ARTICLE IN PRESS

Journal of Monetary Economics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Heterogeneous beliefs and the Phillips curve

Roland Meeks[a,*], Francesca Monti[b]

[a] International Monetary Fund, United States and CAMA, Australia
[b] UCLouvain, Belgium and CEPR, United Kingdom

ARTICLE INFO

JEL classification:
C4
C55
D84
E31

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

ABSTRACT

Heterogeneous beliefs modify the New Keynesian Phillips curve by introducing a term in the cross-section distribution of expectations. To take that model to the data, we develop a novel functional data approach to estimation and inference that accounts for variation in distributions of expectations. We find that this variation may be summarized using a handful of functional factors, and demonstrate their statistical and economic relevance for inflation dynamics. Our results are among the first to highlight the potential benefits to be gained in empirical work from a rigorous treatment of diverse beliefs in the study of macroeconomic outcomes.

1. Introduction

A prominent feature of expectations reported in surveys—whether of households, firms, or professional forecasters—is their cross-sectional dispersion (interpersonal heterogeneity). This dispersion in beliefs, which is often termed ‘disagreement’ following the early work of Mankiw et al. (2003), exhibits considerable time variation, as well as a strong correlation with the ‘consensus’ forecast (the central tendency of the distribution of beliefs; see Rich and Tracy, 2010). Yet curiously when it comes to linking expectations to inflation, the discussion amongst researchers and policymakers typically focuses on the consensus expectation alone, a tacit assumption that the dispersion in beliefs can be safely ignored. The central argument of the present paper is that a systematic treatment of heterogeneous beliefs benefits our understanding of how expectations shape 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] It is not immediately obvious why beliefs in the tails of the distributions should necessarily be considered uninformative. In fact, Reis (2021) presents qualitative evidence for several countries in support of the notion that the second and third moments of belief distributions can help predict inflation, and Tsiaplias (2020) reports similar results in the case of Australia.

* Corresponding author.
E-mail addresses: rmeeks@imf.org (R. Meeks), francesca.monti@uclouvain.be (F. Monti).

[1] The prevalence of forecasts of 0%, 5% and 10% inflation is not an artefact, 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.

https://doi.org/10.1016/j.jmoneco.2023.06.003
Received 16 September 2022; Received in revised form 12 June 2023; Accepted 12 June 2023
Available online 22 June 2023
0304-3932/

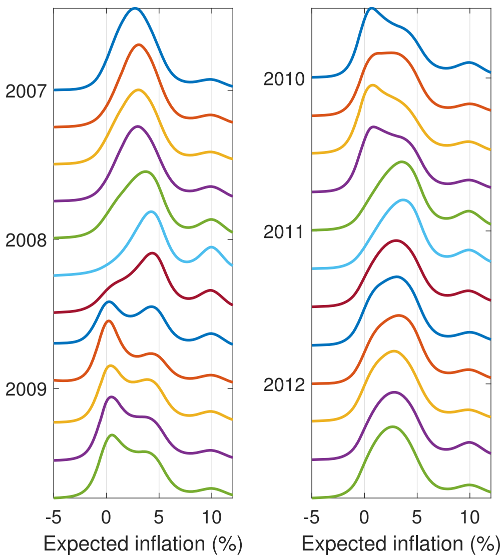

Fig. 1. Cross-section distributions of US household inflation forecasts during the Global Financial Crisis and its aftermath. *Note:* Panels show time series of distributions of individual survey respondents' year-ahead point forecasts. Dates reflect when forecasts were made. Details on the survey data may be found in Section 4.1. Details of the density estimation method may be found in Section C of the Appendix.

Heterogeneity in beliefs has at least two important consequences. First, that the average expectation may not adequately summarize the state of beliefs relevant to the determination of inflation, leading to errors of 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 and amid the recent COVID-19 pandemic and its aftermath. Such episodes are precisely when econometric models tend to break down. The second consequence is that heterogeneous beliefs undermine the microfoundations 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 is not generally preserved when beliefs are averaged, even if the condition is satisfied by the beliefs of each individual agent.[2] This aggregation problem leads to an additional term in the Phillips curve that depends on the cross-section dispersion of beliefs. The potential empirical relevance of this theoretical result remains to be explored.

Our paper provides a rigorous econometric approach that can effectively address both of the aforementioned consequences of heterogeneous expectations on inflation dynamics. It includes both a framework for conducting statistical tests for the presence of distributions in the heterogeneous beliefs Phillips curve, and an evaluation of the quantitative impact of distributional shifts on inflation. Our paper focuses on inflation and inflation expectations, and so primarily relates to the extensive body of work that makes use of 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), Coibion et al. (2018). But the techniques we set out could be of use in other settings where expectations are of central import, for example when estimating households' elasticity of inter-temporal substitution (Attanasio et al., 2020; Crump et al., 2022) or in asset pricing models (see Adam and Nagel, 2023, for a review).

An obvious obstacle to incorporating heterogeneous beliefs into standard macroeconometric models is the very high dimensionality of the survey response data. We approach this problem by considering large cross-sections of survey forecasts as originating from continuous distributions, in which there is an underlying 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 (Bowsher and Meeks, 2008; Ramsay and Silverman, 2005; Tsay, 2016). Applying principal component analysis, suitably adapted for functional data, to the distributional data reveals that a handful of factors can jointly characterize much of the variation therein (Kneip and Utikal, 2001). These factors may then be used to estimate the heterogeneous beliefs model using functional principal component regression (Reiss and Ogden, 2007).

The payoff to our new approach is the discovery of a richer and more nuanced role for expectations in the US inflation process than had previously been recognized.[3] We show that beliefs about future inflation contained in the distributions we study—in addition to the consensus expectation—are powerful in a statistical sense for explaining current inflation, and have effects that are quantitatively

[2] Imposing the law 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. However, this entails some rather strong and therefore unappealing assumptions on agents' beliefs (Adam and Padula, 2010).

[3] We replicate much of the analysis also on UK data, using the newly-collated UK Basix household survey. Results for the UK can be found in the working paper version of this paper, Meeks and Monti (2019), and in Section G of the Appendix.

relevant. We further present novel findings concerning the horizon over which expectations are relevant. In particular, we show that although consensus expectations of near-term inflation are the channel by which beliefs affect price-setting behavior, and that consensus expectations at longer horizon beliefs play a minor role if any, *distributions* of long-horizon beliefs cannot be disregarded. We examine how the inflation narrative is affected by the shifting distributions of near- and long-term beliefs through a close examination of two notable episodes from recent macroeconomic history—the Great Financial Crisis of 2007-09, and the COVID-19 pandemic of 2020.

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 (see 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, other price-setting mechanisms could also be consistent 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. As in related studies, we treat expectations as primitive objects, which of course they are not.

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 approach to bringing the model to the data. Section 4 provides an analysis of the statistical structure present in the time series of distributions of survey forecasts, and analyzes the drivers of observed variation in them. Section 5 contains our headline results in support of the heterogeneous beliefs model, an analysis of the effects of the distribution of long-run beliefs, and consideration of an alternative estimation approach based on distributional moments. Finally, Section 6 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 a 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.

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):

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

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.[4] 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:

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

where:

$$q_{t+1} := \mathsf{E}_N\left\{q_{t+1}^{\star(j)}\right\} \quad \text{and} \quad \overline{\mathsf{E}}_t(x) := \mathsf{E}_N\left[\mathbb{E}_t^{(j)}(x)\right]$$

noting that the future average relative price has been added and subtracted on the right hand side of Eq. (3), and the overbar notation 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

[4] 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; Branch and McGough, 2009, Section 2.1).

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

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

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'.

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:

$$\Delta_t \approx \mathsf{E}_N\left[\gamma\left(\pi_t^e - \overline{\pi}_t^e\right)\right] \quad \text{or} \quad \lim_N \Delta_t \approx \int \gamma\left(\pi_t^e - \overline{\pi}_t^e\right) \mathrm{d}\mathsf{P}_t^c(\pi^e) \tag{6}$$

in the limit as N becomes large. The term P^c denotes the distribution of beliefs around the consensus (see also Section 3.2).[5]

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

$$\pi_t = \alpha\varphi_t + \beta\overline{\pi}_{t,h}^e + \int \boldsymbol{\gamma}\left(\pi^e - \overline{\pi}_t^e\right) \mathrm{d}\mathsf{P}_t^c(\pi^e) + \varepsilon_t \tag{7}$$

This equation is an example of a partial functional linear model or FLM, wherein a scalar quantity (inflation) is related to a functional covariate (the distribution of inflation expectations).[6] 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 estimated to be greater in some range of values for π^e than others, 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 estimation of the FLM, 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-period ahead inflation forecasts π_h^e made in each period the mean of their distribution. In what follows, the notation $\mathsf{p}_{t,h}(\cdot)$ will denote the distribution of h-step ahead point forecasts made at date t. The sequence $\{\mathsf{p}_{t,h}(\cdot)\}_1^T$ is then a functional time series (Bowsher and Meeks,

[5] We sought to validate the approximation in Eq. (6) with a Monte Carlo exercise (reported in Section L of the Appendix), simulating data from a model that generates plausible fluctuations in individual beliefs and testing how well the Δ_t term can be proxied by the functional principal components of the distribution of beliefs about aggregate inflation. We regress Δ_t on the functional principal components and find that the set of functional principal components is strongly associated with it.

[6] A textbook discussion of the functional linear model may be found in Ramsay and Silverman (2005, Ch. 15), while estimation of the partial FLM is discussed in Shin (2009).

[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 (the mean or median, say), and then to ascertain its relationship with inflation.

2008; Horváth and Kokoszka, 2012; Tsay, 2016).

The sample average distribution shape, or functional mean, of h-step ahead point forecasts is then given by $\overline{\mathsf{p}}_h(x) = \sum_t \mathsf{p}^c_{t,h}(x)\ /T$ where $\mathsf{p}^c_{t,h}$ represents the density function over centered forecasts.[8] An alternative robust measure of central tendency is the functional median. 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 (see López-Pintado and Romo, 2009, and Cuevas, 2014, Section 4.3).

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 analogous to classical PCA as applied to empirical covariance matrices (for an even-paced introduction, see Ramsay and Silverman, 2005, Ch. 8; and Appendix A). 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 $\{\mathsf{p}^c_{t,h}\}_0^T$ can be expressed in terms of an expansion in the orthonormal eigenfunction (or empirical) basis $\{\mathsf{e}_k\}$ as $\mathsf{p}^c_{t,h} = \mu_\mathsf{p} + \sum_{k=1}^\infty \langle \mathsf{p}^c_{t,h}, \mathsf{e}_k\rangle \mathsf{e}_k$. The principal component functions $\mathsf{e}(\cdot)$ form an optimal basis for the observations to hand.[9] They satisfy the conditions $\langle \mathsf{e}_k, \mathsf{e}_k\rangle = 1$ 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 analytic solutions to the principal component problem are not generally available, computational approximations are, the details of which are summarized in Appendix A (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. (7) as a familiar multiple regression problem (see Reiss et al., 2017). 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 given above, the functional linear model of Eq. (7) can be rewritten as:[10]

$$\pi_t = \alpha\varphi_t + \beta\overline{\pi}^e_{t,h} + \sum_{k=1}^K \gamma_k s_{k,t} + \varepsilon_t \tag{8}$$

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. (7) as the multiple regression model Eq. (8), 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. (7) is written compactly as:

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

where with a slight abuse of notation $\boldsymbol{\gamma} = (\gamma_1, \cdots, \gamma_K)^\top \in \mathbb{R}^K$ 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 $\widehat{\boldsymbol{\gamma}} = \mathbf{Q}^{-1}\mathbf{M}^\top(\mathbf{I} - \mathbf{P}_Z)\boldsymbol{\pi}$, where $\mathbf{Q} := (\boldsymbol{\Lambda} - \mathbf{M}^\top\mathbf{P}_Z\mathbf{M})$ is the Schur complement of $(\mathbf{Z}^\top\mathbf{Z})$ in $(\mathbf{X}^\top\mathbf{X})$, and $\boldsymbol{\Lambda} = \mathrm{diag}(\lambda_1, \cdots, \lambda_K)$ 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:

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

Then H_0 can be tested using the F-statistic:

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

9 The approximation obtained by truncating the sum in the Karhunen-Loève expansion at a finite number of terms K is optimal in the sense of minimizing an integrated squared error (ISE) criterion, defined as in Eq. (A.1) in Appendix A.

10 Additional details, along with references to the literature, are given in Appendix Section B. See Shin (2009) for a full account of the partial FLM.

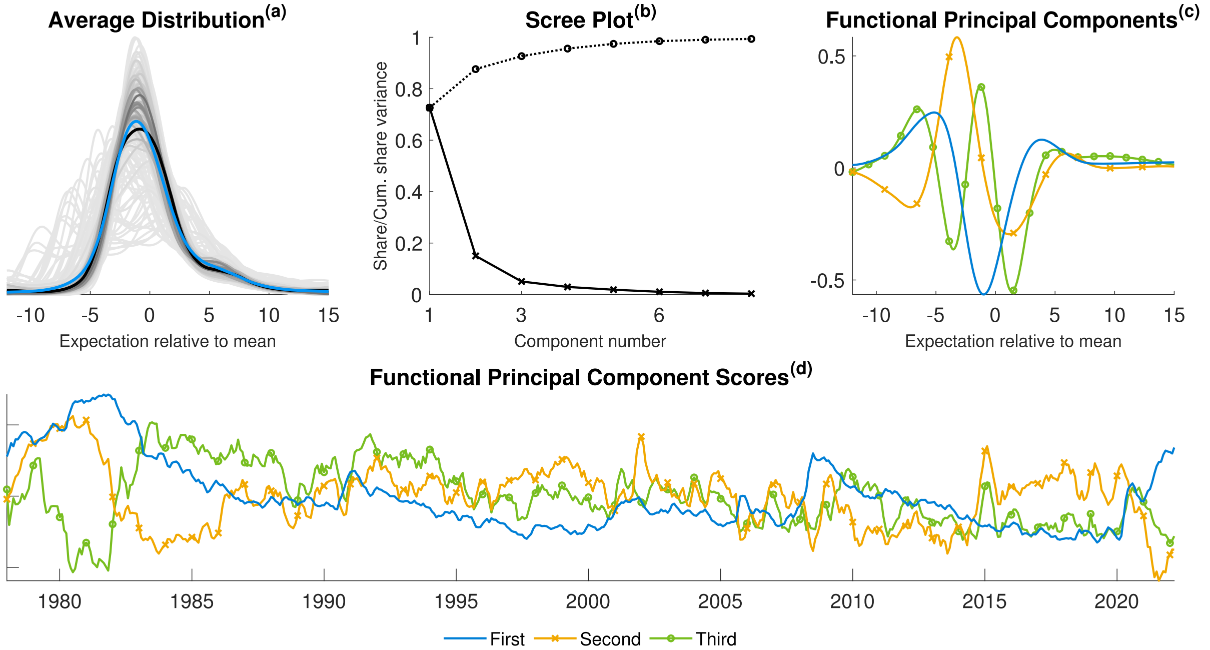

Fig. 2. Statistical summary of the time series of distributions of household inflation expectations. *Notes*: (a) Mean distribution of survey point forecasts (blue), median distribution (black). Distributions are re-centered on average expected inflation each period. Median is the distribution that lies inside the most three-curve bands. All observations overlaid, with darker shaded curves closer to the median distribution. (b) Share of variance (solid line) and cumulated share of variance (dotted line) explained by each eigenvalue. (c) First three eigenfunctions. (d) First three scores normalized so that each has a unit standard deviation, 6 month moving average of monthly data. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

$$T_F = \frac{\boldsymbol{\pi}^\top (\mathbf{P}_X - \mathbf{P}_Z)\boldsymbol{\pi}/K}{\boldsymbol{\pi}^\top (\mathbf{I} - \mathbf{P}_X)\boldsymbol{\pi}/(T-K-N)} \;\overset{approx.}{\sim}\; F_{K,T-K-N} \tag{9}$$

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, and therefore the number of principal component scores that will appear in Eq. (8). Adding more principal components naturally leads to smaller errors of functional approximation (in an integrated square error sense) 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, for which long time series are available. Unfortunately this 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).

[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.

4.1. Data sources

The expectations survey data we use for the US is taken from the Michigan Survey of Consumer Attitudes (MSC). The survey asks questions about 'prices in general' or 'inflation', without specifying a particular measure. Respondents report their expectation for inflation over the following year, and their expectations for the long term. Interviews take place monthly, largely during the month in question, and therefore before that month's consumer price index (CPI) release. A summary of the main features survey data used in this study is given in Appendix Table E.1.

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 penalized maximum likelihood (pML) technique to obtain estimates of that distribution, but any other consistent estimator (such as a kernel estimator) could equally be adopted.[13]

4.3. Average distributions

The functional mean and median of the centered density functions (bold blue and black lines, respectively) are reported in Fig. 2, Panel a, 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 from the Michigan Survey of Consumers is 4.2 percent and its standardized third moment is 0.96, indicating that inflation beliefs are skewed quite strongly to the right.[15] The average distribution displays excess kurtosis of 3.7, indicating that it possesses fatter-than-normal tails.

4.4. Functional principal components

Functional principal component analysis reveals that a handful of features are sufficient to account for the bulk of the variation in the distributions of expectations. Fig. 2, Panel b, displays the eight largest normalized eigenvalues associated with each principal component e_k, and their cumulative sum. It can be seen that to explain 90, 95 or 99 percent of variation in the survey requires 2, 3, or 6 components respectively. The highly structured nature of statistical variation in the surveys will aid our goal to estimate parsimonious models of inflation later on.

We can interpret the components in a statistical sense by scrutinizing the harmonics themselves, shown in Fig. 2, Panel c. The first component—which is associated with the largest eigenvalue—captures changes in the variance of the distributions. To see this, imagine adding the blue function in Panel c to the blue mean function in Panel a; the result would be a more-or-less even shift in distributional mass from the peak of the centered distribution (which is around -1) towards the tails. Similar considerations lead to the conclusion that the second component captures distributional skew. The working paper version of our article provides further analysis to support these associations (Meeks and Monti, 2019), although of course it is worth stressing that as usual in PCA, our functional principal components are identified only up to a rotation.[16]

Time series plots of the first three principal component scores (the 'loadings' on the 'factors') are shown in Fig. 2, Panel d. The scores are mean zero by construction, and their variances are given by the eigenvalues. Several observations are in order. First, the components are highly persistent: on monthly data the first autocorrelations are 0.95, 0.73, and 0.69 for components 1, 2, and 3 respectively. Second, the scores are unconditionally uncorrelated, again by construction, but do exhibit periods of comovement that are particularly evident in the early and late parts of the sample.

An illustration of how functional components combine to produce approximations to observed distributions is displayed in Fig. 3. The actual distribution of inflation expectations in 1979-Q3 (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 to capture gross features such as the distribution's bimodality, while five components deliver a notably improved approximation (in an ISE sense).

[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 approach are described in Part C of the Appendix. 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 Michigan inflation expectations data set.

[14] Our depth calculation sets the number of curves used to form each band to three, as in López-Pintado and Romo, 2009. 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 Section H of the Appendix.

[16] For this reason, and in the absence of an explicit theory of belief formation, the economic interpretation of the factors is problematic. We are grateful to Refet Gürkaynak for highlighting this point.

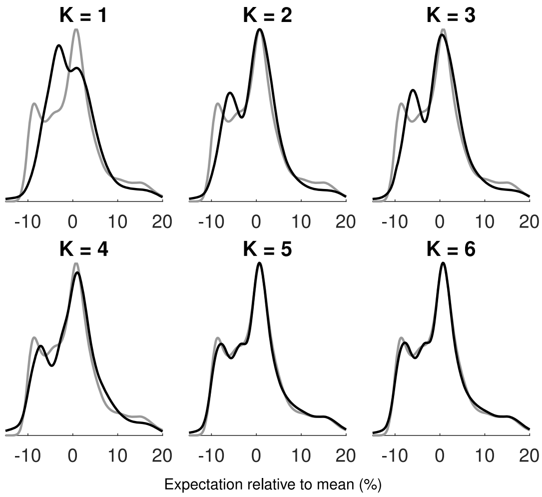

Fig. 3. Density function expansion in the empirical basis for a selected date. *Note:* Michigan survey, September 1979. Grey–observed distribution of expectations. Black–approximation given by $\widehat{p}^{(K)}_{1979\text{-}Q3,4}$, $K = 1, \cdots, 6$, defined in Eq. (A.1). The magnitudes of the associated integrated squared errors are $\log_{10}[ISE^{(K)}_{1979\text{-}Q3,4}] = \{-2.22, -2.71, -2.71, -3.06, -3.83, -3.84\}$.

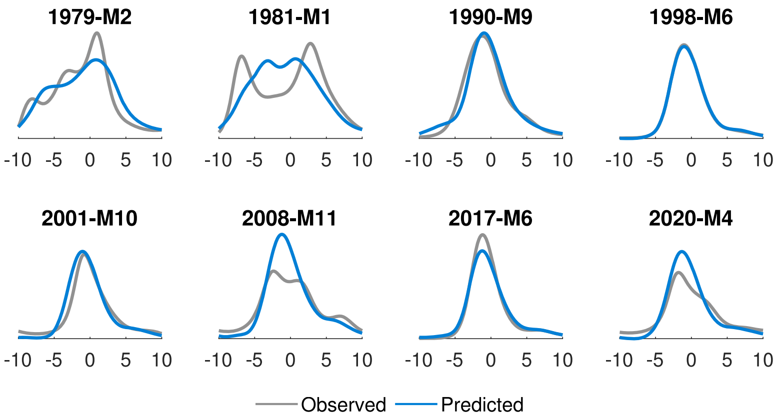

Fig. 4. Distribution of expectations predicted from prices of frequently-purchased consumption goods. *Note*: Shown actual distributions of household expectations versus predictions based on grocery and gas prices. Details of the data used in prediction may be found in Section D of the Appendix. Estimates are obtained using the functional response model displayed in Eq. (D.1).

4.5. Understanding distributional factors

In this section we suggest an approach to understanding shifts in the distribution of expectations around consensus. This aspect of our analysis is relevant to the extent that it offers clues to how policy may (or may not) be able to influence expectations. We pick up on recent findings that expectations are notably shaped by the direct experience of making purchases, and in particular that the prices of frequently purchased items have a strong influence on expectations about aggregate inflation (D'Acunto et al., 2021). We complement that literature by formulating a test of the frequency hypothesis that is based on information contained in the cross-section of expectations and on standard price indices. This allows us to examine a far longer time series of macro expectations than used in previous work.

We turn to a functional response model to determine the extent to which frequently-purchased items influence the cross-section distribution of expectations about overall inflation. We take the dependent variable to be the distribution of beliefs around consensus (p^c) observed at monthly frequency. The (scalar) predictors are the price indices for a range of grocery items, including perishable goods such as fresh fruit and vegetables, eggs, and milk, and of gasoline relative to the overall consumer price index. We include three lags of seven relative price series as predictors.[17]

Our analysis indicates that changes in the relative prices of frequently-purchased items can, in part, explain shifts in the cross-section distribution of household expectations. Over the full sample 1978–2022, the R^2 statistic for our model is 0.60; and the null hypothesis of no effect from the predictors is decisively rejected (the functional response test described in Shen and Faraway, 2004 Theorem 1, has a *p*-value of zero). Fig. 4 illustrates cross-section fit for selected dates. We observe that grocery and gas prices predicted expectations well in the late 1970s and early 1980s, for example—a period of high food price inflation—but at times of generally disrupted conditions, other drivers appear to be more germane—for example, at the start of the pandemic. From the perspective of policy, little can be done about changes in the relative prices of grocery items. But policy may have influence on residual variation in expectations, especially if (as seems likely) these depend upon more general macro conditions.

5. The role of heterogeneous beliefs in the Phillips curve

We now turn to the empirical assessment of our heterogeneous beliefs Phillips curve reported in Eq. (8). All the specifications we estimate measure inflation by the annualized percent change in the seasonally-adjusted CPI; slack by the CBO measure of the unemployment gap; the consensus expectation by the average one-year-ahead expected inflation rate from the Michigan survey; and supply factors by the change in oil prices.[18]

5.1. Estimates of the baseline heterogeneous beliefs model

As a prologue to our principal new results, we first confirm the findings documented in other studies. Table 1 (Cols. 1 and 5) reports an estimate for the slope of the Phillips curve ($\widehat{\alpha}$) of −0.2, meaning that slack of 1% is expected to reduce inflation by 0.2ppt per quarter when expressed at an annual rate.[19] Average expected inflation is estimated to be an important determinant of actual inflation, and continues to matter in the popular ‘hybrid’ Phillips curve (Cols. 3 and 7) in which past rates of inflation also appear in the model. These results are comparable to those in Coibion et al. (2018), for example, as they are based on an equivalent specification and a modestly extended sample.

Does the survey average adequately summarize the expected inflation effect? Our results suggest not. Table 1 (Cols. 2 and 6) offers strong empirical evidence that shifts in the distribution of beliefs around consensus, summarized via the factors set out in Section 4, affect actual inflation. Estimates of Eq. (8) with three components of the distribution of expectations included produce a T_F-statistic with a *p*-value of less than 0.1%, strongly rejecting the mean-only restriction. At the same time, the estimated coefficient on the consensus expectation remains highly significant. Notably, the hypothesis that inflation moves one-for-one with the consensus expectation ($\beta = 1$), a theoretical restriction often imposed in empirical work, cannot be rejected for the heterogeneous belief models (Cols. 2 and 6). Yet it is rejected in the consensus models (Cols. 1 and 5).[20]

The hybrid Phillips curve specification—where both expectations of future inflation and lags of past inflation enter jointly—proves not to be robust to how expectations are modelled. Allowing for heterogeneity eliminates the backward-looking component from the regression (Cols. 4 and 8), while the T_F-statistic indicates that the influence of the distribution of expectations is statistically different from zero. We conclude that the forward-looking Phillips curve is supported by the data, once expectational components are correctly specified.

Some notable insights into the behavior of inflation following the Great Recession of 2007-9 can be gained through the lens of the heterogeneous beliefs model. This episode provides a useful illustration of the potential for misreading expectations in real time, when heterogeneity is ignored. We re-estimated the model on pre-recession data, and contrast the predictions of the baseline model with two alternatives designed to shed light on the mechanics of the model (Fig. 5, panel a).[21] In the first exercise we keep distributional factors constant at their pre-recession (2007:Q3) values. The consensus expectation is allowed to vary. In this case, inflation is predicted to be

[17] Further information on the specification and results of this analysis appear in Section D of the Appendix. We determined that the relative price specification is superior to an alternative based on price inflation using the statistical procedure of Shen and Faraway (2004), but note that it would be equally possible to model uncentered distributions, and in that case rates of inflation might be the preferred predictor.

[18] 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. We include distributed lags in the supply factors in our regressions, and eliminate those variables/lags that are statistically insignificant. This leads us to retain only the contemporaneous change in the oil price (quarterly data) or the current and lag change (monthly data).

[19] 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.

[20] We report estimates of the heterogeneous beliefs model with Survey of Professional Forecasters data in Appendix Section H. Our selection criterion approach leads the first principal component of the distribution of SPF point forecasts to appear in the regression, and as with the Michigan survey data it proves to be statistically significant.

[21] We impose the coefficient on consensus expectations to be unity, a restriction easily accepted by the data; the point estimate on the survey mean is 1.27 on the 1978-Q1 to 2007-Q3 sample, with a robust standard error of 0.18.

Table 1
Baseline heterogeneous beliefs Phillips curve.

	Quarterly				Monthly			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable	CPIQ	CPIQ	CPIQ	CPIQ	CPIM	CPIM	CPIM	CPIM
Unemployment gap	− .211*** (.065)	− .337*** (.063)	− .145*** (.058)	− .291*** (.064)	− .185*** (.066)	− .365*** (.069)	− .165*** (.059)	− .368*** (.069)
Average expectation	1.62*** (.083)	1.23*** (.138)	.959*** (.146)	1.05*** (.154)	1.60*** (.089)	1.13*** (.112)	1.42*** (.129)	1.02*** (.126)
Lagged inflation†	–	–	.400 [.000]	.106 [.235]	–	–	.115 [.023]	.038 [.430]
Number of FPCs	0	3	0	3	0	3	0	2
T_F-statistic	–	19.8 [.000]	–	8.1 [.000]	–	14.2 [.000]	–	21.8 [.000]
Supply factors	y	y	y	y	y	y	y	y
Outlier dummy	y	y	y	y	y	y	y	y
Sample	Q1 1978–Q1 2022				Jan. 1978–Mar. 2022			
R^2	.827	.872	.860	.881	.687	.711	.725	.733
BIC	.655	.441	.560	.490	1.63	1.59	1.55	1.54
AIC	.565	.297	.399	.275	1.57	1.50	1.46	1.44
Number of obs.	177	177	177	177	531	531	531	531

Notes: Estimates of Eq. (8). Dependent variable is the annualized quarter-on-quarter percentage change (CPIQ) or month-on-month percentage change (CPIM) in the seasonally adjusted consumer price index. Newey-West adjusted (5 lags/13 lags resp.) standard errors for t tests appear in parentheses. Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. p-values for the T_F-statistic appear in brackets. All regressions include a constant. Outlier dummies equal 1 in 2008-Q4 (Quarterly) and in Sept. 2005 and Nov. 2008 (Monthly). Supply factors are the quarterly percentage change in the oil price. AIC is the Akaike Information Criterion. BIC is the Bayes Information Criterion. † Lagged inflation: sum of coefficients on four lags of the dependent variable is reported; p-value of Wald statistic for sum equal to zero in brackets.

lower by an average of 1.7ppt during the four quarters 2008:Q3-2009:Q2 than in the baseline. In the second exercise we allow distributions of expectations to shift in line with observation, while holding the consensus expectation constant. The predicted path for inflation lies notably higher than baseline from 2009:Q1 on, and the gap between the two counterfactual paths is as great as 3.6ppt, averaging 2.8ppt in 2009.

The explanation for the patterns just described can be found in Fig. 5, panel b. The consensus expectation ($\times$ symbols) rose through the summer of 2008, as oil prices surged to above \$140 per barrel, but subsequently fell back by more than 4ppt following the September collapse of Lehman Brothers, an investment bank, which marked the onset of a severe phase in the financial crisis. Crucially, though, the downward shift in the consensus expectation coincides with an upward shift in their distributional mass (blue lines). The right tails of the distributions, as approximated by their first three functional principal components, are seen to be longer and thicker over this period than in 2007:Q3 (grey lines); as noted in the Introduction, agents were not all certain at the time that the effects of the crisis and the unconventional (and at the time, unprecedented in the US) monetary policy measures would result in low and stable inflation. The distributional shift therefore 'undoes' some of the shift in the average, putting upward pressure on inflation. Thus the constant-distribution prediction is for lower inflation, while the constant-mean prediction is for higher.

5.2. *Estimates of a multiple-horizon model*

A recent literature (see for example Cecchetti et al., 2017) has questioned whether the importance of expected inflation identified in the type of models studied in Section 5.1 is down to an association between short-run expectations and trend, leading some studies to formulate the Phillips curve entirely in terms of longer-run expectations (Hooper et al., 2020). To investigate this issue, we modify the baseline heterogeneous beliefs Phillips curve Eq. (8) in two ways. First, we introduce the trend component of inflation as measured by the average 5-to-10 year ahead Michigan inflation forecast $\overline{\tau}^e_t$.[22] Second, we allow the distribution of long-run beliefs to enter the Phillips curve alongside the distribution of short-run expectations. These modifications generalize the FLM of Eq. (8) to accommodate multiple functional predictors. Estimation is again via FPCR, with functional principal components of both distributions entering:

$$\pi_t - \overline{\tau}^e_t = \alpha\varphi_t + \beta\left(\overline{\pi}^e_t - \overline{\tau}^e_t\right) + \sum_h \sum_{k_h} \gamma_{h,k} s_{h,k,t} + \varepsilon_t \tag{10}$$

for two horizons $h = \{1, 5 \text{ to } 10\}$, and where the difference between inflation and its trend will be referred to as the 'inflation gap'. The number of included components k_h is allowed to differ by horizon. The coefficient β measures the relative weight on short- versus long-

[22] Data on households' longer-horizon expectations is available from the start of the 1990s as a continuous series. Section J of the Appendix shows that the results are robust to using different measures of long-run inflation expectations. The Appendix also contains a summary of the statistical properties of the distributions of long-run Michigan expectations, which are qualitatively similar to the properties of short-run expectations.

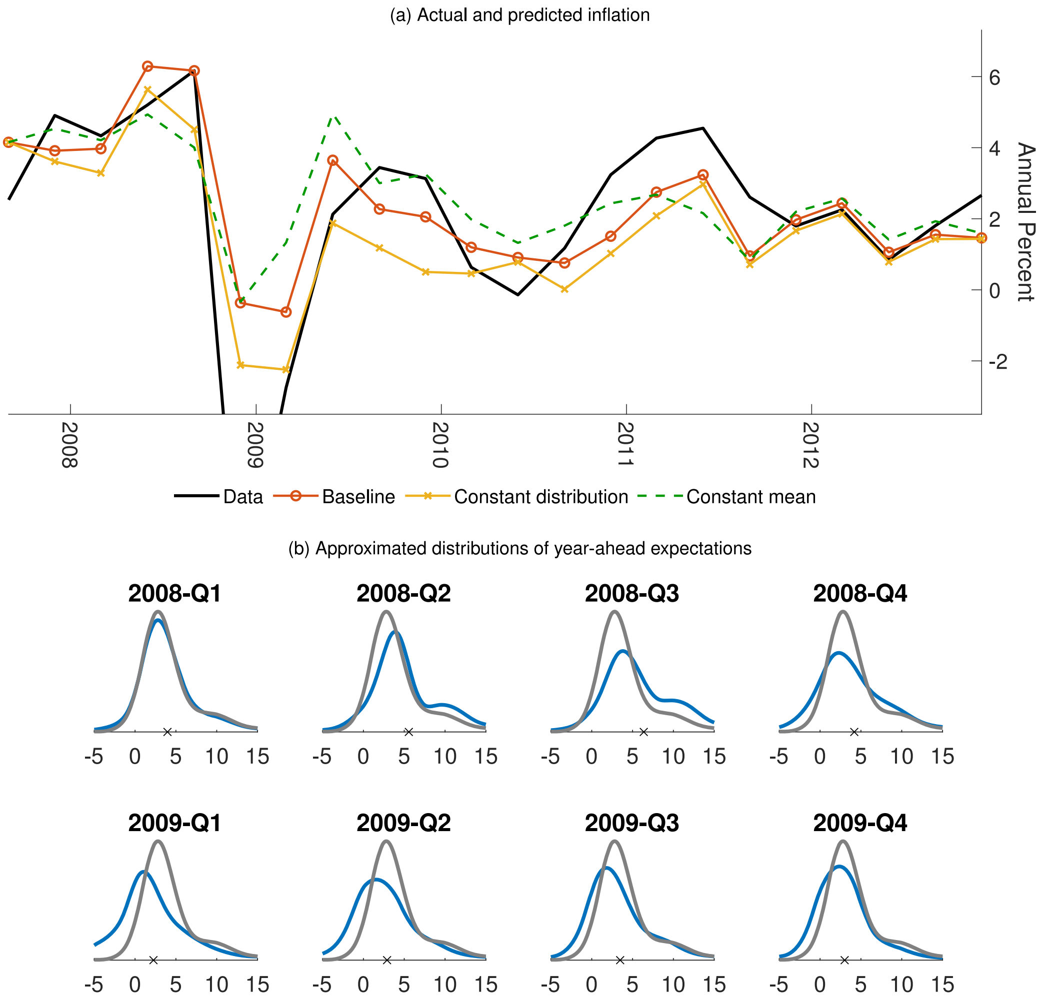

Fig. 5. Heterogeneous expectations and inflation in the Great Recession. *Notes* Panel (a): (i) Data–the quarter-on-quarter annualized percentage change in the seasonally-adjusted CPI; (ii) Baseline–predicted inflation from the heterogeneous beliefs model Eq. (8) with $\beta = 1$ and $K = 3$; (iii) Constant distribution–predicted inflation from the heterogeneous beliefs model when the distributional components $s_{k,t}$ are set to their 2007:Q3 values; (iv) Constant mean–predicted inflation when the consensus expectation is set to its 2007:Q3 value. All estimates are based on data from 1978:Q1 to 2007:Q3. Panel (b): Functional approximations with $K = 3$ in blue plotted against the distribution for 2007:Q3 in grey. The $\times$ symbol indicates the average expectation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

run average expectations: $\beta = 0$ implies inflation moves one-for-one with trend inflation, and does not depend on near-term average expectations; $\beta = 1$ implies a Phillips curve in which inflation is influenced by long-run expectations only to the extent that the distribution of long-run expectations is significant in the regression.

Distributions of both short- and long-run expectations matter for inflation. Table 2 reports estimates of (10) when adding in turn the distributions of 1 year ahead forecasts (Cols. 1 and 5, for the quarterly and monthly regressions respectively), the 5–10 year ahead forecasts (Cols. 2 and 6), and both distributions at the same time (Cols. 3 and 7). Once again, the choice of principal components is guided by the information criteria. Both when using quarterly and monthly data, the functional principal components are significant in all specifications, including where both short- and long-run distributions are included.

This result suggests that the attention given to trend inflation, here measured using beliefs about the long-run, is half right. Inflation is driven by short-run and not by long-run consensus expectations, as seen from the overwhelming acceptance of the null hypothesis that $\beta = 1$ in all of our quarterly models. But although the long-run consensus is of limited relevance to near-term dynamics, long-run beliefs lying away from the consensus do impact inflation, as seen (a) from the significant T_F-statistics when they enter singly (Table 2, Cols. 2 and 6); and (b) from the fact that they are selected by the AIC and BIC when entering joint with the distribution of near-term expectations (Cols. 3 and 7). Imposing the restriction that $\beta = 1$ (Cols. 4 and 8) produces models comparable to those reported in Table 1: Inflation depends on consensus beliefs about short-horizon inflation, but with the additional feature that long-run distributions also impact inflation dynamics.

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

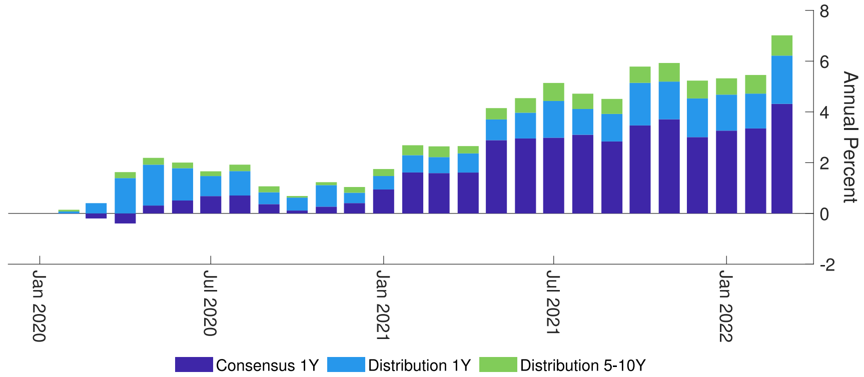

Fig. 6. Predicted effect of household expectations on inflation during the economic recovery from the pandemic. *Note:* Chart shows the predicted effect of the named components of expectations to month-on-month CPI inflation (annual rate). Estimates are as reported in Table 2, Col. 8.

Table 2
Inflation gaps and heterogeneous beliefs.

	Quarterly				Monthly			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent variable	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$	$\pi_t - \overline{\tau}_t^e$
Unemployment gap	−.312*** (.064)	−.294*** (.054)	−.277*** (.065)	−.301*** (.058)	−.353*** (.071)	−.233*** (.069)	−.340*** (.080)	−.316*** (.080)
$\overline{\pi}_t^e - \overline{\tau}_t^e$	.984*** (.590)	.828*** (.174)	1.16*** (.175)	1.0 –	.700*** (.163)	1.02*** (.243)	.758*** (.166)	1.0 –
# short-horizon[(d)]	2	-	2	2	1	-	1	1
# long-horizon[(d)]	-	2	1	1	-	1	1	1
T_F-statistic	17.1 [.000]	13.2 [.000]	12.4 [.000]	12.4 [.000]	28.5 [.000]	19.7 [.000]	14.4 [.000]	13.6 [.000]
Wald-statistic[(e)]	−.092 [.464]	−.992 [.162]	.911 [.818]	–	−1.84 [.034]	.090 [.537]	−1.46 [.073]	–
Supply factors	y	y	y	y	y	y	y	y
Outlier dummy	y	y	y	y	y	y	y	y
Sample	Q2 1990–Q1 2022				Apr. 1990–Mar. 2022			
R^2	.777	.765	.781	.737	.637	.629	.637	.583
BIC	.236	.288	.254	.221	1.44	1.46	1.46	1.45
AIC	.080	.132	.076	.065	1.35	1.37	1.35	1.36
Number of obs.	128	128	128	128	384	384	384	384

Notes: (a) Estimates of Eq. (10). Dependent variable is the annualized quarter-on-quarter percentage change (CPIQ) or month-on-month percentage change (CPIM) in the seasonally adjusted consumer price index. Newey-West adjusted (5 lags/13 lags resp.) standard errors for t tests appear in parentheses. (b) Asterisks denote significance at the 10% (*), 5% (**), and 1% (***) levels. p-values for the T_F-statistic appear in brackets. All regressions include a constant. Outlier dummies equal 1 in 2008-Q4 (Quarterly) and in Sept. 2005 and Nov. 2008 (Monthly). Supply factors are the quarterly percentage change in the oil price. (c) AIC is the Akaike Information Criterion. BIC is the Bayes Information Criterion. (d) Number of included functional principal components of 1-year-ahead/5-10 year ahead distribution (resp.).(e) Test statistic for $H_0 : \beta = 1$ vs. the alternative $H_1 : \beta < 1$

since the global financial crisis. A number of factors underlay its movements: the 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' inflation expectations shifted notably in that same period (Reis, 2021).

The distribution of longer-term expectations started to have a material effect on inflation in mid-2021, when policymakers generally perceived the rise in inflation as being transitory. Seen through the lens of our heterogeneous beliefs model, rising household expectations of inflation accounted for around 7 percentage points (ppt) of the run-up in inflation over the pandemic period, according to our estimates. Fig. 6 shows the regression 'accounting'. The change (relative to January 2020) in the estimated contributions from each of three expectational elements to predicted inflation are shown as stacked bars. The pick-up in near-term consensus expectations mechanically makes up a little over 4ppt of the total. The remaining increase is accounted for by changes in the distributions of expectations. Evidently, not all agents shared the belief that rising inflation would be transitory, and the heterogeneous beliefs model picks up the effect of this disagreement. This suggests that variants of our model could prove to be valuable as a tool to monitor and interpret developments in longer-run expectations.

5.3. Regression on distributional moments

In some familiar circumstances, a probability distribution can be determined from knowledge of its moments.[23] 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 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.[24] In our case, the models specified in Table 1 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 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 model A, 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 model B, the moments-based model (*p*-value of zero). These results provide clear evidence in favor of the FPC model—model A encompasses, but is not encompassed by, model B.

6. 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 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. Our paper offers answers to these problems, and in so doing sheds new light on the role that expectations play in inflation dynamics.

The paper sets out the statistical properties associated with long time series of cross-section distributions of survey forecasts, and proposes a framework for structuring and using the information contained therein. We then propose a straightforward approach to estimation and testing in the presence of heterogeneous beliefs, using the tools of functional data analysis. Our paper establishes both a statistical association between distributions of expectations and inflation, and their quantitative salience. Neither finding has previously been recognized.

In future work, we see the need to develop models that can account jointly for time-varying distributions of beliefs and for the influence of those distributions on inflation. Such an approach could contribute to refining our understanding of the transmission channels underlying the results we have reported.

Data availability

Data will be made available on request.

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 Bank of England, Ca' Foscari University of Venice, the CEPR Paris Symposium, the European Central Bank, the FRB St. Louis Applied Time Series Econometrics workshop, 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 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.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2023.06.003

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

[24] 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.

References

Adam, K., Nagel, S., 2023. Chapter 16 - Expectations data in asset pricing. In: Bachmann, R., Topa, G., van der Klaauw, W. (Eds.), Handbook of Economic Expectations. Academic Press, pp. 477–506.
Adam, K., Padula, M., 2010. Inflation dynamics and subjective expectations in the United States. Econ. Inq. 49 (1), 13–25.
Ang, A., Bekaert, G., Wei, M., 2007. Do macro variables, asset markets, or surveys forecast inflation better? J. Monet. Econ. 54, 1163–1212.
Attanasio, O., Kovacs, A., Molnar, K., 2020. Euler equations, subjective expectations and income shocks. Economica 87 (346), 406–441.
Binder, C.C., 2017. Measuring uncertainty based on rounding: new method and application to inflation expectations. J. Monet. Econ. 90, 1–12.
Bowsher, C.G., Meeks, R., 2008. The dynamics of economic functions: modeling and forecasting the yield curve. J. Am. Stat. Assoc. 103 (484), 1419–1437.
Branch, W.A., McGough, B., 2009. A New Keynesian model with heterogeneous expectations. J. Econ. Dyn. Control 33, 1036–1051.
Cecchetti, S.G., Feroli, M.E., Hooper, P., Kashyap, A.K., Schoenholtz, K.L., 2017. Deflating inflation expectations: the implications of inflation's simple dynamics. Technical Report. University of Chicago Booth School of Business.
Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell us about information rigidities? J. Polit. Econ. 120 (1), 116–159.
Coibion, O., Gorodnichenko, Y., Kamdar, R., 2018. The formation of expectations, inflation and the Phillips curve. J. Econ. Lit. 56 (4), 1447–1491.
Crump, R.K., Eusepi, S., Tambalotti, A., Topa, G., 2022. Subjective intertemporal substitution. J. Monet. Econ. 126, 118–133.
Cuevas, A., 2014. A partial overview of the theory of statistics with functional data. J. Stat. Plan. Inference 147, 1–23.
D'Acunto, F., Malmendier, U., Ospina, J., Weber, M., 2021. Exposure to grocery prices and inflation expectations. J. Polit. Econ. 129 (5).
Davidson, R., MacKinnon, J.G., 1981. Several tests for model specification in the presence of alternative hypotheses. Econometrica 49 (3), 781–793.
Del Negro, M., Lenza, M., Primiceri, G.E., Tambalotti, A., 2020. What's up with the Phillips curve? Brookings Pap. Econ. Act. Spring, 301–357.
Harris, M., Raviv, A., 1993. Differences of opinion make a horse race. Rev. Financ. Stud. 6 (3), 473–506.
Hooper, P., Mishkin, F.S., Sufi, A., 2020. Prospects for inflation in a high pressure economy: is the Phillips curve dead or just hibernating? Res. Econ. 74, 26–62.
Horváth, L., Kokoszka, P., 2012. Inference for Functional Data with Applications. Springer Series in Statistics. Springer-Verlag, New York.
Jolliffe, I.T., 2002. Principal Component Analysis. Springer Series in Statistics, 2nd. Springer-Verlag, New York.
Kneip, A., Utikal, K.J., 2001. Inference for density families using functional principal component analysis. J. Am. Stat. Assoc. 96 (454), 519–532.
Kong, D., Staicu, A.-M., Maity, A., 2016. Classical testing in functional linear models. J. Nonparametr. Stat. 28 (4), 813–838.
López-Pintado, S., Romo, J., 2009. On the concept of depth for functional data. J. Am. Stat. Assoc. 104 (486), 718–734.
Mankiw, N.G., Reis, R., Wolfers, J., 2003. Disagreement about inflation expectations. In: Gertler, M., Rogoff, K. (Eds.), NBER Macroeconomics Annual, Vol. 18. University of Chicago Press, pp. 209–248.
Mavroeidis, S., Plagborg-Møller, M., Stock, J.H., 2014. Empirical evidence on inflation expectations in the new keynesian Phillips curve. J. Econ. Lit. 52 (1), 124–188.
Meeks, R., Monti, F., 2019. Heterogeneous beliefs and the Phillips curve. Staff Working Paper. No. 807. Bank of England. https://www.bankofengland.co.uk/working-paper/2019/heterogeneous-beliefs-and-the-phillips-curve.
Ramsay, J.O., Silverman, B.W., 2005. Functional data Analysis. Springer Series in Statistics, 2nd. Springer-Verlag, New York.
Reis, R., 2021. Losing the inflation anchor. Brookings Pap. Econ. Act. 307–361.
Reiss, P.T., Goldsmith, J., Shang, H.L., Ogden, R.T., 2017. Methods for scalar-on-function regression. Int. Stat. Rev. 85 (2), 228–249.
Reiss, P.T., Ogden, R.T., 2007. Functional principal component regression and functional partial least squares. J. Am. Stat. Assoc. 102 (479), 984–996.
Rich, R., Tracy, J., 2010. The relationships among expected inflation, disagreement and uncertainty: evidence from matched point and density forecasts. Rev. Econ. Stat. 92 (1), 200–207.
Roberts, J.M., 1995. New Keynesian economics and the Phillips curve. J. Money Credit Bank. 27 (4), 975–984.
Shen, Q., Faraway, J., 2004. An *F* test for linear models with functional responses. Stat. Sin. 14, 1239–1257.
Shin, H., 2009. Partial functional linear regression. J. Stat. Plan. Inference 139, 3405–3418.
Tsay, R.S., 2016. Some methods for analysing big dependent data. J. Bus. Econ. Stat. 34 (4), 673–688.
Tsiaplias, S., 2020. Time-varying consumer disagreement and future inflation. J. Econ. Dyn. Control 116, 1–20.