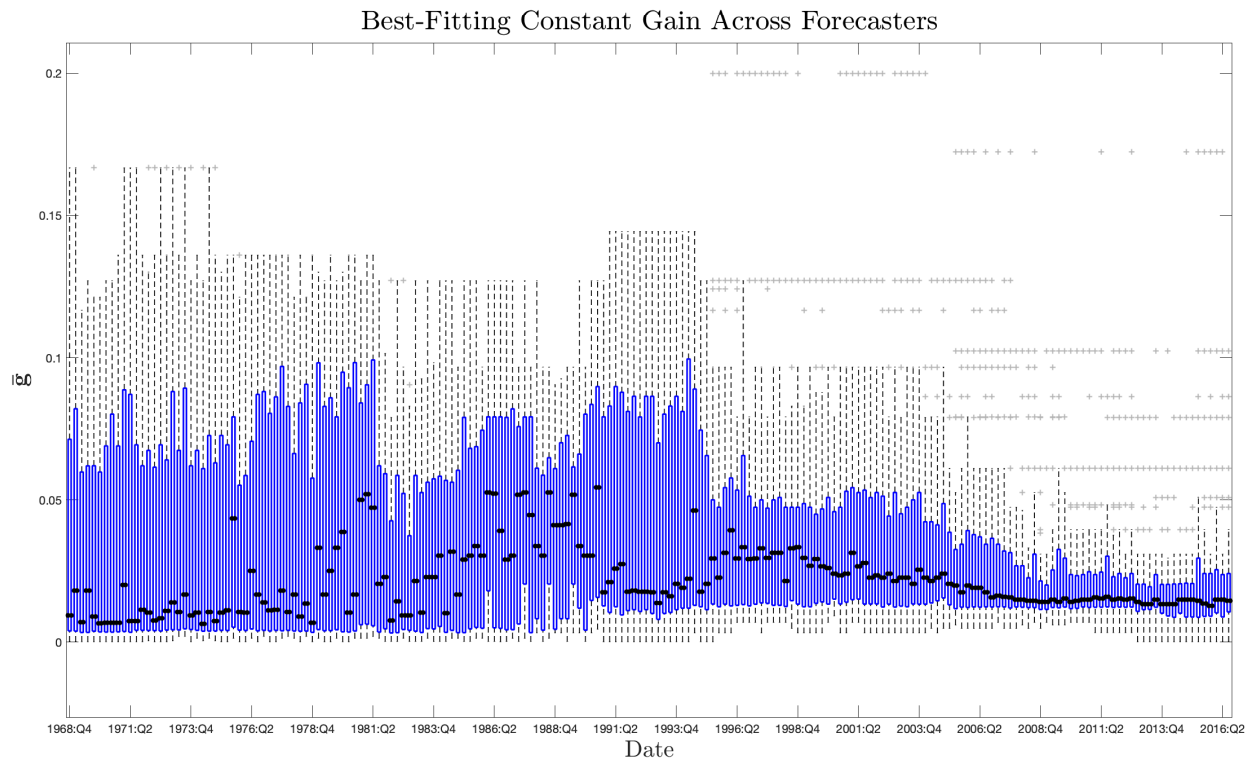


Figure 3: Distribution of the Best-Fitting Constant Gain of Individual Forecasters Each Time Period

Note: The distribution of the best-fitting constant gain of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

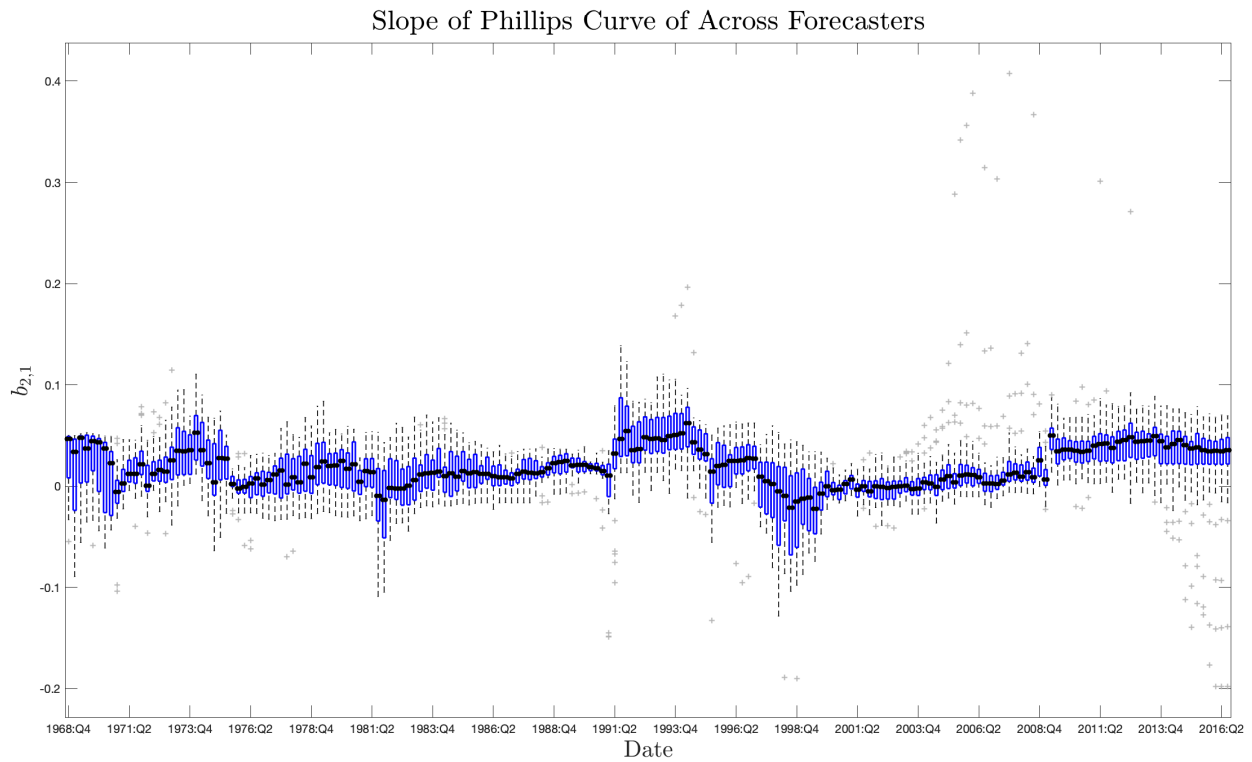


Figure 4: Estimate of Slope of Phillips Curve of Individual Forecasters Each Time Period

Note: The distribution of $b_{2,1}$ of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

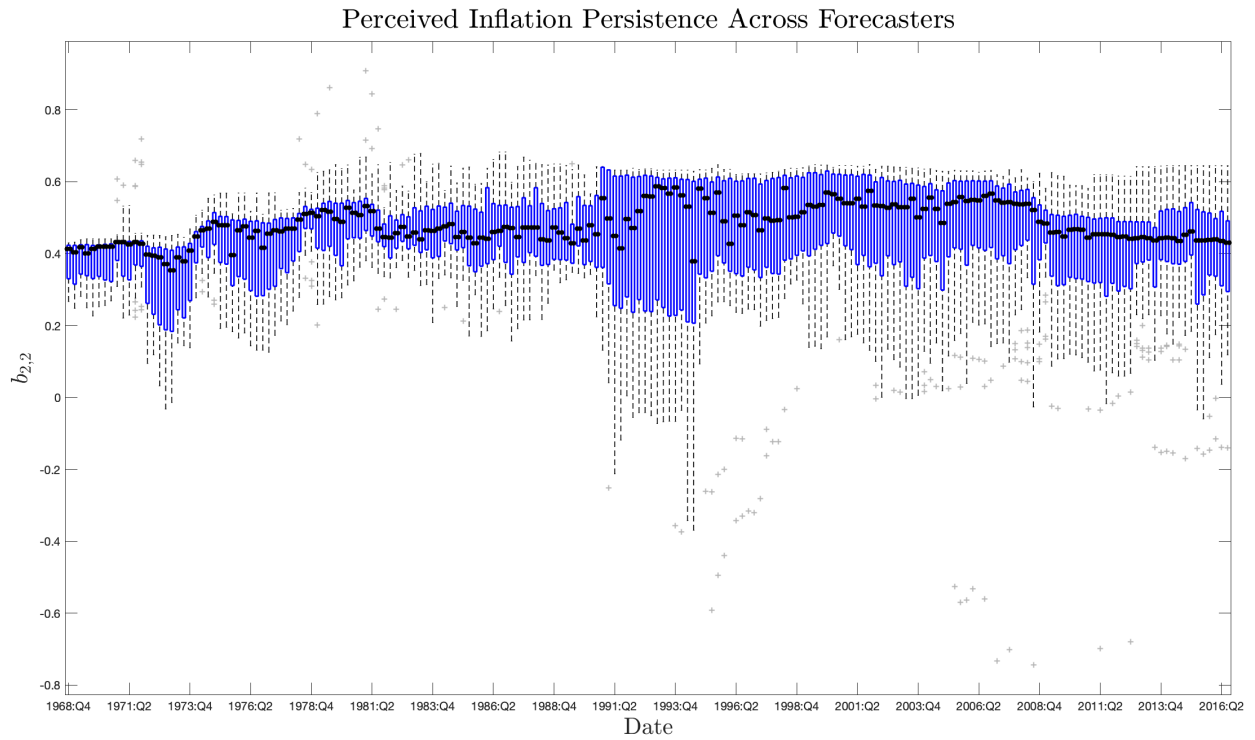


Figure 5: Estimate of Perceived Inflation Persistence Parameter of Individual Forecasters Each Time Period

Note: The distribution of $b_{2,2}$ of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

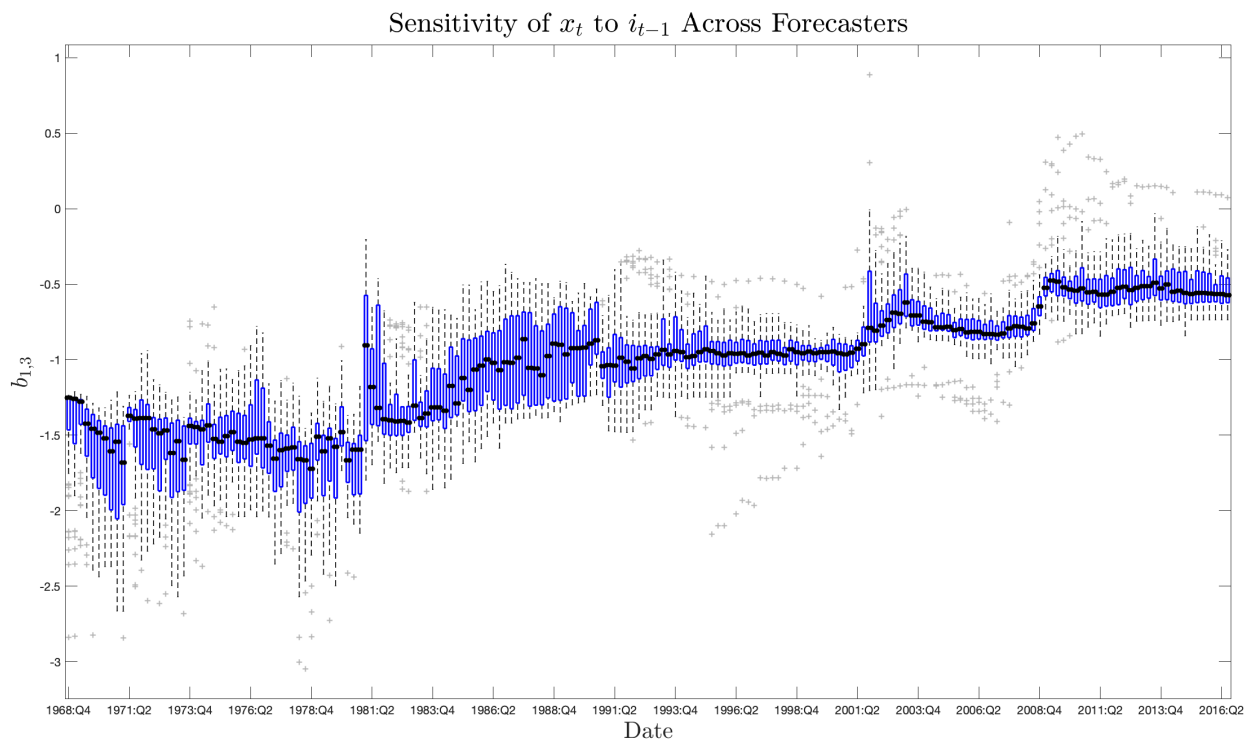


Figure 6: Estimate of Policy Parameter of Individual Forecasters Each Time Period

Note: The distribution of $b_{1,3}$ of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

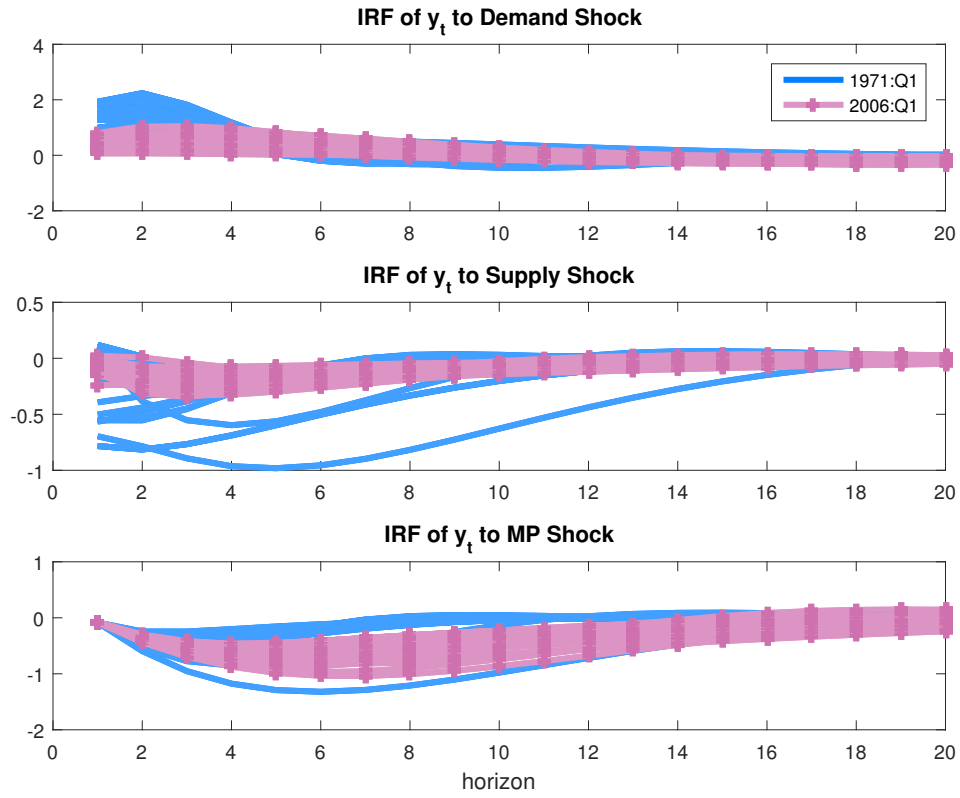


Figure 7: Impulse Response Functions of Output to Demand (Natural Rate), Supply (Cost-Push), and Monetary Policy Shocks, Across Heterogeneous Beliefs.

Note: We fix aggregate beliefs to equal the beliefs of each forecaster j , $j = 1, \dots, N$. The range of impulse responses under heterogeneous beliefs are shown for 1971:Q1 and 2006:Q1.

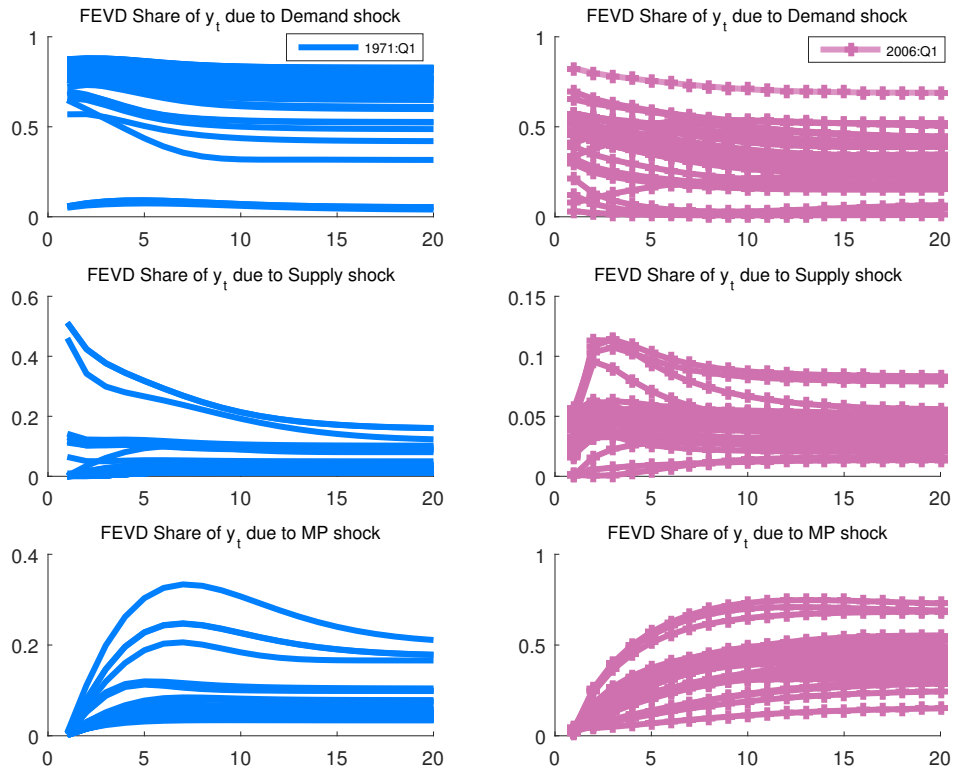


Figure 8: Forecast Error Variance Decomposition. Share of output forecast error variance explained by demand, supply, and monetary policy shocks, across heterogeneous beliefs. Aggregate beliefs in the model are fixed to match, in turn, the beliefs of each forecaster j , $j = 1, \dots, N$. The left panels show the variance shares corresponding to each set of beliefs for 1971:Q1, the right panels those for 2006:Q1.

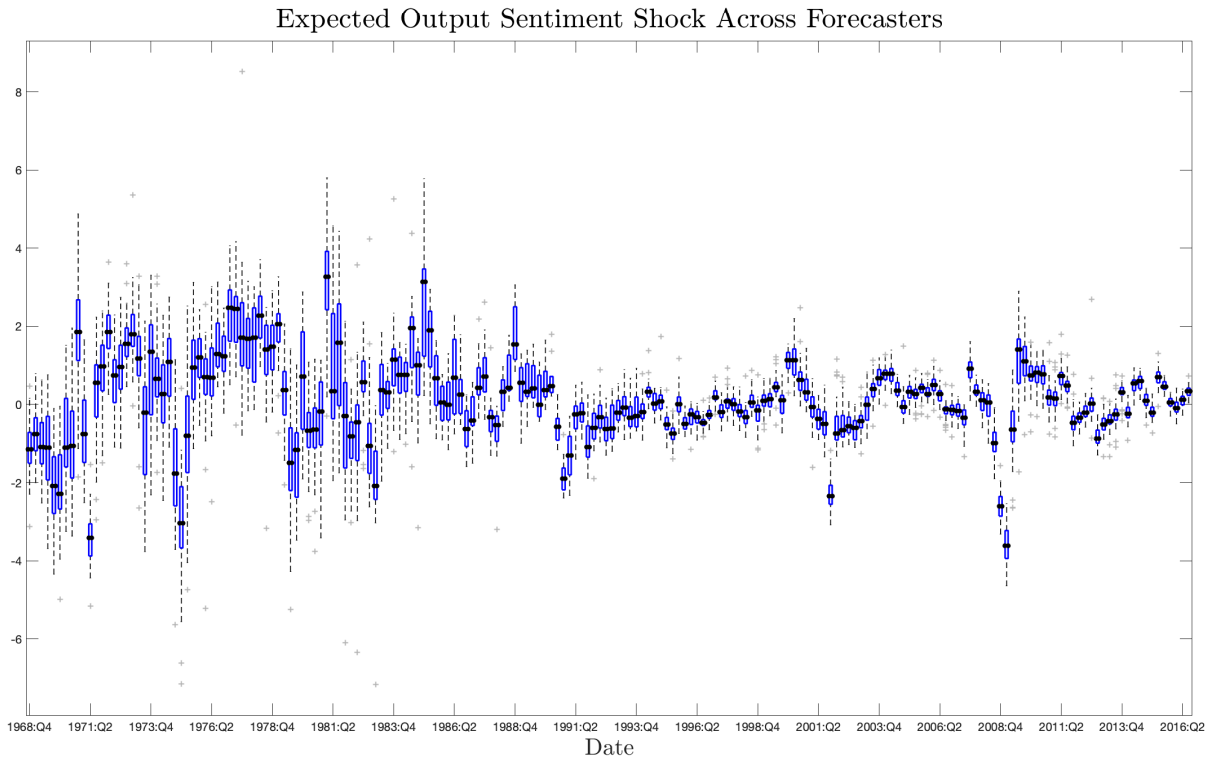


Figure 9: Expected Output Sentiment Shock of Individual Forecasters Each Time Period

Note: The distribution of expected output growth sentiment shock of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

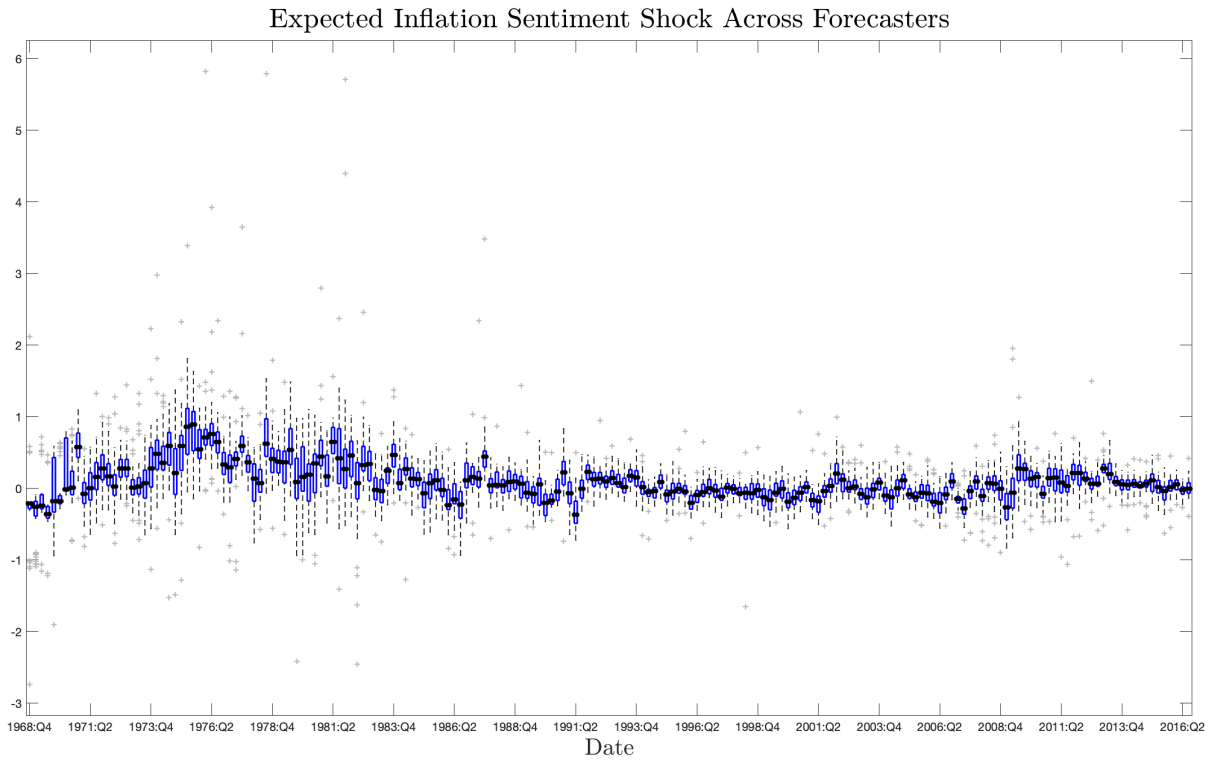


Figure 10: Expected Inflation Sentiment Shock of Individual Forecasters Each Time Period

Note: The distribution of expected inflation sentiment shock of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

Output Sentiment as AR(1) Plus Constant: $s_t^{j,y} = a^y + b^y s_{t-1}^{j,y} + \varepsilon_t^{j,y}$

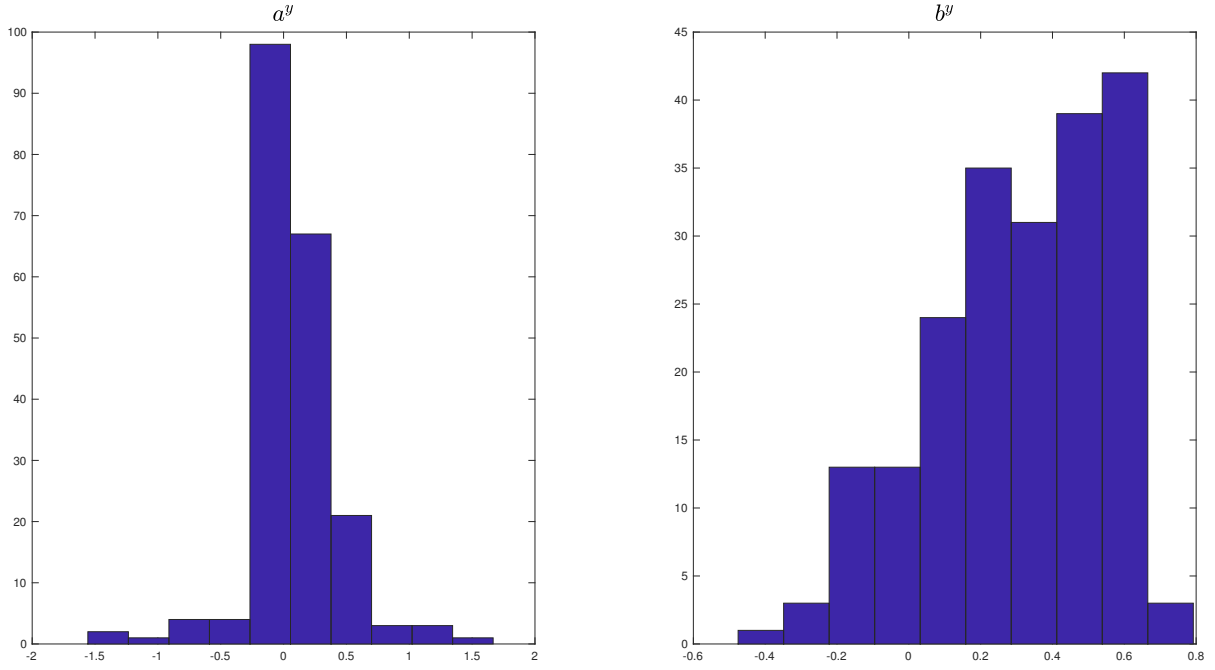


Figure 11: Histogram of Estimated Parameters Across Forecasters If Expected Output Sentiment Shock Evolves as an AR(1) Plus Constant

Inflation Sentiment as AR(1) Plus Constant: $s_t^{j,\pi} = a^\pi + b^\pi s_{t-1}^{j,\pi} + \varepsilon_t^{j,\pi}$

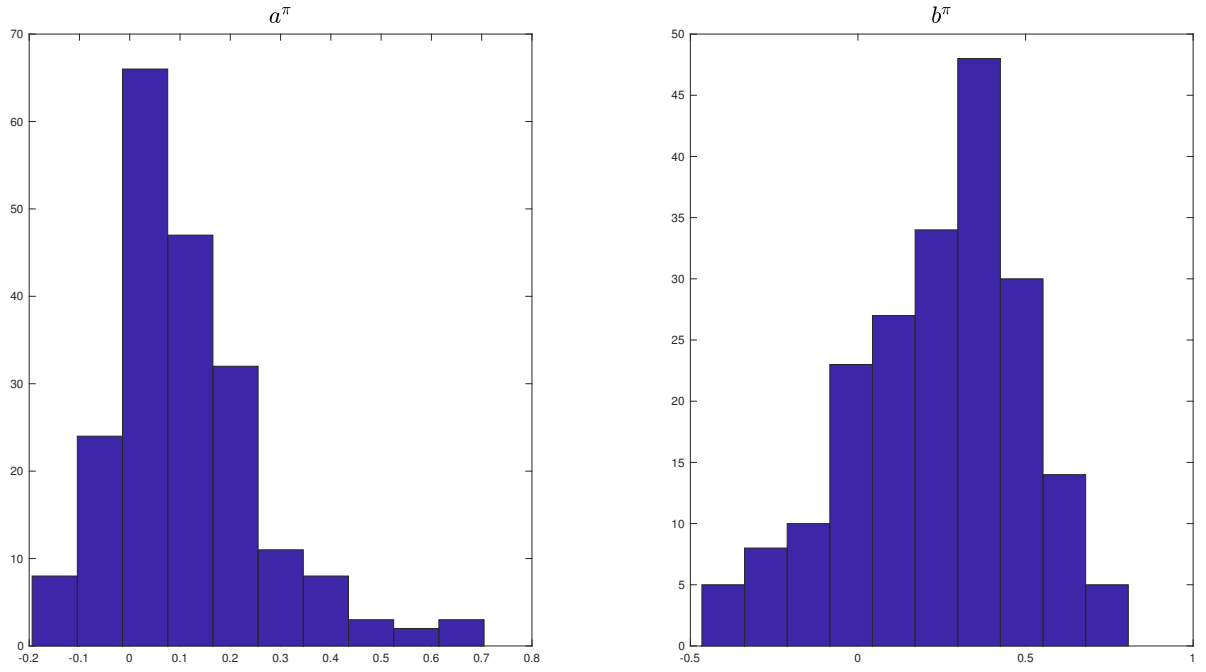


Figure 12: Histogram of Estimated Parameters Across Forecasters If Expected Inflation Sentiment Shock Evolves as an AR(1) Plus Constant

Sentiment as VAR(1) Plus Constant: $s_t^j = \Phi_0 + \Phi_1 s_{t-1}^j + \varepsilon_t^{j,\Phi}$

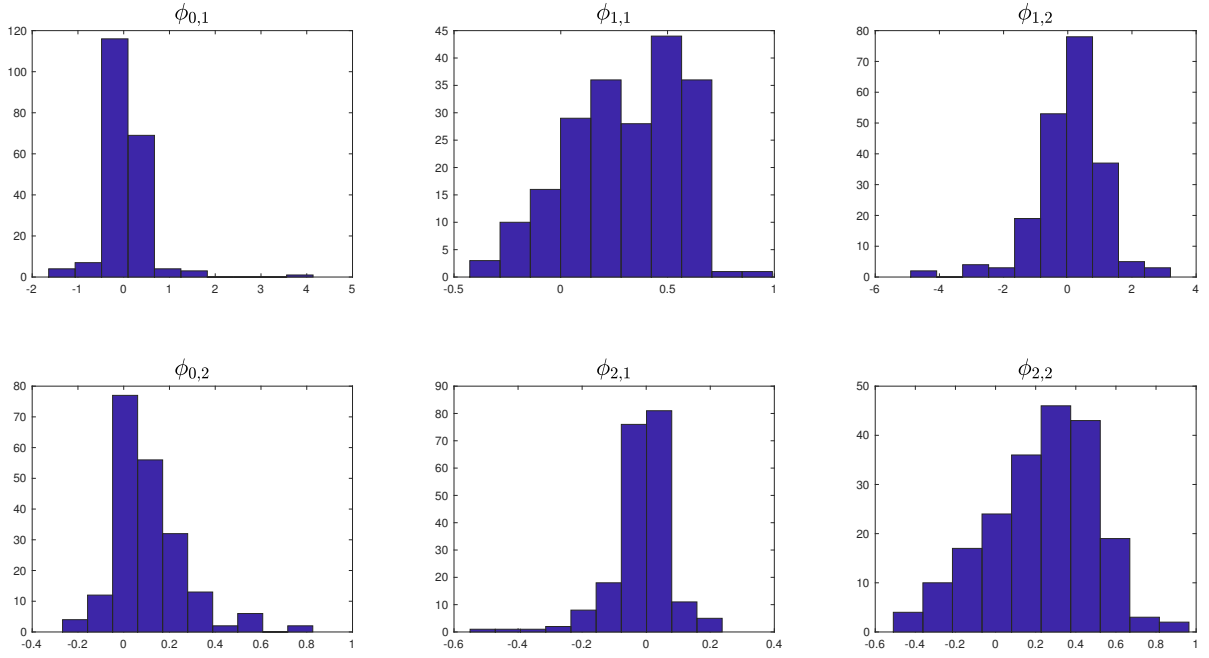


Figure 13: Histogram of Estimated Parameters Across Forecasters If Sentiment Shocks Evolve as a VAR(1) Plus Constant

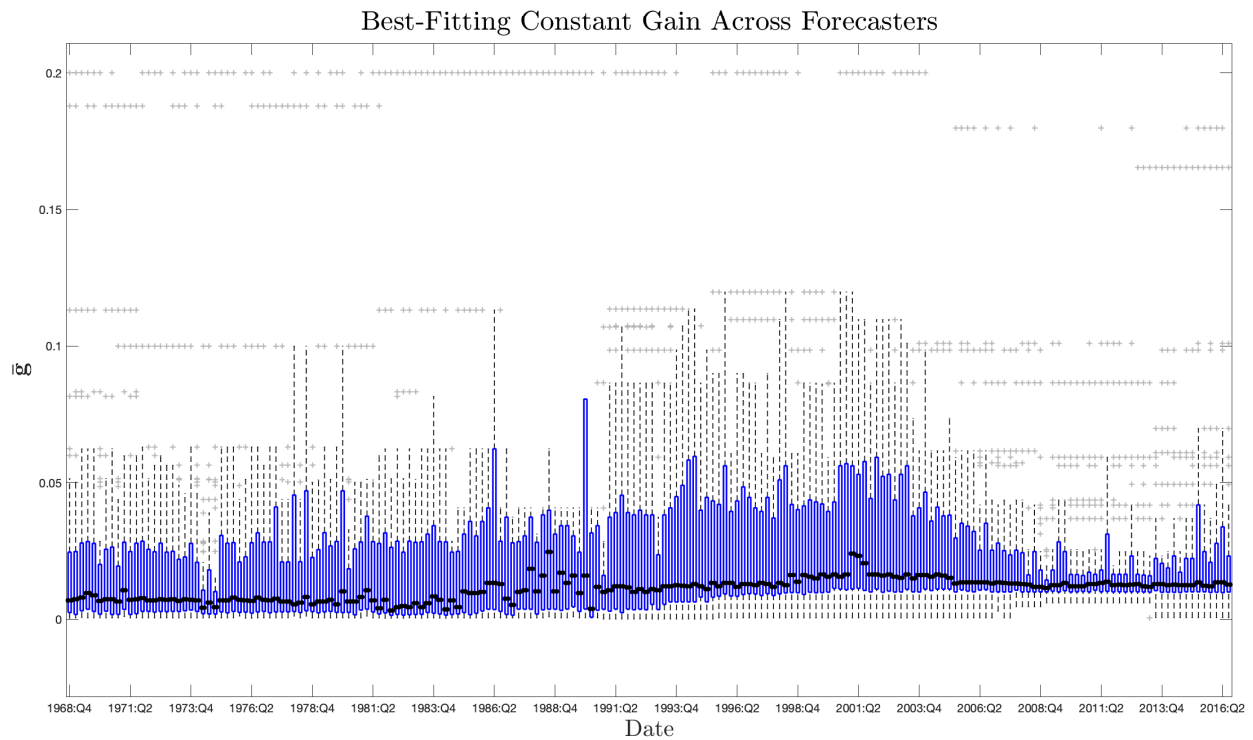


Figure 14: Distribution of the Best-Fitting Constant Gain of Individual Forecasters Each Time Period Under a VAR(1) Plus Constant PLM

Note: The distribution of the best-fitting constant gain of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

Best-Fitting Constant Gain Across Forecasters

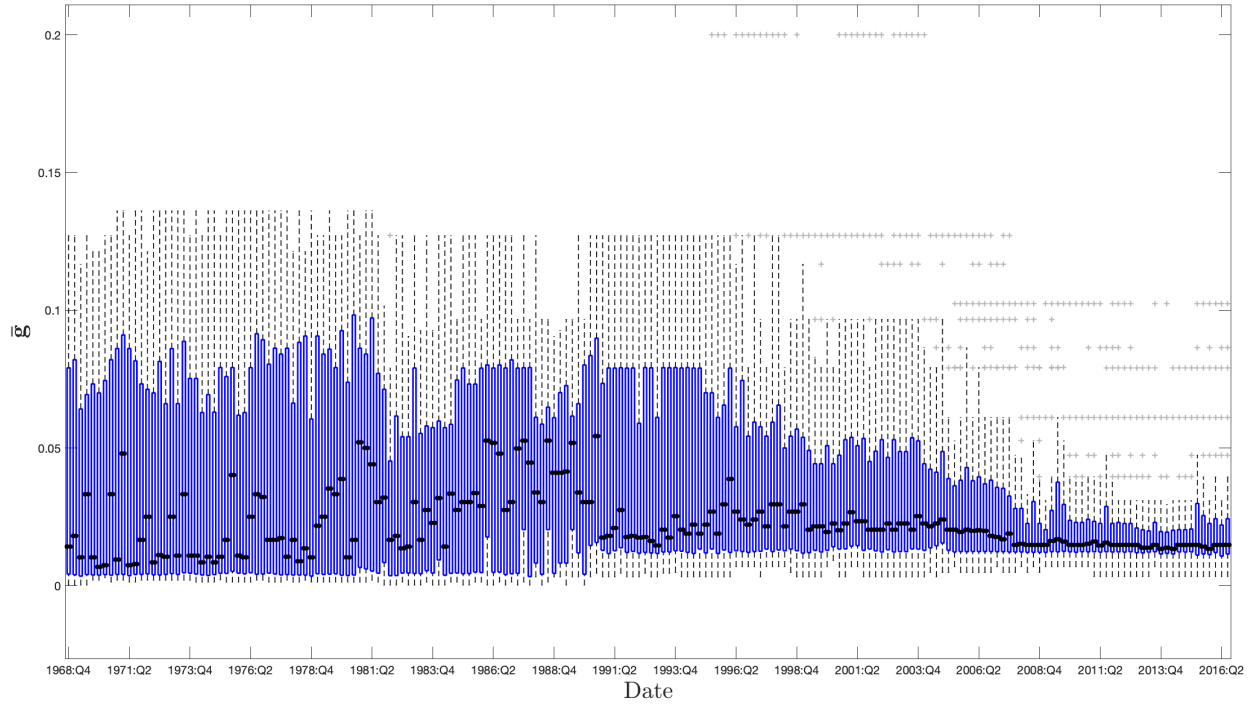


Figure 15: Distribution of the Best-Fitting Constant Gain of Individual Forecasters Each Time Period When SPF Respondents Submit Forecasts for at Least 20 Periods

Note: The distribution of the best-fitting constant gain of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.

Best-Fitting Constant Gain Across Forecasters

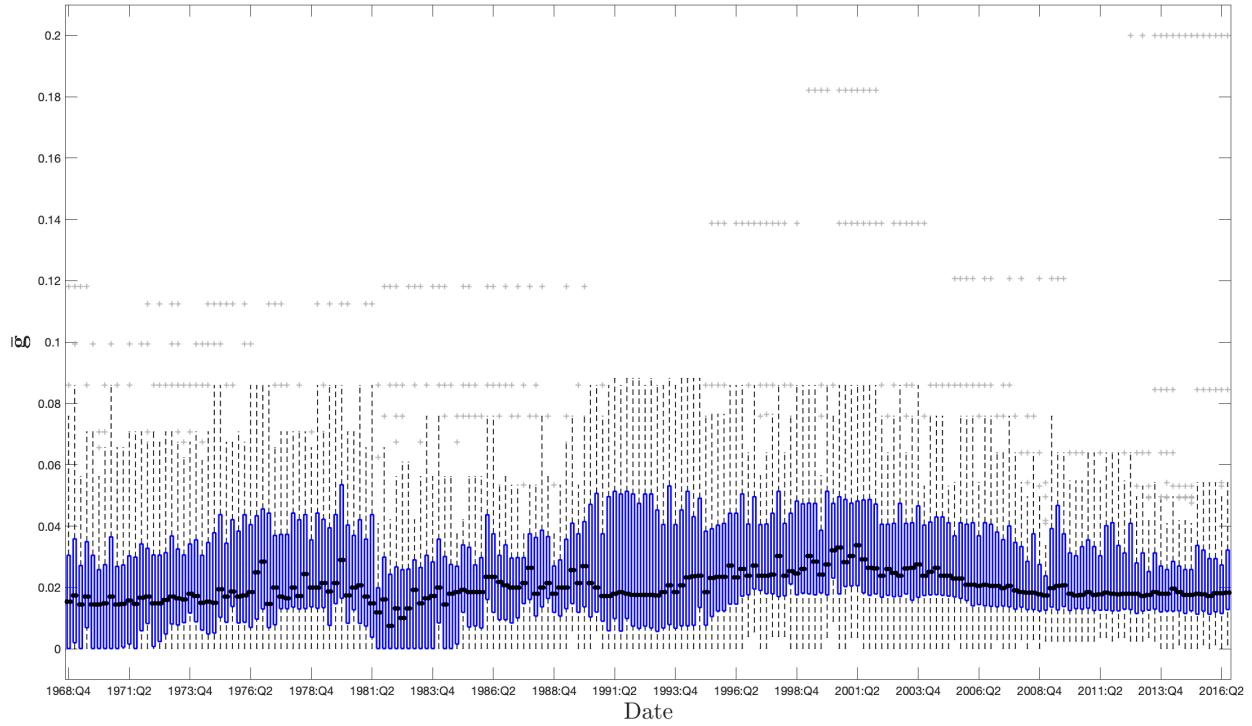


Figure 16: Distribution of the Best-Fitting Constant Gain of Individual Forecasters Each Time Period Under an Alternative Initialization of R_{τ}

Note: The distribution of the best-fitting constant gain of forecasters each time period is represented with a boxplot. The black circle signifies the median, the edges of the blue box the 25th and 75th percentiles, the black whiskers the extreme values not considered outliers, and the grey '+' symbol outliers.