DP17537
(v. 2)

HETEROGENEOUS BELIEFS AND THE PHILLIPS CURVE

Francesca Monti and Roland Meeks

MONETARY ECONOMICS AND FLUCTUATIONS

ISSN 0265-8003

HETEROGENEOUS BELIEFS AND THE PHILLIPS CURVE

Francesca Monti and Roland Meeks

Discussion Paper DP17537
First Published 26 September 2022
This Revision 29 September 2022

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

HETEROGENEOUS BELIEFS AND THE PHILLIPS CURVE

Abstract

Heterogeneous beliefs modify the New Keynesian Phillips curve by introducing a term in the cross-section distribution of expectations. We develop a novel functional data approach to estimation and inference in survey-based Phillips curves that accounts for variation in distributions of expectations, generalizing standard approaches. Our findings demonstrate the statistical and economic importance of heterogeneous beliefs for inflation dynamics, especially during periods of macroeconomic disruption. Our findings hold in similar form across two major economies.

JEL Classification: C4, C55, D84, E31

Keywords: Inflation dynamics, New keynesian phillips curve, Survey expectations, functional principal components, Functional regression

Francesca Monti - francesca.monti@uclouvain.be
CORE-UCLouvain and CEPR

Roland Meeks - rmeeks@imf.org
International Monetary Fund

Heterogeneous beliefs and the Phillips curve

Roland Meeks* and Francesca Monti†

September 9, 2022

Abstract

Heterogeneous beliefs modify the New Keynesian Phillips curve by introducing a term in the cross-section distribution of expectations. We develop a novel functional data approach to estimation and inference in survey-based Phillips curves that accounts for variation in distributions of expectations, generalizing standard approaches. Our findings demonstrate the statistical and economic importance of heterogeneous beliefs for inflation dynamics, especially during periods of macroeconomic disruption. Our findings hold in similar form across two major economies.

Keywords:
Inflation dynamics; New Keynesian Phillips curve; Survey expectations; Functional principal components; Functional regression

JEL: C4, C55, D84, E31

*Corresponding author. International Monetary Fund and CAMA. Email: `rmeeks@imf.org`
†King's College London, UC Louvain, and CEPR. Email: `francesca.monti@uclouvain.be`

1 Introduction

Expectations are widely seen as central to accounts of inflation dynamics. For this reason, surveys of inflation expectations have long enjoyed attention from economists, both for what can be learnt from them about the mechanisms underlying the formation of beliefs, and for the information that they may yield about future inflation itself. A prominent feature of the forecasts reported in surveys—whether of households, firms, or professional forecasters—is the considerable amount of cross-section dispersion (interpersonal heterogeneity) they display. This variation is often termed 'disagreement' following the early work of Mankiw, Reis, and Wolfers (2003). Observed survey disagreement frequently exceeds the variation in their time dimension. Yet curiously when it comes to linking expectations to inflation, the discussion amongst researchers and policymakers typically focuses precisely on time variation in the 'consensus' forecast, so collapsing the data to its cross-sectional average.

The present paper argues that a systematic account of heterogeneity in beliefs can benefit our understanding of the role of expectations in inflation dynamics. Consider the period of the Great Recession. During 2008 and 2009, the average year-ahead point forecast for US inflation reported in the Michigan Survey of Consumers never fell below 2%. But the distribution of responses, displayed in Fig. 1, indicates that during this episode beliefs were clustered in as many as three distinct modes, and that the highest of these was centered on zero inflation for over a year. Many respondents saw inflation running at 10% or more.[1] Similar patterns were observed in the UK. It is not immediately obvious that the average should be considered more reliable than the mode of these distributions, nor why beliefs in the tails of the distributions would be necessarily be considered uninformative.

Heterogeneity in beliefs has at least two important consequences for modelers. First, the average expectation may not adequately summarize the state of beliefs, leading to mismeasurement in econometric models. Unusual shifts in distributions of expectations around the consensus appear to be associated with periods of serious macroeconomic dislocation; as we detail later on, they emerged during the Volker disinflation in the United States, the United Kingdom's 1992 exchange rate crisis, and amid the recent COVID-19 pandemic and its aftermath. Such episodes are precisely when econometric models tend to break down. And second, heterogeneous beliefs undermine the micro-foundations of the widely-used survey-based New Keynesian Phillips curve (Mavroeidis et al., 2014, p. 135). The underlying problem is that, as is widely known, the law of iterated expectations (LIE) is not generally preserved when beliefs are averaged, even if the condition is satisfied by the beliefs of each individual agent.[2] This

[1] The prevalence of forecasts of 0%, 5% and 10% inflation are not artifacts, but a result of known biases towards reporting round numbers when respondents become uncertain (Binder, 2017). Underlying the distributions are thousands of individual observations per quarter, making it unlikely that such features would be 'averaged away' in ever-larger samples.

[2] Imposing the LIE on the cross-section of beliefs allows inflation (which is the change in an aggregate price index) to be written as an expectational difference equation involving only aggregate quantities. This entails some rather strong and therefore unappealing assumptions on agents' beliefs (Adam and Padula, 2010).

Figure 1. Cross-section distributions of inflation forecasts during the financial crisis and its aftermath

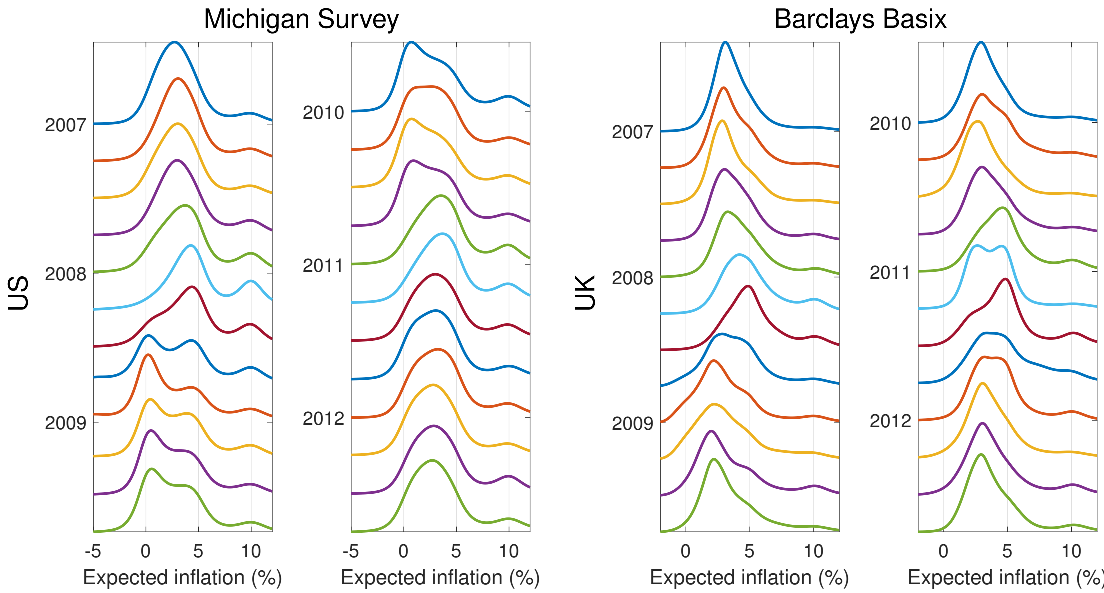

Note: Panels show time series of distributions of individual survey respondents' year-ahead point forecasts from each of the named surveys. Dates reflect when forecasts were made. Details about the surveys may be found in Section 4.1. Details of the density estimation method may be found in Part II of the Supplementary Material.

aggregation problem leads an additional term to appear in the Phillips curve that depends on the cross-section dispersion of beliefs. The empirical relevance of this theoretical result remains to be explored.

The principal contribution of our paper is to offer a new approach to estimation and inference for forward-looking models in environments where beliefs are heterogeneous in unspecified and general ways. This is likely to be the situation faced by most researchers using survey data on expectations in applied macroeconomics. We address the problem of how to formulate an econometric model that relates the scalar time series for inflation to the high-dimensional array of survey responses present in each period. We argue that progress can be made by considering finite but potentially large cross-sections of survey forecasts as originating from continuous distributions, in which there is an underlying close, smooth, or 'functional' relationship between adjacent forecast values. Recognizing that the data can be regarded as a time series of stochastic continuous distributions, or functional time series, allows the techniques of functional data analysis (FDA) to be brought to bear (Ramsay and Silverman, 2005; Horváth and Kokoszka, 2012). Following Kneip and Utikal (2001), we apply functional principal component analysis to time series of distributions, and establish a set of stylized facts that extend what is known from the literature on forecast disagreement. We find that a handful of factors can jointly characterize much of the variation present in the cross-section distributions of beliefs. These functional factors may then be used to estimate the heterogeneous beliefs model using principal component regression (Reiss and Ogden, 2007), thus transforming a scalar-on-function regression problem

Figure 2. Estimated contributions of short- and long-run expectations to recent changes in US inflation

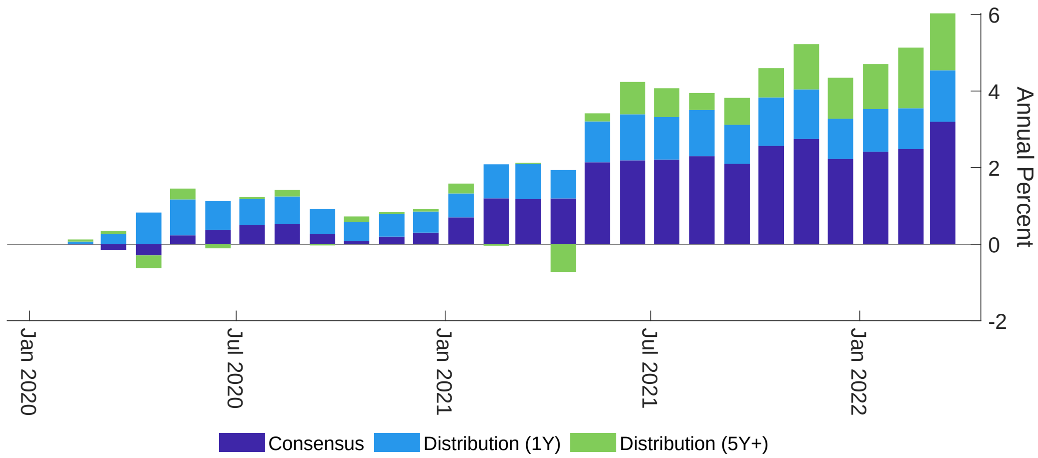

Note: Bars indicate the change relative to January 2020 in the estimated contributions to the month-on-month percentage change in the seasonally-adjusted CPI index (expressed at annual rates) from: (i) Consensus year-ahead expectations; (ii) The distributions of year-ahead expectations; and (iii) The distributions of 5-10 year ahead expectations (all taken from the Michigan Survey of Consumers). For details of the underlying model, refer to Section 7.

to multiple regression analysis. Tests of the standard 'consensus beliefs' model, which is nested in our specification, are readily available (Kong et al., 2016).

The payoff to our new approach is the discovery of a richer and more nuanced role for expectations in the inflation process than had previously been recognized. We estimate the expectations-augmented Phillips curve using complete sets of household inflation forecasts reported in the US Michigan survey, and in a newly-collated UK household survey. We show that in both regions, signals about future inflation contained in the distribution of beliefs—in addition to the consensus expectation—affect current inflation. The effects can be quantitatively substantial. Fig. 2 gives a flavor of our results. It shows the contributions of three components of household expectations to US inflation since the onset of the COVID-19 pandemic estimated using one of the models described in the main body of the paper. The contributions are normalized to zero in January, 2020. The darker bars, which start to rise in early 2021, indicate the increased contribution of consensus beliefs about near-term inflation; their rise contributes 3 percentage points to inflation by March, 2022. The remaining bars show the estimated additional effects from heterogeneous beliefs. As we set out later in the paper, by mid-2021 the distribution of households' beliefs, especially about longer-run inflation, had become skewed to the upside. Seen through the lens of the heterogeneous beliefs model, this shift had a separate and distinct impact on inflation that was of roughly equal magnitude to that of the consensus.

Our focus on expectations in this paper complements, to some degree, the extensive debates on the role of economic slack and other supply-side factors that may drive inflation. One prominent line of argument holds that the Phillips curve has 'flattened' in recent years—meaning that the effect of a tight labor market on inflation is more muted than was the case in the 1980s, for example Del Negro et al. (2020). If that is the case, the need to scrutinize the effects of expectations on inflation is correspondingly greater. Although we motivate our approach by appealing to a model that embeds standard sticky-price reasoning, we do not exclude other possible explanations for our results. And while our approach has the benefit of flexibility, and is able to detect general departures from the consensus beliefs model, we must also acknowledge that its lack of specificity may also be seen as something of a drawback. The semi-parametric flavor of our approach allows us to quantify the effects of distributional shifts, but does not easily lend itself to a particular theoretical interpretation.[3] As in related studies, we treat expectations as primitive objects, which of course they are not.

Related literature.—Our paper lies firmly within the extensive body of work makes use of inflation survey data to improve estimates of the Phillips curve and to aid in inflation forecasting: see amongst many others Roberts (1995), Ang et al. (2007), Brissimis and Magginas (2008), Faust and Wright (2013), Mavroeidis et al. (2014), Binder (2015), Coibion and Gorodnichenko (2015), Coibion et al. (2018), and Coibion et al. (2019). A central theme of this work is that data on expectations helps to resolve some otherwise puzzling shortcomings of the New Keynesian Phillips curve, and to circumvent the econometric traps that can plague estimation under the assumption of rational expectations.[4] However, the present paper is the first to discuss the role played by non-consensus beliefs on inflation.

The paper also connects to work that seeks to characterize the behavior of cross-sectional forecast disagreement (Mankiw et al., 2003). That literature finds that disagreement can persist even in the long-run (Patton and Timmermann, 2010; Andrade et al., 2016), and that it is present in surveys of sophisticated agents such as professional forecasters (Andrade and Le Bihan, 2013). Although a few prior studies have looked beyond disagreement to consider distributions of forecasts (Filardo and Genberg, 2010; Pfajfar and Santoro, 2010), none that we know of have analyzed the factor structure of those distributions as we do here.

Last, the approach we set out further provides a novel application of FDA to a problem in macroeconomics. Previous applications of FDA in econometrics include the work on yield curve forecasting in Bowsher and Meeks (2008), the model of relative price dispersion and inflation in Chaudhuri et al. (2016), and the investigation of cross-market dependence in stock returns in Park and Qian (2012). Tools for exploratory data analysis are presented in Tsay (2016).

Outline.—The remainder of this paper is organized as follows. Section 2 recalls the microfoundations of the Phillips curve under heterogeneous beliefs, with Section 3 setting out our econometric approach to estimating the model. Section 4 provides an analysis of the structure in distributions of survey forecasts. Section 5 contains our headline results in support of

[3] In the earlier version of this paper, we characterize the underlying distributional factors in terms of their statistical properties in greater detail than here (Meeks and Monti, 2019). But in the absence of a theory of belief formation, the economic interpretation of the factors is problematic. We are grateful to Refet Gürkaynak for highlighting this point.

[4] The referenced econometric problems relate mainly to the weakness of the available instruments for future inflation in GMM estimation approaches. The advantages that the use of survey data bring compared to alternative estimation procedures led Mavroeidis et al. (2014, p. 151) to describe the survey approach as having gained a 'commanding' presence in the Phillips curve literature.

the heterogeneous beliefs model, with separate treatment of the US and UK Phillips curves. The robustness of our results to lagged inflation, trend inflation, and an alternative estimation approach appear in Section 6. An analysis of the behavior of inflation during and after the COVID-19 pandemic from the heterogeneous beliefs viewpoint is presented in Section 7. Finally, Section 8 offers concluding comments.

2 Microfoundations with heterogeneous beliefs

A Phillips curve derived from firms' optimizing behavior lies at the core of widely-used New Keynesian macroeconomic models. In this section, we recall the structure of the Phillips curve that is consistent with general forms of heterogeneous beliefs, focusing on the case where there are differences of opinion between agents in the spirit of Harris and Raviv (1993). We take such differences of opinion as given.

2.1 The Phillips curve

Suppose there are N firms operating under monopolistic competition seeking to maximize their profits under time-dependent pricing. For an individual firm j, familiar calculations lead to a log-linear expression for its optimal reset price, $p^{\star(j)}$, in terms of current and expected future nominal marginal costs (log real marginal cost φ plus the log of the aggregate price level p):

$$p_t^{\star(j)} = (1-\beta\theta)\mathbb{E}_t^{(j)}\left[\textstyle\sum_{\tau=t}^{\infty}(\beta\theta)^{\tau-t}(\varphi_\tau + p_\tau)\right] \tag{1}$$

where β is the discount factor applied to future profits and θ is the constant per-period probability that a firm's nominal price remains fixed. Expectations are particular to each firm, as indicated by the j superscript on the expectations operator.

To express the infinite sum on the right of Eq. (1) in the form of a difference equation, it is sufficient to assume the law of iterated expectations applies to the expectation operators $\mathbb{E}_t^{(j)}[\cdot]$ used by each firm j.[5] Applying this assumption, then subtracting the current aggregate price level from both sides gives an expression for the individual firm's optimal real, or relative, price:

$$q_t^{\star(j)} = (1-\beta\theta)\varphi_t + \beta\theta\mathbb{E}_t^{(j)}\left[q_{t+1}^{\star(j)} + \pi_{t+1}\right] \tag{2}$$

where $q_t^{\star(j)} := p_t^{\star(j)} - p_t$, and $\pi_{t+1} := p_{t+1} - p_t$ is the rate of price inflation.

To obtain an aggregate relationship, sum Eq. (2) over the cross-section of firms. Using an obvious notation for the average over N units:

$$q_t = (1-\beta\theta)\varphi_t + \beta\theta\overline{\mathbb{E}}_t[\pi_{t+1} + q_{t+1}] + \beta\theta\mathsf{E}_N\left[\mathbb{E}_t^{(j)}\left(q_{t+1}^{\star(j)} - q_{t+1}\right)\right] \tag{3}$$

noting that the future average relative price $q_{t+1} := \mathsf{E}_N\left\{q_{t+1}^{\star(j)}\right\}$ has been added and subtracted on

[5] Although not innocuous, this assumption retains tractability without requiring stronger assumptions on the dependence (or lack thereof) of beliefs in the cross-section (see Coibion et al., 2018, p. 1466-7; Branch and McGough, 2009, Section 2.1).

the right hand side, and use of the notation:

$$\overline{\mathbb{E}}_t(x) := \mathsf{E}_N\left[\mathbb{E}_t^{(j)}(x)\right]$$

for the average expectation of variables x that do not depend on j ('aggregates').

We can now obtain an expression for the *heterogeneous beliefs* Phillips curve. Under the sticky price assumption, the aggregate price level is a combination of past and current prices and inflation is proportional to the average relative reset price:

$$\pi_t = \frac{(1-\theta)}{\theta} q_t$$

with the coefficient of proportionality depending on the frequency of price resets. Substituting into Eq. (3), we find:

$$\pi_t = \kappa\varphi_t + \beta\overline{\mathbb{E}}_t(\pi_{t+1}) + \beta(1-\theta)\Delta_t \tag{4}$$

where

$$\Delta_t := \mathsf{E}_N\left[\mathbb{E}_t^{(j)} q_{t+1}^{\star(j)} - \mathbb{E}_t^{(j)} q_{t+1}\right] \tag{5}$$

and the slope of the Phillips curve $\kappa := (1-\beta\theta)(1-\theta)/\theta$. There are two points of departure from the conventional New Keynesian Phillips curve: First, an average expectation replaces the conventional rational expectation; And second, the term Δ_t representing differences of opinion averages over the gaps between each agent's forecast of next period's optimal reset price and their forecast of the 'consensus'.

2.2 A two-type example

To better understand the conditions under which differences of opinion might matter for aggregate outcomes, consider an example featuring two types of agent, who differ only in their beliefs. Setting $N = 2$ in Eq. (5) and noting that q_t is given by the average of the reset prices $q_t^{\star(1)}$ and $q_t^{\star(2)}$, we find that:

$$\Delta_t = \frac{1}{2}\left[\mathbb{E}_t^{(1)}\left(q_{t+1}^{\star(1)} - q_{t+1}^{\star(2)}\right) + \mathbb{E}_t^{(2)}\left(q_{t+1}^{\star(2)} - q_{t+1}^{\star(1)}\right)\right]$$

Differences of opinion manifest as gaps between the expectations that type 1 agents hold about their own reset prices and those that they hold about type 2 agents' reset prices (and similarly for type 2 agents). In order for $\Delta_t \neq 0$, the different beliefs of type 1 and type 2 agents must not simply cancel out. Notice that full rational expectations is a sufficient but not a necessary condition for $\Delta_t = 0$, since we can have $\mathbb{E}_t^{(j)} q_{t+1}^{\star(j)} \neq \mathbb{E}_t^{(i)} q_{t+1}^{\star(j)}$, violating rational expectations, but still find $\Delta_t = 0$.

To formalize differences in beliefs, proceed as follows. Let there be a single endogenous state variable, φ_t, which follows a stationary AR(1) process driven by a white noise shock ν. Agents are capable of computing mathematical expectations, so the LIE holds for the beliefs each agent type. Both agent types know the φ_t process, and correctly surmise that a linear function of the state variables $\{\varphi_t, \nu_t\}$ appropriately captures the decision rule that should be followed when selecting their optimal reset price. But agents entertain different models of the economy: they differ in their beliefs about these decision rules. (We do not take a stand on why they differ.) Type j's expectation of type i's reset price is given by:

$$\mathbb{E}^{(j)} q_t^{\star(i)} = b_{ij}\, \varphi_t \tag{6}$$

where the coefficients b_{ij} differ by agent type. To keep matters simple, suppose agents agree on how their own reset prices should be determined, such that $b_{11} = b_{22} = b$.[6] However, agents disagree concerning the other type's decision rule. Type 1 agents hold that when type 2 agents choose their reset price they use a decision rule that is different from their own, in which case $b_{21} \neq b$ (and similarly $b_{12} \neq b$). So the differences of beliefs term Eq. (5) is non-zero as:

$$\Delta_t = \frac{1}{2}\left[\mathbb{E}_t^{(1)}\left(q_{t+1}^{\star(1)} - q_{t+1}^{\star(2)}\right) + \mathbb{E}_t^{(2)}\left(q_{t+1}^{\star(2)} - q_{t+1}^{\star(1)}\right)\right] = \left[b - \frac{1}{2}\left(b_{21} + b_{12}\right)\right]\varphi_t$$

If the beliefs of type 1 agents about the decision rule of type 2 agents did coincide exactly with the rule that they themselves apply, then $b_{21} = b$ (and similarly that $b_{12} = b$), it would follow immediately that $\Delta_t = 0$, and also that the LIE would hold 'in the cross-section'—that is, $\mathbb{E}_t^{(1)}[\mathbb{E}_t^{(2)} q_{t+1}^{\star(1)}] = \mathbb{E}_t^{(2)} q_{t+1}^{\star(1)}$. One can straightforwardly show that $\Delta_t = 0$ if and only if LIE holds in the cross section (see also Adam and Padula, 2010, Condition 1).

3 Estimation and testing with heterogeneous beliefs

Having established the potential theoretical relevance of heterogeneous beliefs for the Phillips curve, we now consider the problem of estimating and testing the model against the canonical 'consensus beliefs' alternative.

3.1 Approximating the model

If surveys elicited information on individual firms' price plans along with their expectations about general inflation, it would be possible to estimate Eq. (4) directly. However, such data has not generally been collected over sufficiently long periods for time series analysis. To make progress, we must therefore make an approximating assumption based on the information most commonly to hand, namely expectations of general inflation.

In absence of any particular guidance from theory, we assume that the expectational gap in Eq. (5) can be approximated, via some unknown function γ, by the gap between the cross-section of the individual expectations of aggregate inflation and the consensus. Adopting the shorthand notation π^e and $\overline{\pi}^e$ for time-t individual and consensus expectations about the rate of aggregate inflation respectively, we may therefore write:

$$\Delta_t \approx \mathsf{E}_N\left[\gamma\left(\pi_t^e - \overline{\pi}_t^e\right)\right]$$

[6] The coefficients the agents adopt in their own decision rules may, but need not, be those that would apply under rational expectations.

or in the limit as N becomes large:

$$\lim_{N} \Delta_t \approx \int \gamma(\pi_t^e - \overline{\pi}_t^e) \mathrm{d}\mathsf{P}_t^{\mathsf{c}}(\pi^e) \tag{7}$$

with P^{c} denoting the distribution of beliefs around the consensus (see also Section 3.2).

We sought to validate the approximation in Eq. (7) with a Monte Carlo exercise (reported in Part VII of the Supplementary Material). In brief, we simulate data from a model that generates plausible fluctuations in individual beliefs and test how well the Δ_t term can be proxied by the functional principal components of the distribution of beliefs about aggregate inflation. To do so, we regress Δ_t on the functional principal components and find that the set of functional principal components is strongly associated with it. Therefore, in the simulation we conducted, the functional principal components of the cross-sectional distributions of inflation expectations carry relevant information on the unobserved term Δ_t.

With our approximation in place, the estimable version of the heterogeneous beliefs model Eq. (4) is given by:

$$\pi_t = \alpha\varphi_t + \beta\overline{\pi}_{t,h}^e + \int \boldsymbol{\gamma}(\pi_t^e - \overline{\pi}_t^e) \mathrm{d}\mathsf{P}_t^{\mathsf{c}}(\pi^e) + \varepsilon_t \tag{8}$$

This equation is an example of a functional linear model or FLM (Ramsay and Silverman, 2005, Ch. 15), wherein a scalar quantity (inflation) is related to a functional covariate (the distribution of inflation expectations). The effect of the coefficient function is to place weight on parts of the distribution of expectations that associate strongly with inflation, and to down-weight parts that do not. Put differently, when $|\boldsymbol{\gamma}|$ is large in some range of values for π^e, expectations in that region of the distribution have greater influence on inflation.[7]

3.2 Working with functional data

As a prelude to our discussion concerning the estimation of FLMs, this section briefly sets out the necessary functional data concepts. Amongst the most fundamental of these is that of the average distributional shape. To find an interpretable average, we first align the distribution functions around a common feature (a process is known as 'registration', see Ramsay and Silverman, 2005, Ch. 7). A natural choice is the consensus forecast, so we center (i.e. horizontally translate) each distribution by subtracting from the h-step ahead inflation forecasts π_h^e made in each period the mean of their distribution $\mathsf{p}_{t,h}$. The sample average distribution shape, or functional mean, of h-step ahead point forecasts is then given by:[8]

$$\overline{\mathsf{p}}_h(x) = \frac{1}{T}\sum_{t=1}^{T} \mathsf{p}_{t,h}^{\mathsf{c}}(x) \tag{9}$$

[7]One way to think about the resulting estimates $(\beta, \boldsymbol{\gamma})$ is that they provide the means to construct a single index of expectations. That index is the transformation of the survey data that is most the closely related to actual inflation. The commonly-pursued alternative is to first construct an expectations index (to assume a $\boldsymbol{\gamma}$), and then to ascertain its relationship with inflation.

[8]The expectation of a random function $\mathsf{p}(x)$ is defined as the ordinary expectation taken pointwise for $x \in [a, b]$. For discussion on the concept of functional expectation, see Cuevas (2014, Section 3.1).

where $\mathsf{p}^{\mathsf{c}}_{t,h}$ represents the density function over centered forecasts.

An alternative robust measure of central tendency is the functional median. The median is taken to be the function with maximal band depth, as in López-Pintado and Romo (2009). Given an empirical distribution of functional objects $\mathbb{P}_T$ and a particular function p, depth is a function $D(\mathbb{P}_T, \mathsf{p}) \geq 0$ indicating how far 'inside' that distribution p lies. A measure of depth therefore provides an ordering of the data, with the usual notion of the median being the function that lies the 'deepest' within the set.[9]

Functional principal component analysis (FPCA) is a standard technique for summarizing general functional variation, and was applied to probability density functions by Kneip and Utikal (2001). FPCA is closely analagous to classical PCA as applied to empirical covariance matrices (for an even-paced introduction, see Ramsay and Silverman, 2005, Ch. 8). The representation of a function in terms of its principal component functions (synonymously 'eigenfunctions') is known as the Karhunen-Loève expansion. It tells us that the functional data $\{\mathsf{p}^{\mathsf{c}}_{t,h}\}_0^T$ can be expressed in terms of an expansion in the orthonormal eigenfunction (or empirical) basis $\{\mathsf{e}_k\}$ as $\mathsf{p}^{\mathsf{c}}_{t,h} = \mu_{\mathsf{p}} + \sum_{k=1}^{\infty} \langle \mathsf{p}^{\mathsf{c}}_{t,h}, \mathsf{e}_k \rangle \mathsf{e}_k$. The principal component functions form an optimal basis for the observations to hand. Optimality in this context means that, for a given K, the linear approximation $\hat{\mathsf{p}}^{(K)}_{h,t}$ minimizes the integrated squared error criterion:

$$\mathsf{ISE}^{(K)}_{t,h} = \int \left\{ \left(\hat{\mathsf{p}}^{(K)}_{t,h} - \overline{\mathsf{p}}_h \right) - \left(\mathsf{p}^{\mathsf{c}}_{t,h} - \overline{\mathsf{p}}_h \right) \right\}^2 \mathrm{d}x, \quad \text{where} \quad \hat{\mathsf{p}}^{(K)}_{t,h} = \overline{\mathsf{p}}_h + \sum_{k=1}^{K} s_{kt} \mathsf{e}_k \tag{10}$$

averaged over all t, subject to the constraint that the functions $\mathsf{e}(\cdot)$ satisfy $\langle \mathsf{e}_k, \mathsf{e}_k \rangle = 1$ and $\langle \mathsf{e}_k, \mathsf{e}_j \rangle = 0$, $k \neq j$ where $\langle \cdot, \cdot \rangle$ denotes the usual inner product for square-integrable functions. The principal component scores, which play a central role in estimation, are given by $s_{kt} = \langle \mathsf{p}_t, \mathsf{e}_k \rangle$. Although exact solutions to the principal component problem are not generally available, computational approximations are, the details of which are summarized in Appendix B (see also Tsay, 2016, Section 3.3).

3.3 Estimation

A variety of approaches have been proposed for the estimation of functional linear models. As foreshadowed in the preceding section, we adopt the popular functional principal component regression approach. It recasts the functional regression Eq. (8) as a familiar multiple regression problem (see Reiss et al., 2017).[10] To understand the approach, notice that expanding the functional coefficient γ in the same basis $\{\mathsf{e}_k\}$ as the functional data allows us to write $\gamma = \sum_{k=1}^{\infty} \langle \gamma, \mathsf{e}_k \rangle \mathsf{e}_k$. Then using the properties of the e_k, see Eq. (A.1), the functional linear model

[9]The concept of band depth is based on the graph of a function on the plane. A band can be thought of as the envelope delimited by n such graphs. The band depth of a given curve p_0 is given by the proportion of times that it falls inside the bands formed by taking all possible combinations of n curves. For example, if $n = 2$ and $T = 10$, there would be 45 pairs of curves (bands), and if the graph of p_0 lay entirely inside 9 of those bands its depth would be 0.2. See Cuevas (2014, Section 4.3) for further discussion.

[10]Additional details, along with references to the literature, are given in Appendix A.

of Eq. (8) can be rewritten as:

$$\pi_t = \alpha\varphi_t + \beta\overline{\pi}^e_{t,h} + \sum_{k=1}^{K} \gamma_k s_{k,t} + \varepsilon_t \tag{11}$$

where the γ_k are scalar coefficients to be estimated, and the functional principal component scores $s_{k,t}$ obtained from the FPCA appear as covariates.

Having recast the functional linear model Eq. (8) as the multiple regression model Eq. (11), estimation proceeds as follows. Denote the $(T \times 1)$ vector formed by stacking the dependent variable by $\boldsymbol{\pi}$, and the $(T \times K)$ matrix of orthogonal principal component scores $s_{k,t}$ by $\mathbf{M}$. The N additional (scalar) regressors, including a vector of mean expectations, are collected in the $(T \times N)$ matrix $\mathbf{Z}$. Then conditional on the truncation level K and the true principal component scores, the heterogeneous beliefs Phillips curve model Eq. (8) is written compactly as:

$$\boldsymbol{\pi} = \mathbf{M}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\beta} + \boldsymbol{\varepsilon}, \qquad \boldsymbol{\varepsilon} \sim \mathsf{N}(0, \sigma^2\mathbf{I})$$

where with a slight abuse of notation $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_K)^\top \in \mathbb{R}^K$ collects the scalar coefficients on the scores. Let $\mathbf{X} = [\mathbf{Z}, \mathbf{M}]$ be the $T \times (N + K)$ matrix of regressors, and define the projection matrices $\mathbf{P}_X = \mathbf{X}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top$ and $\mathbf{P}_Z = \mathbf{Z}(\mathbf{Z}^\top\mathbf{Z})^{-1}\mathbf{Z}^\top$. Then the maximum likelihood estimator of the coefficients on the functional principal component scores is:

$$\hat{\boldsymbol{\gamma}} = \mathbf{Q}^{-1}\mathbf{M}^\top(\mathbf{I} - \mathbf{P}_Z)\boldsymbol{\pi} \tag{12}$$

where $\mathbf{Q} \coloneqq (\boldsymbol{\Lambda} - \mathbf{M}^\top\mathbf{P}_Z\mathbf{M})$ is the Schur complement of $(\mathbf{Z}^\top\mathbf{Z})$ in $(\mathbf{X}^\top\mathbf{X})$, and $\boldsymbol{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_K)$ contains the first K size-ordered eigenvalues corresponding to the scores arrayed in the columns of $\mathbf{M}$.

3.4 Testing

To establish whether an association exists between current inflation and the distribution of inflation forecasts, we employ the classical testing procedure of Kong et al. (2016). A natural null hypothesis is that $\boldsymbol{\gamma}(\pi^e) = 0$, which recalling that the distributions p^{c} are mean zero by construction, corresponds to the special case where only the average forecast matters for inflation, or equivalently that a functional effect on inflation is absent. A test of the hypothesis $H_0 : \boldsymbol{\gamma}(\pi^e) = 0$ for all π^e is equivalent to:

$$H_0 : \gamma_1 = \gamma_2 = \cdots = \gamma_K = 0 \quad \text{vs.} \quad H_a : \gamma_j \neq 0 \quad \text{for at least one } j,\ 1 \leq j \leq K$$

Then H_0 can be tested using the F-statistic:

$$T_F = \frac{\boldsymbol{\pi}^\top(\mathbf{P}_X - \mathbf{P}_Z)\boldsymbol{\pi}/K}{\boldsymbol{\pi}^\top(\mathbf{I} - \mathbf{P}_X)\boldsymbol{\pi}/(T - K - N)} \overset{approx.}{\sim} F_{K,T-K-N} \tag{13}$$

where $F_{K,T-K-N}$ denotes the F distribution with degrees of freedom depending on the number of functional principal components K and the number of scalar regressors N (Kong et al., Theorem 3.1).

3.5 Model selection

An outstanding question is how to select the truncation level K of the Karhunen-Loève expansion in Eq. (11). Adding more principal components naturally leads to smaller errors of functional approximation (in the sense of ISE) in every time period. One simple approach to model selection is then to select only those components for which the cumulative share of variance (in the functional explanatory variable) is below some threshold value, often set at 95% or 99%. But a low variance share for a particular component does not necessarily imply that it is unimportant in the regression model (see the discussion in Jolliffe, 2002, Section 8.2).[11] Conversely, including a large number of FPC scores in the regression risks over-fitting the data. In the subsequent analysis, we therefore consider alternative values of K, based either on the simple cumulative eigenvalue criterion, or an information criterion, which takes account of both fit and parameterization.

4 Survey expectations data

Our analysis will make use of household survey data. Long time series of household data are available for both the US and the UK, which unfortunately is not the case with surveys of firm expectations, which are presumably of greatest relevance to pricing decisions. Household beliefs are considered to be a better proxy for those of firms than are the beliefs of professional forecasters (Coibion and Gorodnichenko, 2012).

4.1 Data sources

The expectations survey data we use for the US is taken from the Michigan Survey of Consumer Attitudes (MSC), and for the UK from the Barclays survey of inflation expectations (Basix). To the best of our knowledge, we are the first to make research use of the full Basix data set. The two surveys ask similar questions about 'prices in general' or 'inflation', without specifying a particular measure. Each asks respondents to report their expectation for inflation over the following year, and their expectations for at least one other horizon. Interviews for the Michigan Survey take place monthly, largely during the month in question (and therefore before that month's CPI release). The quarterly Basix interviews usually take place late in the second month to early in the third month of each quarter. A summary of the main features survey data used in this study is given in Tab. I.1 of the Supplementary Material.

[11] Kneip and Utikal (2001) develop asymptotic inference for selecting principal components of density functions, and Tsay (2016) proposes a cross-validation procedure based on the Hellinger distance. Faraway states in his comment on Kneip and Utikal (2001) that: "In other situations, selection of dimension [the number of components] is a secondary consideration to some [primary] purpose—typically prediction. The dimension should be chosen to obtain good predictions ... It is important to optimize the secondary selection with respect to the primary objective and not some criterion associated with the secondary objective". His arguments motivate our use of information criterion.

4.2 Estimating distributions

The first step in our analysis to transform the discrete cross-section of point expectations reported by survey respondents into continuous distribution functions.[12] In each survey quarter, we adopt a nonparametric technique to obtain consistent estimates of that distribution.[13] The notation $\mathsf{p}_{t,h}(\cdot)$ will denote the distribution of h-step ahead point forecasts made at date t. The sequence $\{\mathsf{p}_{t,h}(\cdot)\}_0^T$ is then a functional time series (Bowsher and Meeks, 2008; Tsay, 2016), and a sub-sample of that time series was displayed in Fig. 1.

4.3 Average distributions

The average shapes of the distribution functions display remarkable similarities in US and UK expectations. Fig. 3 reports the functional mean and median of the centered density functions for both surveys (bold blue and black lines, respectively), overlaid with the cross-sectional densities for every time period in our sample (thin lines). For the latter, lighter colours correspond to observations further (in the sense of band depth) from the median.[14] The standard deviation of the belief distribution is 4.2 percent in the US sample (Fig. 3, Col. 1), somewhat higher than the 2.3 percent seen in the UK (Fig. 3, Col. 2), since the former includes observations from the high-inflation period of the late 1970s while the latter does not, owing to the shorter sample at our disposal. The standardized third moment of the Michigan distribution is .96, and for the Basix distribution is 1.3, indicating that inflation beliefs are skewed quite strongly to the right.[15] The average distributions have excess kurtosis of 3.7 and 3.2, for the US and UK respectively, indicating that they possess fatter-than-normal tails.

4.4 Principal components

Functional principal component analysis of the inflation expectations data reveals that a handful of features are sufficient to account for the bulk of the variation in distributions. To illustrate how common functional components combine to produce approximations to observed distributions (see Section 3.2), consider the estimated distribution of US inflation expectations in 1979-Q3 displayed in Fig. 4. The distribution (grey line) is shown along with its Karhunen-Loève expansion (black line) for successively larger numbers of empirical basis functions K. The distributional shape is highly complex, but the first two components alone are sufficient

[12]Some form of initial data processing is typical in the analysis of functional data (Ramsay and Silverman, 2005, Ch. 1.5), as observations are seldom continuous even if the underlying processes are best thought of that way.

[13]Details of the penalized maximum likelihood (pML) approach we adopt are described in Part II of the Supplementary Material. In the case of the Michigan survey, we discard extreme observations prior to density estimation, using the same truncation rule as those who construct the commonly-used set of summary statistics associated with the data set. For further details on working with Michigan survey data, see Curtin (1996).

[14]Our depth calculation sets the number of curves used to form each band to three, as in López-Pintado and Romo. In practice, we truncate the range of the density functions before computing band depth to avoid regions of the tails which are close to zero. This prevents multiple small curve crossings in regions of near-zero density which would tend to reduce the depth of all functions.

[15]By contrast, the average distribution of professional forecaster beliefs are almost perfectly symmetric about the mean; see the Supplementary Material, Part III.

Figure 3. Mean and median cross-section distributions for year-ahead inflation forecasts

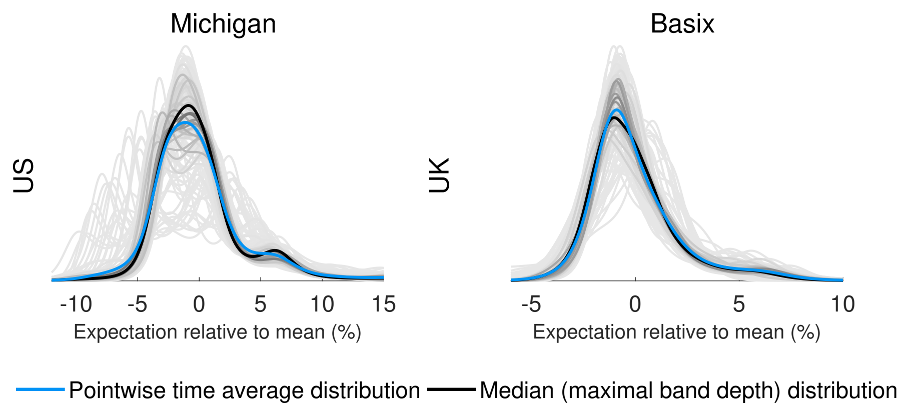

Note: For each survey, panels overlay quarterly distributions of responses for all dates (see text for details). The average expectation at each date has been subtracted to ensure every distribution is mean zero. Darker shaded curves are closer to the median distribution, where the median is the distribution that lies inside the most three-curve bands.

to capture gross features such as the distribution's bimodality, while five components deliver a notably improved approximation (in an ISE sense). A more systematic approach looks across all time periods to compute the average share of variation explained by K components. The scree plot (Fig. 5) displays the ten largest normalized eigenvalues associated with each e_k (left panel) and their cumulative sums (right panel). It can be seen that to explain 90, 95 or 99 percent of variation in either survey requires 2, 3, or 6 components respectively.

Time series variation in approximation accuracy for both surveys is displayed in the complementary Fig. 6. For the Michigan survey (left panel), there is something of a downward trend in the errors between 1978 and 1985, as the observed distributional shapes go from complex and multi-modal towards being close to average. Capturing shapes that are closer to the functional mean naturally requires fewer components. It can be seen that there are some periods—for example, in 1995—where one component alone produces approximately the same magnitude of error as three components. But there are also periods where the two additional components reduce the approximation error by more than an order of magnitude—for example, in 2012. Similar observations apply for the Basix survey (right panel).

5 The role of heterogeneous beliefs in the Phillips curve

We now turn to the empirical assessment of the heterogeneous beliefs version of the Phillips curve. We ask whether heterogeneity 'matters' for inflation dynamics or whether, as is commonly assumed, the consensus survey forecast captures all the relevant effects of expectations on inflation. Our answer to this question will take the form of statistical tests for the presence of distributions of expectations in the more general Phillips curve regression (11), and a quantitative assessment of their economic materiality. We consider identical models and estimation

Figure 4. Density function expansion in the empirical basis for a selected date

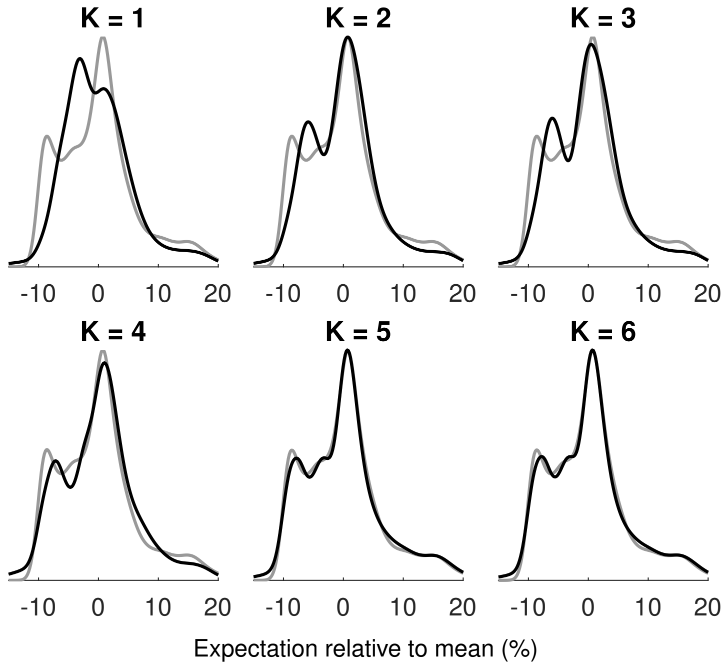

Note: Michigan survey, 1979-Q3. Grey line–observed distribution of expectations. Black line–approximation given by $\hat{\mathsf{p}}^{(K)}_{1979\text{-}Q3,4}$, $K = 1, \ldots, 6$, defined in Eq. (10). The magnitudes of the associated integrated squared errors are $\log_{10}[\mathsf{ISE}^{(K)}_{1979\text{-}Q3,4}] = \{-2.22, -2.71, -2.71, -3.06, -3.83, -3.84\}$.

methods for the United States and United Kingdom. The following section documents the robustness of our finding to various alternative specifications that have been widely examined in the literature.

5.1 United States

The baseline specification for the US Phillips curve relates the annualized quarterly percent change in the seasonally-adjusted consumer price index (CPI) to the CBO measure of the unemployment gap, the average one-year-ahead expected inflation rate from the Michigan survey, and an outlier dummy variable.[16] The importance of the average survey expectation that has been documented in other studies is confirmed by the results (Tab. 1, Col. 1). The slope of the Phillips curve ($\hat{\alpha}$) is around -0.3 (and is significant at 1%), which means that when unemployment is 1 percent below its natural rate, quarterly inflation is 0.3 percent lower (at an annual rate).[17] The substance of these results is very similar to that reported in Coibion et al. (2018), as they are based on an equivalent specification and a sample that is only modestly

[16] The average value of the CPI in a given quarter is used to compute the quarter-on-quarter inflation rate. We use expectations reported in the first month of the quarter, which may incorporate information about last quarter's inflation rate, but cannot incorporate any data for the current quarter. This practice helps to ameliorate concerns over endogeneity bias in the expectations data, but results based on full-quarter responses are very similar.

[17] We find no robust evidence supporting non-linearities in the Phillips curve, for example when we allow for different slopes when unemployment is above or below its natural rate.

Figure 5. Shares and cumulative shares of functional variation explained by the leading K principal components

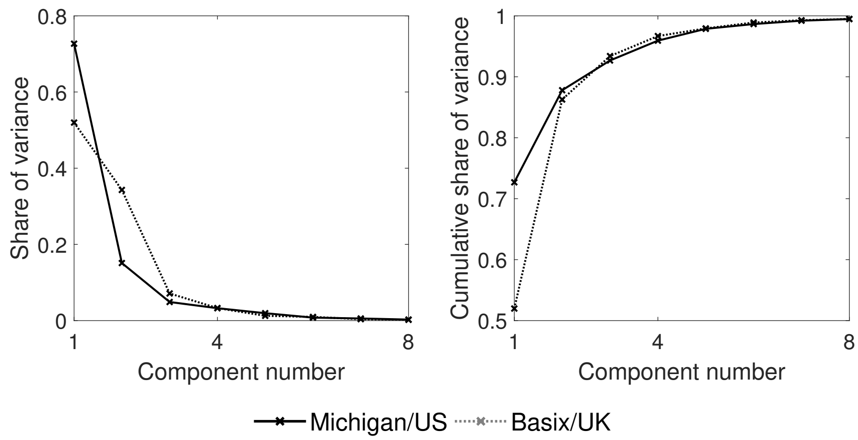

Note: Scree plot gives the normalized sums of eigenvalues of the covariance operator.

extended.

While average expectations hold considerable explanatory power for inflation, heterogeneity in expectations matters too. The T_F-statistic reported in Tab. 1 (Cols. 2–3) indicates that the aggregation function γ is strongly significant in our Phillips curve regressions. The BIC selects three principal components of the belief distributions, but remarkably the penalty for the model with six components is no larger than that for the model with none.[18] The p-values of the functional T_F-statistic are below 0.1%, both when three components are used and when six are used. At the same time, the estimated coefficient on the consensus expectation remains highly significant. Our results are robust to including supply factors (Col. 4).[19] Moreover, formal tests provide no evidence against the stability of parameters on the expectations terms over this sample (see Supplementary Material, Section VI).[20]

Our results provide evidence of a robust association between cross-sectional distributions of expectations and inflation. The estimated coefficient curve indicates that a shift in probability mass towards the upper tail of the distribution produces a positive response in inflation (and vice versa, see Fig. IV in the Supplementary Material). But the coefficient curves themselves can be difficult to interpret. An assessment of the effect size is more easily appreciated by

[18]In models with SPF data, the first component is selected. It has a high correlation with disagreement, and is significant at the 1% level. Overall we find that the household model encompasses the professional forecaster model, in line with the findings reported in Coibion and Gorodnichenko (2015). For further details, see the Supplementary Material, Part III.

[19]Because supply shocks have at times driven inflation and demand—summarized by the unemployment gap—in opposite directions, if omitted, supply factors may impart a downward bias to the coefficient on slack. We include distributed lags in the supply factors in our regressions, and eliminate those variables/lags that are statistically insignificant. For the US, this leads us to retain only the contemporaneous change in the oil price; for the UK, the change in the sterling price of oil and its first lag are retained, along with the change in the relative price of imported goods.

[20]We are also able to assess the stability of the US inflation model using data up to 2022-Q1. The Chow break-point test for 2019-Q4 has an F-statistic of 1.46, which with an $F(9, 160)$ distribution corresponds to a p-value of .17. We are therefore unable to reject the null of no structural change at the onset of the COVID-19 pandemic. See Section 7 for further discussion.

Table 1. Baseline heterogeneous beliefs Phillips curve

	US/Michigan				UK/Basix			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable	CPI	CPI	CPI	CPI	CPI	CPI	CPI	CPI
Unemployment gap	$-.267^{**}$ (.103)	$-.283^{***}$ (.095)	$-.266^{**}$ (.104)	$-.354^{***}$ (.066)	.048 (.131)	$-.233^{**}$ (.109)	$-.237^{**}$ (.108)	$-.176^{*}$ (.100)
Average expectation	1.71^{***} (.104)	1.54^{***} (.131)	1.79^{***} (.168)	1.23^{***} (.095)	1.06^{***} (.129)	$.732^{***}$ (.199)	$.756^{***}$ (.239)	$.890^{***}$ (.187)
Number of FPCs	–	3	6	3	–	3	6	3
T_F-statistic	–	8.32 [.000]	5.26 [.000]	19.22 [.000]	–	8.48 [.000]	5.51 [.000]	5.56 [.001]
Supply factors	n	n	n	y	n	n	n	y
Outlier dummy	y	y	y	y	y	y	y	y
Sample	1978Q1–2017Q4				1986Q4–2017Q4			
R^2	.773	.805	.813	.870	.674	.733	.743	.767
BIC	.914	.858	.914	.486	.652	.571	.645	.509
AIC	.837	.724	.721	.332	.539	.390	.396	.282
DW test (p-value)	.092	.578	.430	.270	.001	.164	.164	.062
Number of obs.	160	160	160	160	125	125	125	125

Note: Estimates of Eq. (11). Dependent variable is the seasonally adjusted annualized quarter-on-quarter percentage change in the consumer price index. Newey-West adjusted (5 lags) standard errors for t test (scalar covariates) appear in parentheses. The p-values for the T_F-statistic appear in brackets. All regressions include a constant. Outlier dummies equal 1 in 2008-Q4 (US and UK) and in 1991-Q2 (UK). Supply factors are the quarterly percentage change in the oil price (US) or the sterling oil price lagged one quarter (UK), and the quarterly percentage change in import prices (UK). Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels.

Figure 6. Integrated square errors of K-component approximations to the distribution of inflation forecasts

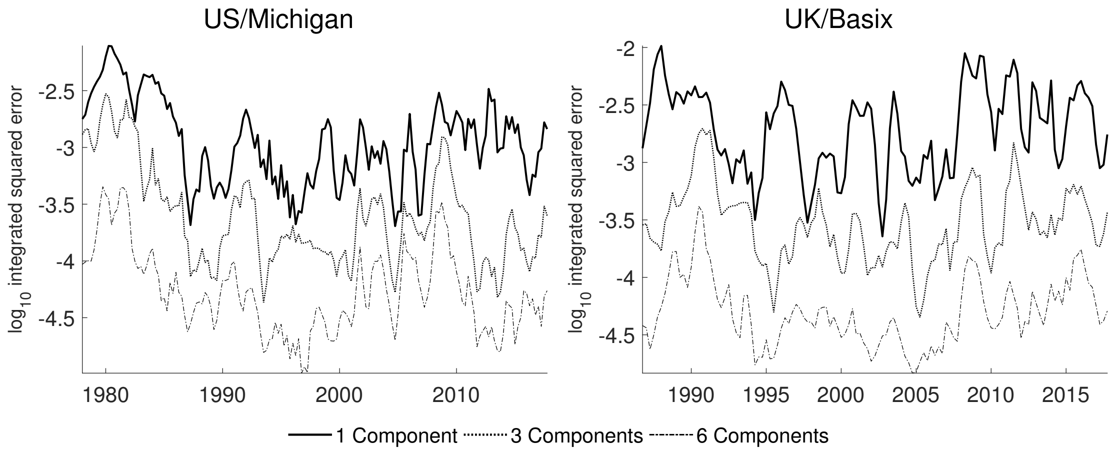

Note: The $\log_{10}$ integrated square error, Eq. (10), associated with a $K = \{1, 3, 6\}$ component expansion of the observed distributions of forecasts. US data is from the Michigan survey; UK data from the Basix survey. Centered three-quarter moving average.

considering the contributions to the regression fitted value from the functional covariates. Fig. 7 reports these (light bars) alongside the contribution from average beliefs (dark bars), in terms of deviations from the sample average to improve legibility. The top panel (for the US) indicates that heterogeneous beliefs have often had a quantitatively relevant effect on inflation over the past 40 years.

A notable insight that can be gained by looking at inflation through the lens of the heterogeneous beliefs model is that the overall effect of expectations was considerably less supportive of inflation following the Great Recession of 2007-9 than is commonly thought. As remarked in the Introduction, the consensus household expectation held up well. But expectations as a whole became more symmetric (less right skewed) and showed lower disagreement than average, imparting a somewhat more disinflationary impulse than average. (Note that since the mean consensus expectation is strongly positive, its overall contribution is always positive, so never implies deflation.) This observation vitiates somewhat the conclusion reached in Coibion and Gorodnichenko (2015), who used the same underlying data, but summarized expectations using the cross-section average alone in their model.

5.2 United Kingdom

We estimated identically-specified Phillips curves on UK data, again using year-ahead expectations. Because no official measures of the natural rate of unemployment exist for the UK for the sample period in question, we compute one by fitting a cubic spline to the raw unemployment data using OLS (Poirier, 1973). Our measure of the unemployment gap is the residual from that regression.[21] Estimates of the Phillips curve that exploit our newly-

[21] Unemployment gap measures based on natural rates estimates constructed using more sophisticated methods, including filter-based methods, were closely comparable to those produced via our spline approach. Moreover,

Figure 7. Estimated contribution of expectations to inflation in the heterogeneous beliefs model

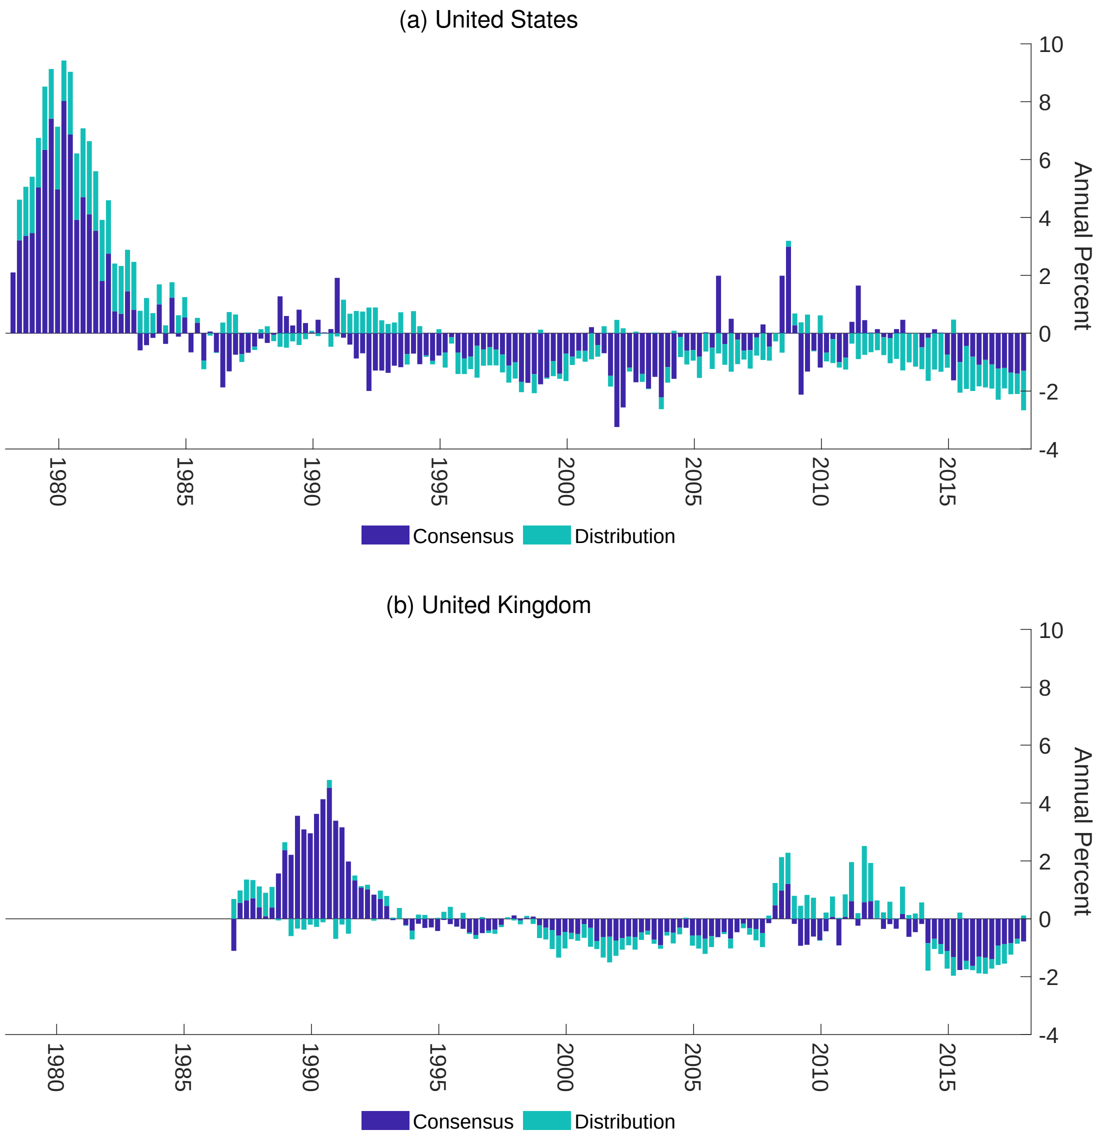

Note: Height of bars represent the contribution to the fitted value for inflation obtained from the regression in Eq. (11). Consensus: the fitted contribution of deviations in the expected year-ahead rate of inflation from its sample mean. Distribution: the fitted contribution of the functional principal components of the distribution of expected year-ahead rate of inflation.

constructed household survey data series (Basix) are also reported in Tab. 1 (Cols. 5–8). The standard variant featuring average expectations only (Col. 5) has a positive ('incorrect') but insignificant slope. The coefficient on average expectations is almost identical to unity.

To the standard specification, we once more add functional principal components from the full distribution of survey responses. The BIC selects three components, but even the model with six components is preferred to that with none. Estimates given in Cols. (6–7) show values for the T_F-statistic that indicate a high level of statistical significance for both three and six components. The additional information also yields a more interpretable model: When the distribution is included, the coefficient on the unemployment gap becomes sizeable, correctly-signed, and significant. Including supply factors does not change the nature of the results (Col. 8).

The impact of heterogeneous expectations on fitted UK inflation appears particularly sizeable around the time of the financial crisis, Fig. 7 (bottom panel). The contribution of the distribution of inflation expectations also leads us to revise our narrative of inflation drivers over the period of devaluation-driven of inflation in 2011; in particular, the reassuring stability of the consensus expectation masked the positive contribution to inflation being made by the distributional factors. The notable features of the UK distributions in this period, shown in Fig. 1 (right panel), are the upward shifts in modal beliefs, and the increased probability mass concentrated at higher rates of inflation. The functional regression assesses a positive contribution to inflation from these distributional movements. Indeed, expectations contributed far more to CPI inflation than did the direct effects from oil and other import prices.

6 Robustness

This section considers the robustness of our findings to two common alternative specifications: a 'hybrid' Phillips curve with both forward- and backward-looking elements; and an inflation 'gap' Phillips curve in which fluctuations around a long-run trend, often associated with expectations about long-run inflation, are considered. We further consider the performance of an obvious alternative to our econometric approach, regression on higher moments.

6.1 Is inflation backward-looking?

An important question in monetary economics is the extent to which inflation depends on its own past values. In a purely backward-looking model, disinflating the economy is costly, because unemployment must be driven high enough for long enough to 'wring out' inflation from the system. In a purely forward-looking model, anticipated disinflations need not be costly at all. Backward-looking inflation behaviour is commonly identified with one of two potential mechanisms. The first is simply that expectations themselves are formed, at least in

constructing the unemployment gap using a spline-interpolated version of the OECD's annual natural rate series, and using that in our regressions, produced estimates of the Phillips curve slope that were very similar to those reported in Table 1.

Table 2. Backward- and forward-looking components in inflation

	US/Michigan			UK/Basix		
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable	CPI	CPI	CPI	CPI	CPI	CPI
Unemployment gap	−.059 (.093)	−.177** (.079)	−.309*** (.062)	−.267* (.119)	−.027 (.101)	−.180* (.095)
Lagged inflation	.701*** (.044)	.231*** (.067)	.083 (.063)	.449*** (.072)	.111 (.074)	.073 (.059)
Average expectation	–	1.20*** (.143)	1.61*** (.192)	–	.918*** (.133)	.814*** (.188)
Number of FPCs	–	–	3	–	–	3
T_F-statistic	–	–	14.28 [.000]	–	–	4.99 [.003]
Supply factors	y	y	y	y	y	y
Outlier dummy	y	y	y	y	y	y
Sample	1978Q1–2017Q4			1986Q4–2017Q4		
R^2	.743	.836	.872	.612	.740	.770
BIC	1.07	.652	.497	.903	.544	.536
AIC	.975	.537	.324	.745	.363	.287
DW (p-value)	.448	.581	.648	.821	.094	.285
Number of obs.	160	160	160	125	125	125

Note: Estimates of Eq. (14). For further explanatory notes, see Tab. 1.

part, in a backward-looking manner. The second mechanism relates to the intrinsic persistence of the inflation process, rather than to the persistence of expectations (or indeed, any of the other determinants of inflation), arising for example due to price indexation (Fuhrer, 2011).

We investigate the extent and sources of backward-looking behaviour using the Phillips curve framework set out above. To our baseline specification, we add an additional term in lagged inflation to produce a hybrid Phillips curve:

$$\pi_t = \alpha\varphi_t + \beta\,\overline{\pi}^e_{t,h} + \delta\pi_{t-1} + \int \gamma(\pi^e)\mathrm{d}\mathsf{P}^{\mathsf{c}}_{t,h}(\pi^e) + \varepsilon_t \tag{14}$$

In Tab. 2 we show the results of adding the expectation terms $\overline{\pi}^e_{t,4}$ and $\mathsf{p}^{\mathsf{c}}_{t,4}$ one at a time to a purely backwards-looking model (one with $\beta = 0$, $\gamma = 0$).

When lagged inflation appears without any forward looking terms in the Phillips curve, its coefficient is large and significant for both the Michigan and Basix models (Tab. 2, Cols. 1 and 4). However, this result is not robust. In both cases, the coefficient on lagged inflation is upward biased because of its positive correlation with the omitted consensus expectation. Adding the average survey expectation substantially reduces the magnitude of the coefficient, consistent with the findings reported by Fuhrer (2017). For the US (Col. 2), the weight on the backward-looking term falls by a factor of three, although it remains significant. The result is a hybrid

Phillips curve, with forward- and backward-looking components both having a statistically significant role. For the UK (Col. 5), adding average expectations results in backward-looking terms becoming economically and statistically indistinguishable from zero. As a result, the other parameter estimates are close to those in Tab. 1 (Col. 5).

Allowing for heterogeneity in inflation expectations eliminates the backward-looking component from the Michigan regression (Tab. 2, Col. 3). The T_F-statistic indicates that the coefficient function is statistically different from zero. Furthermore, omitting the information contained in the distribution of beliefs about future inflation leads to an upward bias in the backward-looking coefficient δ in Eq. (14) even after adding average expectations. For the Basix regression (Col. 6), lagged inflation remains irrelevant, and the distribution function is strongly significant. We also observe that the version with forward-looking terms is preferred by the BIC over the purely backwards-looking version in both regions.

The results in Tab. 2 indicate that intrinsic persistence is not an important feature of the inflation process. But might the distributional effects we identify be capturing lagged effects from consensus expectations? It turns out not. The previous quarter's consensus expectation enters the UK regression at 5% significance, but is insignificant when the distribution of expectations also appears (not shown). For the US, the lagged consensus expectation is never significant. In summary, we find that backward-looking inflation behavior arising either from intrinsic persistence or the persistence in consensus expectations themselves can be rejected in favor of the heterogeneous beliefs model.

6.2 Do trends drive out expectations?

A recent literature recognizes the importance of accounting for time-varying trend inflation when thinking about cyclical inflation dynamics. Cogley and Sbordone (2008) present a microfounded Phillips curve that features inflation trends, and fit it to US data; and leading statistical approaches to modeling and forecasting inflation formulate the inflation process in 'gap' form, that is, in terms of deviations from trend (Stock and Watson, 2007; Faust and Wright, 2013), where trend is typically identified with longer-run survey expectations. A potential concern is that the apparent importance of expected inflation found using the type of models studied in Section 5 may be down to an association between short-run expectations and trend. For example, Cecchetti et al. (2017) argue for the unimportance of short run expectations on this basis, at least in periods where monetary policy was well run. Such perceptions have led some studies to formulate the Phillips curve entirely in terms of longer-run expectations (Hooper et al., 2020).

We modify the baseline heterogeneous beliefs Phillips curve Eq. (8) in two ways. (The following explanation applies to the model for the US; the case of the UK is treated in the Supplementary Material.) First, we introduce the trend component of inflation as measured by the Michigan consensus 5-to-10 year ahead inflation forecast $\overline{\tau}_t^e$ to form an inflation 'gap'.

Table 3. Inflation gaps and heterogeneous beliefs

	US/Michigan			
	(1)	(2)	(3)	(4)
Dependent variable	CPI gap	CPI gap	CPI gap	CPI gap
Unemployment gap	$-.179^{**}$ (.088)	$-.275^{***}$ (.070)	$-.268^{***}$ (.041)	$-.241^{***}$ (.050)
Average expectation gap	$.520^{***}$ (.191)	$.892^{***}$ (.190)	$.601^{***}$ (.089)	$.923^{***}$ (.164)
Number of FPCs (1Y)	–	2	–	2
Number of FPCs (5Y)	–	–	2	3
T_F-statistic	–	8.41 [.000]	10.6 [.000]	5.63 [.000]
Supply factors	y	y	y	y
Outlier dummy	y	y	y	y
Sample	1990Q2–2017Q4			
R^2	.695	.737	.746	.761
BIC	.282	.217	.180	.248
AIC	.160	.046	.010	.004
DW test (p-value)	.130	.850	.336	.386
Number of obs.	111	111	111	111

Note: Estimates of Eq. (15). Dependent variable is the seasonally adjusted annualized quarter-on-quarter percentage change in the consumer price index less the mean expected inflation rate 5-10 year ahead from the Michigan survey. The 'average expectation gap' is the mean expected rate of inflation 1 year ahead less the mean expected rate of inflation 5-10 year ahead. For further explanatory notes, see Tab. 1.

Second, we allow the distribution of beliefs about trend inflation to enter the Phillips curve alongside beliefs about short-run expectations. The modified model generalizes the FLM of Eq. (8) to accommodate multiple functional predictors:

$$\pi_t - \overline{\tau}_t^e = \alpha\varphi_t + \beta(\overline{\pi}_t^e - \overline{\tau}_t^e) + \sum_h \int \gamma_h(\pi^e) \mathrm{dP}_{t,h}^{\mathsf{c}}(\pi^e) + \varepsilon_t \tag{15}$$

for $h = \{4, 20\text{-}40\}$ quarters. The coefficient β measures the relative weight on short- and long-run expectations: $\beta = 0$ implies inflation moves one-for-one with trend inflation; $\beta = 1$ implies the standard model in levels (as opposed to gaps), where inflation depends only upon short-run expectations.[22] Estimation of Eq. (15) is again via FPCR, with functional principal components of both distributions entering. The results of adding in turn the distributions of 1 year ahead forecasts, 5-10 year ahead forecasts, and both distributions are given in Tab. 3. Once again, the choice of principal components is guided by the information criteria.[23]

Distributions of both short- and long-run expectations matter for inflation. When functional effects are absent (Tab. 3, Col. 1), a levels specification can be ruled out as the point estimate for β is statistically different from unity. An intermediate case holds, in which the inflation gap responds to near-term consensus expectations. However, the T_F-statistic indicates that distributional effects on inflation are present. When the distribution of one year ahead expectations is included (Cols. 2 and 4), the point estimate on the consensus expectations gap is indistinguishable from unity, and the standard specification pertains.

The most notable new element of the results is the importance the data affords to the distribution of long-term expectations (Tab. 3, Cols. 3 and 4). The selection criteria point to a specification that includes heterogeneity in beliefs concerning both short and long forecast horizons, with two FPC drawn from the year-ahead distribution and three from the 5-10 year ahead distribution. By contrast, consensus beliefs about trend inflation receive essentially no weight. This result suggests that the attention given to trend inflation, in this case measured using beliefs about the long-run, are half right; although the long-run consensus is of limited relevance to near-term dynamics, long-run beliefs lying away from the consensus impact inflation.

6.3 Can moments substitute for functional principal components?

In some familiar circumstances, a probability distribution can be determined from knowledge of its moments.[24] This is the case, for example, for distributions from the exponential family. The present section therefore investigates the econometric merits of a potential alternative to functional principal component regression in which functional effects are summarized

[22] Model (15) is similar in spirit to Models (8) and (9) of Faust and Wright (2013). Those authors use lagged inflation to proxy forward-looking behaviour rather than directly including survey expectations as a covariate.

[23] Results using monthly data are presented in Tab. V.1 of the Supplementary Material. There are no material changes to the monthly results discussed here.

[24] The quoted inversion is known as the ‘problem of moments’. A correspondence between moments and distributions need not exist, or be unique.

using the standardized moments of belief distributions. Our approach is to assess whether the FPCR specification encompasses the moments specification (or vice versa) by means of the Davidson and MacKinnon (1981) *J*-test for non-nested models. Given competing models A and B, the *J*-statistic tests the null that given model A, no extra explanatory power is provided by B.[25] In our case, the models specified in Tab. 1 (Col. 4 and 8. for the US and UK) will constitute model A. The competing model B will substitute the standard deviation, and standardized third and fourth moments, for the functional terms in model A.

The results for the US indicate that the moments-based model does not hold any additional explanatory power for inflation once the information in the FPCs has been accounted for. The *J*-statistic in the case where the FPC specification is the null has a *p*-value of 0.5 (when three components are present). At the same time, the FPC model adds significant explanatory power to the moments-based model (*p*-value of zero). These results provide clear evidence in favor of the FPC model—it encompasses, but is not encompassed by, a regression on moments. For the UK, the results are less clear-cut. Each model contains some information that the other does not (the null of the *J*-test is rejected at 1% significance for both models in turn).[26] In sum, researchers should bear in mind that although we favor FPCR, the merits of other approaches can be considered on a case-by-case basis.

7 Inflation and the pandemic

In the two years or so following the onset of the COVID-19 pandemic in early 2020, US inflation was more volatile that at any time since the global financial crisis. A number of factors underlay its movements: the headline rate of US unemployment increased by more than 10 percentage points in the space of two months; prices of some primary commodities, notably oil, collapsed, then strongly rebounded; consumers' expenditure patterns underwent rapid shifts in response to changes in both the demand and supply of goods and services; and survey measures of households' consensus inflation expectations, which were initially stable, then increased notably as the economy recovered during 2021 (Fig. 8, Panel (a)).

We investigate the behavior of inflation during this period through the lens of our heterogeneous beliefs model. As policymakers needed to monitor macroeconomic out-turns at high frequency during the pandemic, we re-estimated our quarterly models at a monthly frequency using data up to March 2022. Shifts in longer-run expectations may have been especially germane in this period. Fig. 8 (Panel (b)) compares the distribution of 5-10 year ahead Michigan expectations during the economic recovery of late 2021 and early 2022 to the distribution that pertained immediately prior to the pandemic. A material upside skew can be observed, along

[25]The Davidson and MacKinnon (1981) procedure tests for encompassing, although they do not explicitly state that fact. The test is implemented by means of an 'artificial' regression that combines model A with the fitted values from model B. A significant *J*-statistic (which is the *t*-statistic on the fitted values in the aforementioned regression) indicates a rejection of the null hypothesis that A encompasses B.

[26]The FPCR approach dominates once again when we increase the number of PCs to six.

Figure 8. US inflation and household expectations during the economic recovery from the pandemic

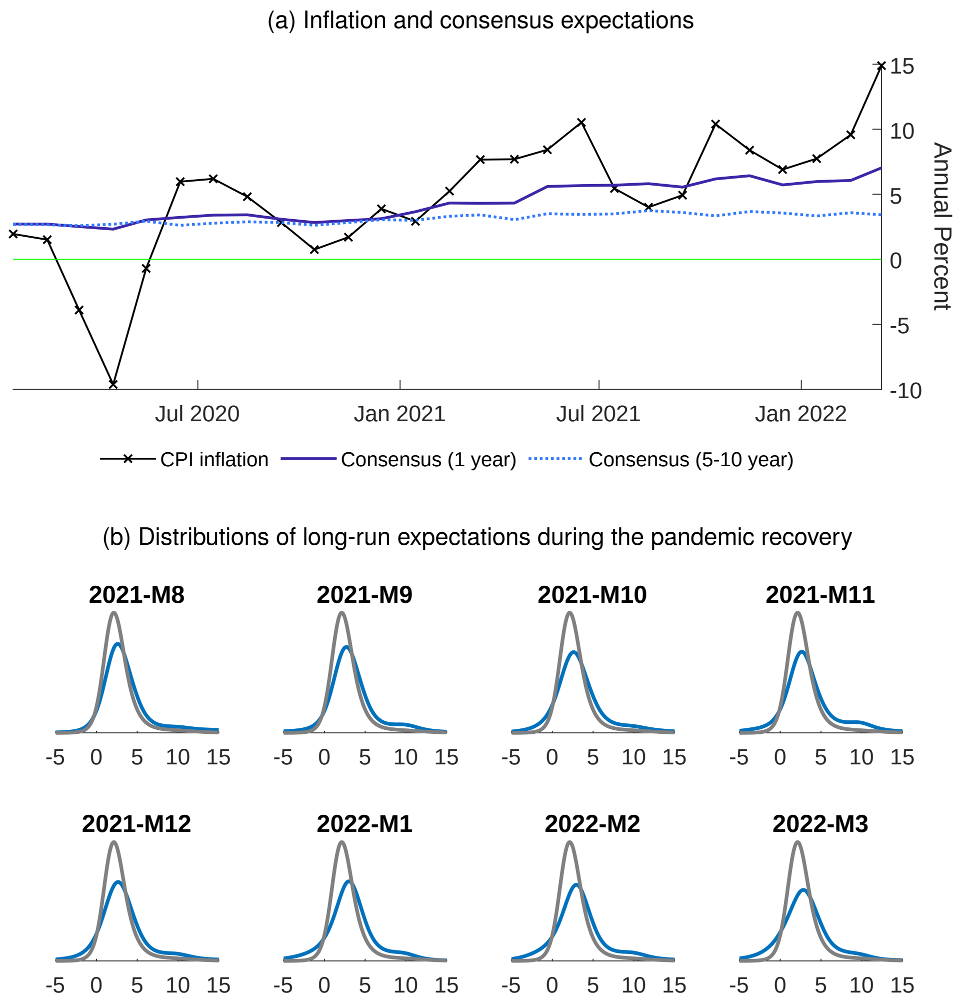

Note: Panel (a) shows month-on-month CPI inflation (annual rate), along with mean Michigan Survey inflation expectations at 1 year and 5-10 years ahead. Panel (b) shows the distribution of Michigan Survey inflation expectations at 5-10 years ahead (blue) and the same distribution as of January, 2020 (grey). Dates reflect when forecasts were made.

with a notable rise in disagreement. To capture these unusual movements, we adopt the specification selected in Section 6.2, with the year-ahead consensus and distributions of shorter- and longer-run expectations both featuring. The specification finds strong support in the data: With one and three FPCs (respectively), the T_F-statistic is 7 (for a p-value of zero).

Household expectations account for around 6 percentage points of the run-up in inflation between January 2020 and March 2022, according to our model (Fig. 2). The pick-up in near-term consensus expectations makes up around half of this total. The remaining 3 percentage points is accounted for by changes in the distributions of expectations, relative to January 2020. Monetary policymakers are typically sensitive to the behavior of longer-run expectations due to the importance attached to anchoring such beliefs around the inflation target. The shifting distribution of longer-run beliefs illustrated in 8 (Panel (b)) started to have a material effect on inflation in June 2021, when policy institutions generally perceived the rise in inflation as being transitory in nature. Notably, the long-run expectations effect approximately doubled to March 2022. The heterogeneous beliefs model could prove valuable as a tool for policymakers to monitor and interpret future inflation developments.

8 Conclusion

Expectations are widely believed to be a core driver of inflation, and consensus expectations from surveys are frequently used in empirical work. However, surveys reveal that at any given time, respondents hold a broad range of beliefs about future inflation. The heterogeneity in beliefs present in the survey data poses problems both of measurement and of potential inconsistency between theoretical and empirical work.

This paper has proposed a straightforward approach to estimating and testing the relevance of heterogeneous beliefs in an otherwise standard empirical Phillips curve. Our principal findings lend support to a model in which expectations that differ from the 'consensus' or average forecast are determinants of inflation. The contribution of heterogeneous beliefs to inflation is estimated to be largest at times of macroeconomic disruption, when disagreement between agents is particularly marked: for example due to a changes in monetary regime, the global financial crisis, and the Covid-19 pandemic.

Our paper establishes both a statistical association between distributions of expectations and inflation, and their quantitative salience. Neither finding has previously been recognized. However, the rather general nature of our approach does not lend itself to discriminating between alternative theories of belief formation. Work to develop models that can account jointly for time-varying distributions of beliefs and for the influence of those distributions on inflation is needed to refine our understanding of the findings we report.

Acknowledgements

We are indebted to John Barrdear, Andy Filardo, Ferre De Graeve, Domenico Giannone, Francesco Grigoli, Jesper Lindé, Marco Del Negro, Paul Hubert, Mike Kiley, Emmanuel Moench, Giovanni Ricco, Paolo Surico, Mark Watson, and especially to Refet Gürkaynak and Ricardo Reis for helpful comments and discussions. We also thank conference and workshop participants at the CEPR Paris Symposium, Bank of England, the European Central Bank, Glasgow Business School, Lancaster Business School, London Business School, the Oxford-Indiana macroeconomic policy conference, the Sciences-Po/OFCE workshop on empirical monetary economics, the FRB St. Louis Applied Time Series Econometrics workshop, the SKEMA Business School workshop on inflation forecasting, the Royal Economic Society annual conference, Queen Mary University of London, UCLouvain, and Warwick Business School. An earlier version of this paper was circulated as Bank of England Staff Working Paper No. 807 (June, 2019). The views expressed in this paper are those of the authors, and should not be taken to represent the views of the IMF, its Executive Board, or IMF management.

References

Klaus Adam and Mario Padula. Inflation dynamics and subjective expectations in the United States. *Economic Inquiry*, 49(1):13–25, January 2010.

Philippe Andrade and Hervé Le Bihan. Inattentive professional forecasters. *Journal of Monetary Economics*, 60(8):967–982, November 2013.

Philippe Andrade, Richard K. Crump, Stefano Eusepi, and Emanuel Moench. Fundamental disagreement. *Journal of Monetary Economics*, 83:106–128, October 2016.

Andrew Ang, Geert Bekaert, and Min Wei. Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54:1163–1212, 2007.

Carola Conces Binder. Whose expectations augment the Phillips curve? *Economics Letters*, 136: 35–38, November 2015.

Carola Conces Binder. Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90:1–12, October 2017.

Clive G. Bowsher and Roland Meeks. The dynamics of economic functions: Modeling and forecasting the yield curve. *Journal of the American Statistical Association*, 103(484):1419–1437, 2008.

William A. Branch and Bruce McGough. A New Keynesian model with heterogeneous expectations. *Journal of Economic Dynamics and Control*, 33:1036–1051, 2009.

Sophocles N. Brissimis and Nicholas S. Magginas. Inflation forecasts and the New Keynesian Phillips curve. *International Journal of Central Banking*, 4(2):1–22, June 2008.

Stephen G. Cecchetti, Michael E. Feroli, Peter Hooper, Anil K. Kashyap, and Kermit L. Schoenholtz. Deflating inflation expectations: The implications of inflation's simple dynamics. Technical report, University of Chicago Booth School of Business, March 2017.

Kausik Chaudhuri, Minjoo Kim, and Yongcheol Shin. Forecasting distributions of inflation rates: the functional auto-regressive approach. *Journal of the Royal Statistical Society, Series A*, 179(1):65–102, 2016.

Timothy Cogley and Argia M. Sbordone. Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve. *American Economic Review*, 98(5):2101–2126, 2008.

Olivier Coibion and Yuriy Gorodnichenko. What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159, February 2012.

Olivier Coibion and Yuriy Gorodnichenko. Is the Phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1): 197–232, 2015.

Olivier Coibion, Yuriy Gorodnichenko, and Rupal Kamdar. The formation of expectations, inflation and the Phillips curve. *Journal of Economic Literature*, 56(4):1447–1491, December 2018.

Olivier Coibion, Yuriy Gorodnichenko, and Mauricio Ulate. Is inflation just around the corner? the phillips curve and global inflationary pressures. *AEA Papers and Proceedings*, 109:465–69, 2019.

Antonio Cuevas. A partial overview of the theory of statistics with functional data. *Journal of Statistical Planning and Inference*, 147:1–23, 2014.

Richard T. Curtin. Procedure to estimate price expectations. Surveys of Consumers, University of Michigan, January 1996.

Russell Davidson and James G. MacKinnon. Several tests for model specification in the presence of alternative hypotheses. *Econometrica*, 49(3):781–793, 1981.

Marco Del Negro, Michele Lenza, Giogio E. Primiceri, and Andrea Tambalotti. What's up with the Phillips curve? *Brookings Papers on Economic Activity*, Spring:301–357, 2020.

Julian Faraway. Comment on kneip and utikal. *Journal of the American Statistical Association*, 96 (454):537, June 2001.

Jon Faust and Jonathan H. Wright. Forecasting inflation. In Graham Elliott and Alan Timmermann, editors, *Handbook of Economic Forecasting*, volume 2A, pages 2–56. Elsevier, 2013.

Andrew Filardo and Hans Genberg. Targeting inflation in Asia and the Pacific: Lessons from the recent past. In D. Cobham, O. Eitrheim, S. Gerlach, and J. Qvigstad, editors, *Twenty Years of Inflation Targeting: Lessons Learnt and Future Prospects*, pages 144–171. Cambridge, 2010.

Jeffrey C Fuhrer. The role of expectations in us inflation dynamics. Technical report, Working Papers, 2011.

Jeffrey C. Fuhrer. Expectations as a source of persistence: Evidence from survey expectations in a dynamic macro model. *Journal of Monetary Economics*, 86:22–35, 2017.

Milton Harris and Artur Raviv. Differences of opinion make a horse race. *Review of Financial Studies*, 6(3):473–506, 1993.

Peter Hooper, Frederic S. Mishkin, and Amir Sufi. Prospects for inflation in a high pressure economy: Is the Phillips curve dead or just hibernating? *Research in Economics*, 74:26–62, 2020.

Lajos Horváth and Piotr Kokoszka. *Inference for Functional Data with Applications*. Springer Series in Statistics. Springer-Verlag, New York, 2012.

Ian T. Jolliffe. *Principal Component Analysis*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2002.

Alois Kneip and Klaus J. Utikal. Inference for density families using functional principal component analysis. *Journal of the American Statistical Association*, 96(454):519–532, June 2001.

Dehan Kong, Ana-Maria Staicu, and Arnab Maity. Classical testing in functional linear models. *Journal of Nonparametric Statistics*, 28(4):813–838, 2016.

Sara López-Pintado and Juan Romo. On the concept of depth for functional data. *Journal of the American Statistical Association*, 104(486):718–734, 2009.

N. Gregory Mankiw, Ricardo Reis, and Justin Wolfers. Disagreement about inflation expectations. In Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual*, volume 18, pages 209–248. University of Chicago Press, 2003.

Sophocles Mavroeidis, Mikkel Plagborg-Møller, and James H. Stock. Empirical evidence on inflation expectations in the New Keynesian Phillips curve. *Journal of Economic Literature*, 52 (1):124–188, March 2014.

Roland Meeks and Francesca Monti. Heterogeneous beliefs and the Phillips curve. Staff Working Paper 807, Bank of England, June 2019.

Joon K. Park and Junhui Qian. Functional regression of continuous state distributions. *Journal of Econometrics*, 167:397–412, 2012.

Andrew J. Patton and Allan Timmermann. Why do forecasters disagree? lessons from the term structure of cross-section dispersion. *Journal of Monetary Economics*, 57:803–820, 2010.

Damjan Pfajfar and Emiliano Santoro. Heterogeneity, learning and information stickiness in inflation expectations. *Journal of Economic Behavior & Organization*, 75(3):426–444, September 2010.

Dale J. Poirier. Piecewise regression using cubic spline. *Journal of the American Statistical Association*, 68(343):515–524, 1973. ISSN 01621459. URL `http://www.jstor.org/stable/2284770`.

James O. Ramsay and Bernard W. Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2005.

Philip T. Reiss and R. Todd Ogden. Functional principal component regression and functional partial least squares. *Journal of the American Statistical Association*, 102(479):984–996, September 2007.

Philip T. Reiss, Jeff Goldsmith, Han Lin Shang, and R. Todd Ogden. Methods for scalar-on-function regression. *International Statistical Review*, 85(2):228–249, 2017.

John M. Roberts. New Keynesian economics and the Phillips curve. *Journal of Money, Credit and Banking*, 27(4):975–984, November 1995.

James H. Stock and Mark W. Watson. Why has US inflation become harder to forecast? *Journal of Money Credit and Banking*, 39:3–33, 2007.

Ruey S. Tsay. Some methods for analysing big dependent data. *Journal of Business and Economic Statistics*, 34(4):673–688, October 2016.

A An introduction to functional regression

This section provides a condensed primer on functional regression. The literature on estimation of the functional linear model is extensive. An excellent treatment of functional principal component regression may be found in Reiss and Ogden (2007), with Reiss et al. (2017) providing an up-to-date survey. A textbook treatment of estimation and inference in the functional linear model is given by Horváth and Kokoszka (2012), while the particular approach to inference we adopt is due to Kong et al. (2016).

Although various formalizations of functional data are found in the literature (Cuevas, 2014, Section 2.3), we follow common practice and take X to be a measurable function in a sample space $L^2(\mathcal{I}), \mathcal{I} \subset \mathbb{R}$ defined on a probability space $(\Omega, \mathcal{F}, \mathrm{P})$. The real-valued scalar random variable Y is defined on the same probability space as X. We have a sample $(y_t, \mathsf{x}_t), t = 1, \ldots T$ drawn from (Y, X). The scalar-on-function (SOF) regression model is defined as:

$$y_t = m_y + \int \gamma(\iota)\mathsf{x}_t(\iota)\mathrm{d}\iota + \varepsilon_t, \qquad \varepsilon_t \sim \text{i.i.d.}(0, \sigma^2)$$

where γ is a square integrable function, $\|\gamma^2\| < \infty$, and ε is independent of x. Here and elsewhere integration is over $\mathcal{I}$. We express the functional regressor in terms of its Karhunen-Loève expansion, truncated at the Kth term:

$$\mathsf{x}_t(i) = \sum_{k=1}^{K} s_{kt}\mathsf{e}_k(i)$$

where the principal component scores $s_{kt} = \langle \mathsf{x}_t, \mathsf{e}_k \rangle$ satisfy $\mathbb{E}[s_{kt}] = 0$, $\mathbb{E}[s_{kt}^2] = \lambda_k$, and $\mathbb{E}[s_{kt}s_{k't}] = 0, k \neq k'$. As we observe only T curves, there are at most $T-1$ non-zero eigenvalues, so we must choose $K \leq T-1$. Expand the coefficient function in the same basis to obtain:

$$\gamma(i) = \sum_{k'=1}^{K} \gamma_{k'}\mathsf{e}_{k'}(i)$$

We may then express the integral in the SOF model as:

$$\begin{aligned}
\int \left(\sum_{k'=1}^{K} \gamma_{k'}\mathsf{e}_{k'}(i)\right)\left(\sum_{k=1}^{K} s_{kt}\mathsf{e}_k(i)\right)\mathrm{d}i &= \sum_{k=1}^{K} \gamma_k s_{kt} \int \mathsf{e}_k(i)^2\mathrm{d}i \\
&= \sum_{k=1}^{K} \gamma_k s_{kt}
\end{aligned}$$

where the first line follows from $\langle \mathsf{e}_k, \mathsf{e}_{k'} \rangle = 0, k \neq k'$, and the second line follows from $\|\mathsf{e}_k\| = 1$. Making the above substitution, the SOF model may be written as a multiple regression:

$$y_t = m_y + \sum_{k=1}^{K} \gamma_k s_{kt} + \varepsilon_t \tag{A.1}$$

The normal equations for the γs are then immediately seen to be:

$$0 = \sum_{t=1}^{T} s_{jt}\left\{(y_t - m_y) - \sum_{k=1}^{K} \gamma_k s_{kt}\right\}, \quad j = 1, \ldots, K$$

Recalling that the scores are orthogonal, and that the variance of the jth score is equal to the jth eigenvalue, it is easy to see that:

$$\hat{\gamma}_j = \frac{c_{y,s_k}}{\lambda_j} \tag{A.2}$$

where $c_{y,s_k} = \sum_t (y_t - m_y) s_{jt}$ is the sample covariance between the dependent variable and the jth score. It follows that our estimate of the functional coefficient will be given by:

$$\hat{\gamma}(i) = \sum_{k=1}^{K} \frac{c_{y,s_k}}{\lambda_j} \mathsf{e}_k(i) \tag{A.3}$$

As we have seen, SOF regression using FPCs reduces to multiple regression, so extending the model to include scalar covariates, as in our application, is rather routine.

B Computing functional principal components

This section gives the computational results necessary to compute the functional principal components used throughout this paper. The basic approach is to replace functions with linear combinations of basis functions. The material, which is standard, draws on Ramsay and Silverman (2005, Section 8.4).

The FPCA problem is to find to minimize the integrated squared error criterion:

$$\mathsf{ISE}_{t,h}^{(K)} = \int \left\{ \left(\hat{\mathsf{p}}_{t,h}^{(K)} - \bar{\mathsf{p}}_h \right) - \left(\mathsf{p}_{t,h}^{\mathsf{c}} - \bar{\mathsf{p}}_h \right) \right\}^2 \mathrm{d}\iota, \quad \text{where} \quad \hat{\mathsf{p}}_{t,h}^{(K)} = \bar{\mathsf{p}}_h + \sum_{k=1}^{K} s_{kt} \mathsf{e}_k \tag{B.1}$$

averaged over all t, subject to the constraint that the functions $\mathsf{e}(\cdot)$ satisfy $\|\mathsf{e}_k\| = 1$ for all k, and $\langle \mathsf{e}_k, \mathsf{e}_j \rangle = 0, k \neq j$.

Let the functions $\{\mathsf{x}_t(i)\}_1^T$ be defined as in Appendix A. The eigenequation of the covariance operator $V(\mathsf{x})(\cdot)$ is:

$$\int v(i,j) \mathsf{e}_k(i) \mathrm{d}j = \lambda_k \mathsf{e}_k(i) \tag{B.2}$$

Now let the basis expansion of the x_t be:

$$\mathsf{x}_t(i) = \sum_{k=1}^{K} c_{tk} \phi_k(i)$$

or, stacking by t:

$$\mathbf{x}(i) = \mathbf{C}\boldsymbol{\phi}(i), \qquad \underset{(T\times K)}{\mathbf{C}} = [c_{tk}] \quad \text{and} \quad \underset{(K\times 1)}{\boldsymbol{\phi}} = [\phi_k]$$

We may then express the sample covariance function as:

$$v(i,j) = (T-1)^{-1} \boldsymbol{\phi}(i)^\top \mathbf{C}^\top \mathbf{C} \boldsymbol{\phi}(j) \tag{B.3}$$

Assume that the eigenfunctions have the basis expansion:

$$\mathsf{e}(i) = \sum_{k=1}^{K} b_k \phi_k(i) = \boldsymbol{\phi}(i)^\top \mathbf{b}, \qquad \underset{(K\times 1)}{\mathbf{b}} = [b_k]$$

Then substituting (B.3) into (B.2), the eigenequation may be written:

$$(T-1)^{-1} \boldsymbol{\phi}(i)^\top \mathbf{C}^\top \mathbf{C} \mathbf{W} \mathbf{b} = \lambda \boldsymbol{\phi}(i)^\top \mathbf{b} \tag{B.4}$$

where the symmetric $(K \times K)$ matrix $\mathbf{W} = \int \boldsymbol{\phi}(i)\boldsymbol{\phi}(i)^\top$ is a matrix of inner products of the basis functions $\boldsymbol{\phi}_k(\cdot)$, and λ is the eigenvalue corresponding to e. Observing that (B.4) must hold for all i implies that a solution to (B.2) may be obtained from the solution to the symmetric matrix eigenvalue problem:

$$(T-1)^{-1}\mathbf{W}^{1/2}\mathbf{C}^\top\mathbf{C}\mathbf{W}^{1/2}\mathbf{u} = \lambda\mathbf{u}, \qquad \mathbf{u} = \mathbf{W}^{-1/2}\mathbf{b}$$

using standard methods. For an alternative approach that applies standard PCA to the grid of G values $\{\mathsf{p}_{t,h}(x_i)|i = 1,\ldots,G; t = 1,\ldots,T\}$, see Tsay (2016, Section 3.3).