Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Monitoring Macroeconomic Prospects with a Meta VAR-E Dashboard

CAMA Working Paper 10/2026
February 2026

Kevin Lee
University of Nottingham
Centre for Applied Macroeconomic Analysis, ANU

Kalvinder Shields
University of Melbourne
Centre for Applied Macroeconomic Analysis, ANU

Abstract

The time series properties of output and price inflation can be accurately captured using VAR-E's, Vector-Autoregressive models of actual and expected measures of the series where the latter are provided by surveys. The paper proposes a method for estimating VAR-E's that accommodate individuals' real-time understanding of the macroeconomy and which deliver forecasts in a way that is useful to decision-makers. It notes the sort of statistics and figures that might be reported in a 'dashboard' to monitor the health of the macroeconomy, and this is illustrated using the actual and expected data produced by the Bank of England's Decision-Maker Panel.

Keywords

learning, expectations, surveys, forecasts, decision-making

JEL Classification

C32, D84, E31, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Monitoring Macroeconomic Prospects with a Meta VAR-E Dashboard*

Kevin Lee† and Kalvinder Shields††

December 2025

Abstract

The time series properties of output and price inflation can be accurately captured using VAR-E's, Vector-Autoregressive models of actual and expected measures of the series where the latter are provided by surveys. The paper proposes a method for estimating VAR-E's that accommodate individuals' real-time understanding of the macroeconomy and which deliver forecasts in a way that is useful to decision-makers. It notes the sort of statistics and figures that might be reported in a 'dashboard' to monitor the health of the macroeconomy, and this is illustrated using the actual and expected data produced by the Bank of England's Decision-Maker Panel.

Keywords: Learning, Expectations, Surveys, Forecasts, Decision-Making.
JEL Classification: C32, D84, E31, E32

*† University of Nottingham, UK; †† University of Melbourne, Australia. Version dated December 2025. Paper prepared for CFE-CMStatistics 2025 Conference sessions in honour of Hashem Pesaran. Financial support from the Australian Research Council (Discovery Grant DP230100959) is gratefully acknowledged. Corresponding author: Kevin Lee, School of Economics, University of Nottingham, Nottingham NG7 2RD. E-mail: kevin.lee@nottingham.ac.uk.

1 Introduction

Macroeconomic forecasts are used to help policy-makers, businesses and households make decisions but the act of model-building and forecasting is frequently called into question, citing flawed theoretical foundations, difficulties in choosing the variables to be included in the model and its specification, time-variation in the underlying structures, and so on. In this paper, we present a Vector Autoregression model involving the data on the actual and expected values of output and inflation (a 'VAR-E model'), obtained using a 'meta-modelling' approach to estimation, and use this to produce a limited set of forecasts - in the form of a 'dashboard' - that might be used to judge the state of the economy in real-time. We argue that, despite the the complexity of the macroeconomy, the difficulties in macromodelling, the challenges in conveying the uncertainties surrounding forecasts, and the variety of reasons why individuals are interested in the macroeconomy, it is possible to estimate a model and produce a dashboard that is useful and credible/reliable for a broad range of decision makers. And we illustrate the approach by producing a dashboard for the UK economy, constructed at the time of writing, based on the expectations data produced by the Bank of England's Decision Maker Panel (DMP) survey.

We judge our approach following the general framework of Pesaran and Smith (1985) which proposes three criteria to evaluate the success of a macromodelling exercise: its relevance, consistency and adequacy. *'Relevance'* considers the extent to which the exercise satisfies the purpose for which it is designed. In this case, we are interested in delivering forecasts in a form that is useful for a wide range of decision-makers and our approach acknowledges the ways in which decision-makers use information - often in the form of 'rule-of-thumb' heuristics - and proposes an appropriate way to report model-based forecasts. *'Consistency'* refers to the requirement that the model's properties are compatible with known physical, institutional or financial constraints. To some extent, our approach side-steps this criterion, as it is based on a VAR and can rely on the Wold representation theorem - interpreted as saying that every stationary time series can be approximated by an autoregressive one - as a justification. But our approach specifically involves a VAR in the actual and expected values of a variable and this pairing has a number of features, elaborated below, which ensure the model is consistent with a reasonable view of functioning of the macroeconomy. *'Adequacy'* refers to the acceptability of the modelling in statistical terms. The use of a VAR-E means econometric problems are less likely to arise from choice of variables, treatment of endogeneity/exogeneity, and so on. But there is considerable time-variation in macroeconomic relationships, generated not least by individuals' changing understanding of the fundamental nature of the relationships and the shocks to them, and this needs to be addressed in any macroeconomic forecasting exercise. Our approach does this directly through out meta modelling estimation method.

The remainder of the paper is set out as follows. Section 2 describes the modelling framework underlying our work, elaborating the advantages of VAR-E models in this context, explaining the meta-modelling estimation method, and commenting on the ways in which forecasts are used in decision-making and the implications for presenting forecasts. Section 3 applies the approach to the data on UK sales growth and price inflation obtained from the Bank of England's DMP. Section 4 concludes.

2 Forecasting with a Meta VAR-E Model

Our empirical work is based on bivariate VAR-E models explaining, in turn, actual and expected price inflation, and actual and expected output growth. The annual inflation rate in month t is $\pi_t = p_t - p_{t-12}$, where p_t is (the logarithm of) the price level in t, and the expected one-year-ahead annual inflation rate, as published in a survey dated in month t, is denoted ${}_t\pi_{t+12} = {}_tp_{t+12} - p_t$. Similar expressions can be written for output so that, for example, output growth $g_t = y_t - y_{t-12}$ where y_t is (the logarithm of) output in t. In We assume that both prices and output are difference-stationary series. Interest focuses on the difference-stationary inflation series π_t for prices while, for output, both the growth and the level y_t (and its relationship with trend output) are of interest.

Assuming a first order process for exposition purposes, the basic model explaining, for example, the observations on inflation dated at $t = 1, ..., T$ is given by

$$\begin{bmatrix} \pi_t \\ {}_t\pi_{t+12} \end{bmatrix} = \alpha + \beta \begin{bmatrix} \pi_{t-1} \\ {}_{t-1}\pi_{t+11} \end{bmatrix} + \begin{bmatrix} \varepsilon_t^a \\ \varepsilon_t^e \end{bmatrix} \tag{1}$$

where α and β are 2×1 and 2×2 parameter matrices, $\varepsilon_t = (\varepsilon_t^a, \varepsilon_t^e)' \sim N(0, \mathbf{\Sigma})$, and $\mathbf{\Sigma}$ is the 2×2 matrix describing the uncertainty surrounding these residual shocks. Note that, if re-written in terms of Δp_t and $\Delta_t p_{t+12}$, the model in (1) takes the form of a cointegrating VAR, effectively (and sensibly) imposing convergence between the actual and expected price level in the long run. Of course, similar comments apply to output levels in the VAR-E of actual and expected output growth.

2.1 The benefits of the bivariate VAR-E

The VAR-E model has a number of features that make it eminently useful for characterising and providing forecasts for output growth and inflation (and other macro variables). The statistical features include:

- Parsimoneous: the bivariate model is obviously very concise in its characterisation of the variable of interest. This makes it very easy to work with, to summarise and describe its properties and to use in simulation and counter-factual exercises. The treatment of actual and expected inflation and growth as stationary variables builds in the assumption that neither inflation, nor growth, nor expectational errors grow without bounds, consistent with reasonable views on the functioning of the macroeconomy.

- Adaptable: given its simplicity, it is relatively easy to adapt the VAR-E to accommodate extensions to the modelling framework. Two extensions come readily to mind. First, the bivariate VAR can easily allow for time-variation in its parameters and for this time variation to be conveyed simply. As elaborated below, this is important where individuals experience structural change in the macroeconomy. Second, the bivariate VAR can be easily embedded in a global context if there are comparable VAR-E's available for m, say, inter-related economies. The inclusion of lagged values of global inflation and global expected inflation (measured as a weighted average of inflation in the other economies) as additional regressors in each country's VAR-E allows the models to be stacked and re-written as a $2m$-variable VAR to produce a simple characterisation of the global system.[1]

- Meaningful trends: the VAR-E can be used to generate the estimated unconditional means of the actual and expected inflation and growth series expected to be achieved over the long-run; e.g. $[\tilde{\pi}_t,\ {}_t\tilde{\pi}_{t+12}]' = (\mathbf{I_2} - \beta)^{-\mathbf{1}}\alpha$ in (1). For output, we are also interested in the long-horizon effects of shocks on the level shown by the Beveridge-Nelson (BN) trend. This captures the accumulated effects of the shocks and measures the value to which output will converge in the absence of any future shocks and once the effects of shocks that have been experienced in the past have worked themselves through. The 'output gap' between current output level and the BN trend is a key macroeconomic indicator showing the extent to which the economy is operating above or below capacity.

- Flexible dynamics: The bivariate VAR of order 1 given in (1) can capture relatively complex dynamics. To see this, note that (1) can be re-written in univariate terms for either one of the variables, substituting out the successive lagged values of the expected variable to obtain an expression in terms of the actual values only for example. In the case of this bivariate VAR(1), the corresponding univariate model for the actual variable would be an ARMA(2,1) with the constructed shocks derived as a function of the two shocks $\varepsilon^a_{T,t}$ and $\varepsilon^e_{T,t}$. A univariate AR model would require an infinite number of lags to reproduce the properties of the ARMA model and any finite order AR model would have to be an approximation of the ARMA specification. Moreover, this approximation could be very misleading: for example, a 1% shock to output characterised by the AR(1) model $\Delta y_t = \rho \Delta y_{t-1} + \varepsilon_t$ will cause y_t to be $\frac{1}{1-\rho}$% higher in the long-run than in the absence of the shock; if the model is stable, so that $-1 < \rho < 1$, the measured persistence always exceeds $\frac{1}{2}$ even if the variable is close to stationary in levels (with the MA parameter close to -1 in the ARMA(2,1) specification. A univariate AR specification would usually require many lags to approximate a variable that is close-to-stationary.

[1] The stacking of systems in this way follows the GVAR methodology popularised in the di Mauro and Pesaran (2013) text and illustrated in the context of output growth in Garratt et al. (2018).

- Accommodating permanent and transitory shocks: The statistical flexibility of the bivariate model has clear benefits when a variable experiences both permanent and transitory shocks (as is the case for output levels, for example). Obviously, a univariate model will not be able to properly distinguish the effects of the two types of shock and, in attempting to characterise the series as a function of the constructed shocks, it will overstate the persistence of the genuinely-permanent shocks. The two types of shock can be identified explicitly in the bivariate model, following the procedure of Blanchard and Quah (1989) and they provide a direct measure of the signal:noise ratio of interest when considering the size of permanent and transitory shocks in the business cycle, for example.[2]

The use of direct measures of expectations means the VAR-E has the further advantage that it can exploit the information known by individuals - as reflected by their survey responses - and investigate the role of expectation formation in macro dynamics:

- Comprehensive: the use of direct measures of expectations in the bivariate VAR means that, despite its size, it provides a comprehensive characterisation of the influences on a variable. For example, if expectations are formed according to FIRE, the fit of the model for the variable could not be improved by including any other variable as the (appropriately lagged) expectation would summarise fully the influences on the variable from the rest of the economy. And, in so far as expectations are formed relatively autonomously, providing their own impetus to the variable beyond that suggested by RE, the direct measure of expectation will pick up those influences too.[3]

- Testing expectations-formation: the bivariate VAR provides a forum to consider the inter-relatedness of a variable's outcomes and their expectations. This provides an explicit way to investigate the nature of expectation formation; for example, the assumption of FIRE, or of sticky information a la Coibion and Gorodnovich, or noisy-information models, can be readily accommodated through imposed restrictions on the parameters of the VAR-E and the validity of the restrictions provides a straightforward test of the assumptions.[4]

- Interpretable: the VAR-E can often be given a straightforward structural interpretation working at an appropriate level of abstraction. For example, a structural model explaining price inflation according to a NKPC and assuming expectations are formed partly according to RE and partly autonomously, can be readily written as a VAR-E in inflation. Further, the estimated unconditional mean of expected inflation provides a natural measure of the 'reference inflation rate' around which

[2] See Lee and Mahony (2025) for further discussion.

[3] See Pesaran (1987) and Pesaran and Weale (2006) for discussion of the use of expectations data obtained from surveys in macromodelling.

[4] Garratt et al. (2016) illustrate these tests in the context of US output expectations.

long-run expectations will be based, and the sensitivity of the statistic to shocks to actual inflation provides a useful measure of the extent to which inflation is or is not anchored to this rate.[5]

- Capturing individuals' insights: The distinction between permanent and transitory shocks obtained through the Blanchard-Quah decomposition mentioned above means shocks can be given further structural interpretation in the VAR-E because the survey respondents are implicitly reporting how much of the effect of a transitory shock is expected to have dissipated at the forecast horizon. This can be very useful information when judging the persistence of shocks to inflation and the need for monetary policy reactions, for example, or in identifying economically-meaningful shocks according to whether they are short-lived or long-lived (a distinction that is often more useful than the more contentious identifying assumptions often found in the literature based on the source of shocks).[6]

2.2 Structural change and learning in a 'meta VAR-E' model

As mentioned above, an important adaptation of the bivariate VAR-E of (1) allows for time-variation in the parameters of (1), α, β and Σ. This will be important if there are structural breaks in the relationships underlying the determination of the actual and expected series and their interactions. In the case of output growth, the experiences of the GFC and the covid pandemic illustrate well the changes that can occur over time in the nature and persistence of shocks to output. And time-variation in the determinants of output growth are at the heart of the debates over the extent and causes of episodes of slowdown in productivity growth across advanced economies (relating to the introduction of digital technologies, demographic changes, the ebb and flow of globalisation and international trade, and so on). Similarly, the most recent period of high inflation - prompted by supply chain disruptions following the pandemic, and energy price increases following Russia's invasion of Ukraine - represent a clear change from the period of low and stable inflation of the 1990's and early 2000's. Such episodes could influence the frequency with which price setters change their prices and/or the frequency with which they update their information sets (in an economy with information rigidities), with consequent changes in inflation dynamics. And, depending on how the monetary authorities react, the episodes could influence individuals' views on their ability and willingness to control inflation and hence long-run inflation expectations. These changes illustrate that any model of output growth and price inflation and their expectations are likely to benefit from a flexibility that accommodates changing circumstances.

There are many approaches taken in the applied literature to deal with structural

[5] See Lee, Shields and Nguyen (2025) for further details.

[6] See Garratt et al. (2016) for a description of how BN trend measures can be purged of the effects of short-lived shocks to obtain 'news-adjusted' output gap measures. And see Lee et al. (2024) for an analysis of how the different types of shock influence business cycle dynammics.

instabilities but our preferred approach is based on a particular type of model-averaging. Here, at any point in time T, many versions of the VAR-E model of (1) are estimated based on different sample sizes $T-j, ..., T$ for $j = j_{min}, .., J$, and these are combined with weights to obtain a 'meta VAR-E model'. More specifically, denoting the VAR-E model model estimated at T for $t = T-j, ..., T$ by M_{Tj} - with corresponding parameters α_{Tj}, β_{Tj} and Σ_{Tj} - our meta model is defined by

$$\overline{\mathrm{M}}_T = [M_{Tj},\ w_{Tj} \text{ for } j = 1, ..., J] \tag{2}$$

where weights w_{Tj} evolve over time according to the following:

$$w_{Tj} \to \begin{cases} w_{T+1,j+1} \text{ if there is no evidence of a break in } M_{T+1,j+1} \text{ at } T+1 \\ w_{T+1,j-s} \text{ for } s = 0, 1, .., j+1 \text{ if there is evidence of a break.} \end{cases} \tag{3}$$

The transition of weights in (3) captures the idea that, at time T, an individual places a weight of w_{Tj} on the model estimated over $t = T-j, ..., T$. On the arrival of news in $T+1$, the individual re-estimates the model over $t = 1, ..., T+1$ and tests whether there has been a structural break (using a Chow test of predictive ability, say in the univariate case or a likelihood ratio test in multivariate model). If there is no evidence of a structural break, the weight shifts to the model $M_{T+1,j+1}$, estimated over the period with one extra observation $t = T-j, ..., T+1$. If there is evidence of a break, the weights are distributed equally over all of the models based on the statistically-preferred shorter samples $M_{T+1,j-s}$ for $s = 0, 1, .., j+1$. The properties of the model can be usefully summarised by the evolution over time of the weighted average of the parameters; e.g. $\overline{\beta}_T = \sum_{j=1}^{J} w_{Tj}\beta_{Tj}$, while the 'average duration' statistic $d_T = \sum_{j=1}^{T} j * w_{Tj}$ shows the average length of the regimes in operation at time T.[7]

An important feature of this model is its flexibility. In times of relative stability, the samples of all the underlying sub-models grow by one period per period so there is a gradual, evolution of their parameters (averaging over a set of recursively estimated models). But where there is instability, the meta model can change very abruptly if the evidence is that all of the sub-models are subject to breaks, or more incrementally if some sub-models appear untenable while there is no evidence to reject others. This latter property - arising from the averaging across models - provides a more nuanced and flexible approach to structural breaks than can be found when looking for a breaks in a single model.

Perhaps even more importantly, the meta modelling approach is able to capture an important feature of macro dynamics involving learning. Learning is the process of acquir-

[7]The advantages of the meta modelling approach in forecasting are elaborated in Pesaran and Timmerman (1995), where there is a trade-off in forecast performance between using long samples to improve the precision of parameter estimates and using short samples to minimise the risk of including breaks. But the approach is equally useful in inference as illustrated in the models of inflation, growth, and exchange rate and interest rate deternmination of Lee et al. (2015), Aristidou et al. (2022), Lee et al. (2025a,b), and Lee and Mahony (2025).

ing new understanding and skills and is particularly important in the face of structural change. The standard approach to accommodating learning in the literature is based on rational decision-making in a Bayesian context and assumes individuals make optimal use of all the information available to them. This is usually captured empirically through time-variation in parameters obtained in real time in recursively estimated models. However, in practice, individuals' perceptions might be dominated by 'learning rigidities' involving slow-to-adjust rules of thumb (heuristics) to describe their understanding of the world. Here, individuals sometimes appear to ignore relevant available information and rely on previously-formed mental models of the world. Examples of rules generating learning rigidities are 'inertia' or 'anchoring' heuristics, where individuals make choices disproportionately influenced by previously held views, 'reinforcement' heuristics where individuals maintain an unjustified consistency with successful past decisions, and 'availability' heuristics where individuals make judgements according to the ease with which relevant ideas are brought to mind and so base views on relatively extreme or significant past experiences.[8] The meta modelling approach captures this aspect of learning very well, with the weights in (3) representing an individual's degree of belief in a sub-model and with agents continuing to believe in the model unless there is significant evidence with which to abandon the belief. The 'protection' of the null in the transition tests captures the reluctance with which agents give up their existing views and reflects well the possible presence of learning rigidities in decision-making. Further, in the context of various models of inflation, tests of the meta model against alternative models of structural change have shown that this characterisation of learning - and associated macro dynamics - fits the data well.[9]

2.3 Decision-making and the macro context

The meta model provides a straightforward characterisation of individuals' understanding of the macro context in real time, capturing the interactions between the actual and expected series through the $\overline{\beta}_T$ coefficients, providing estimates of the 'steady-state' values of the series, $\overline{\widetilde{\Pi}}_T = \sum_{j=1}^{T} w_{Tj} * ([\mathbf{I}_2 - \beta_{Tj}]^{-1} * \alpha_{Tj})$ and reflecting the uncertainty surrounding the series at each point in time through $\overline{\Sigma}_T = \sum_{j=1}^{T} w_{Tj} * \Sigma_{Tj}$.

But the model also provides a powerful vehicle for generating forecasts and, more specifically, density forecasts and event probability forecasts.[10] Given the simplicity of the VAR-E, and focusing on just one sub-model M_{Tj} at first, it is straightforward to generate $s_{Tj} * H$ random draws from a bivariate Normal distribution with zero mean and covariance matrix Σ_{Tj} where $s_{Tj} = w_{Tj} * S$ is a fraction of the total number of simulations, S, we will consider in total. We can use these these random draws with the estimated

[8]See Evans and Honkepohja's (2001) text for a discussion of the Bayesian approach to learning, and Tversky and Kahneman's (1974) paper for a seminal discussion of heuristics.

[9]See Lee and Shields (2025) and Lee et al. (2025) for evidence relating to the US and Australian economies.

[10]See Garratt et al. (2003, 2006) for a detailed discussion on the use of event probability forecasting.

model to dynamically simulate s_{Tj} alternative sets of future values of $[\pi^{(j,s)}_{T+h},\ _{T+h}\pi^{(j,s)}_{T+12+h}]'$ for $h = 1, ..., H$ and $s = 1, ..., s_{Tj}$. If interest was solely on this sub-model, the range of values taken by the simulated sets of values at $t + h$ provides a direct representation of the forecast density for actual and expected inflation at $t + h$ for that model. And, for any (possible complicated) event defined by the future values - e.g. that actual inflation falls below 3% in each of the next six months and below 2% for the following six months - the estimated forecast probability of the event is simply the fraction of times out of the s_{Tj} simulations that the event is observed to occur. But, of course, this procedure can be repeated for each of the sub-models in the meta model $\overline{M}_T$, providing S simulated futures in total. Again, the range of values taken across the simulations describes the density forecast at each horizon while the fraction of times an event occurs shows the event probability forecast, but now the simulated futures take into account the full range of sub-models, with corresponding weights, that are considered relevant at the date the forecasts are made.[11]

There remains a question, however, of how the forecasts produced by the model should be presented to be most useful for individuals making decisions in real time.[12] In some situations, the simulated futures provided by the meta VAR-E model provide all the contextual information that is required for a fully informed individual to make decisions. A decision maker will typically want to maximise a payoff that depends on the macro context at some point or points in the future as well as the value of a decision variable that has to be chosen now out of a set of alternative choices. If the payoff function is well defined, then an understanding of all the future outcomes that could potentially occur - with associated probabilities - allows the decision-maker to choose the value of the decision variable that maximises their expected payoff: effectively, for every value of the decision-variable, the decision-maker can evaluate the payoff achieved for each of the possible future outcome, and then choose the value of the decision-variable that gives the maximum payoff on average across those simulated futures. The implication for the presentation of the forecasts is that the meta model should be used to simulate a large number of potential futures and the forecaster should simply present these simulated futures to decision-makers for use in their own optimisation problem.[13,14]

This approach to decision-making - taken by fully-informed, unboundedly-rational in-

[11]The simulated futures take into account the stochastic uncertainty - represented by the draws from the Σ_{Tj} - and the model uncertainty - through the weights used in defining the number of simulations from each sub-model - associated with the meta model, but they do not accommodate the parameter uncertainty arising from the model estimation. This could be addressed in principle, through further simulation exercises, but would add complexity in practice.

[12]See Granger and Pesaran (2000) and Granger and Machina (2006) for general discussion of the use of forecasting in decision-making.

[13]The simulated futures can be conveniently summarised for the decision-maker by the *point forecast* at a particular horizon only in the very specific case where the macro context is linear in variables and the payoff is quadratic in a loss function involving the decision variable and just one macro variable dated at that specific horizon - the so called 'LQ problem'.

[14]An example of this sort of analysis, relating to exchange rate forecasts, is provided in Garratt and Lee (2010).

dividuals - might not always describe how decisions are taken in practice though, where time, use of information, and mental and computational resources might be limited. We have acknowledged that a variety of 'slow-to-adjust' rules of thumb are likely to introduce learning rigidities that influence macro dynamics and the modelling of the macro context. But it is often argued that decision-making more generally is also driven by 'fast and frugal' heuristics, requiring limited computation and using only limited information and, in some circumstances, delivering outcomes that are as useful as strategies using all available information. A simple example of a heuristic of this sort is the 'Take-the-Best One Reason' decision strategy in which the individual focuses on a single cue event which has been most successful in indicating the correct versus incorrect decision in the past, and bases her current decision on that single cue (moving on to the second most successful cue only if the first best does not distinguish between decision choices). Or, in more complicated settings, an individual might adopt 'elimination' heuristics in which successive cues are used to eliminate more and more alternatives until a single decision can be decided upon.[15]

In the context of fast and frugal heuristics, the presentation of forecasts would more usefully focus on describing the likelihood of specific, easily-understood events to use as cues. Events based on the future improvement/deterioration in a variable, or whether a variable will be above/below an easily understood benchmark, say, provide natural cues for decisions involving macro variables. Probabilistic statements on events of interest can improve the relationship between the forecaster and user of the forecasts as individuals will prefer to receive true but imperfect statements on future prospects than apparently precise statements which are likely to be untrue (as is often the case when emphasis is placed on point forecasts, for example). Attention should be paid to the sophistication of the reader, so that the degree of detail in the probabilistic statements that are presented might be graded in terms of 'Generalist', 'Specialist' and 'Expert'. And attention should be paid to the presentation method, employing numerical measures, along with clearly-defined narratives, graphs and charts as appropriate. These recommendations can, of course, be implemented in presenting forecasts from our VAR-E models.[16]

3 A Meta VAR-E Dashboard for the UK

The issues discussed above can be illustrated in the context of a Meta VAR-E Dashboard developed for the UK economy. The dashboard makes use of data from the Decision-Maker Panel (DMP) survey of Chief Financial Officers from small, medium and large UK businesses. The DMP was set up in August 2016 by the Bank of England and the University of Nottingham and its sampling frame consists of active UK businesses with ten or more employees and includes third-sector and private-sector organisations of various

[15]See Gigerenzer and Todd (1999) and Brandstatter et al. (2006) for discussion.

[16]See Speigellhalter et al. (2000), Dhami and Mandel (2022), and Galvão and Mitchell (2023), for example, for discussions of the importance of the ways in which forecast and other types of uncertainty are conveyed.

sizes and industries. Prior to the June 2024 update, the DMP survey sampling frame contained approximately 73,000 firms and it now regularly receives around 2,500 monthly responses. The DMP asks special questions on topical issues, but has also produced monthly data on reported outcomes and expectations on a variety of macro variables since 2016, including aggregated data on realised past-year and expected year-ahead sales growth and realised past-year and expected year-ahead price inflation in the surveyed firms.

The questions about prices refer to the current period, 12 months ago and 12 months ahead. For the expectations, the survey asks firms to estimate growth in prices over the next year under five distinct scenarios (lowest, low, medium, high, and highest) where the firms define the growth for each scenario themselves and then assign a probability to each scenario. Firms therefore describe their own range of possible outcomes and their likelihood and these figures are used to calculate their expected growth. (See Bunn et al (2024) for details.) Specifically, realised inflation is based on the question:

> 'Looking back, from 12 months ago to now, what was the approximate % change in the average price you charge, considering all products and services?'.

and expected price growth data are based on the question:

> 'Looking ahead, from now to 12 months from now, what approximate % change in your average price would you expect in each of the following scenarios: lowest, low, middle, high and highest?'.

where respondents were then asked to assign a probability to each scenario (values must sum to 100%).

The questions about sales pertain to the previous quarter and the corresponding quarter a year ago and a year from now. Realised sales growth data are based on the question:

> 'Looking back over the year from first/second/third/fourth quarter of this year, by what % amount has your sales revenue changed since the same quarter a year ago?'

The question on expected sales growth follows the same format as the expected inflation question. To derive measures of real economic activity , we assume firms interpret the 'calendar quarter' to mean the 'last 3 months' and deflate the monthly sales growth figure by the average of the last three months' inflation figures.[17] Our measure of the level of real activity, y_t, is therefore the average level of real sales during the previous three months; we shall call this the 'output level' in what follows. Our output growth measure is the annual growth in this output measure. Corresponding calculations are used to derive the expected output growth series.[18]

[17] Reported sales for the first and second months of a calendar quarter differ from each other, and from the third month of the previous quarter, so there is no evidence that respondents interpret 'calendar quarter' to mean Jan-Mar, April-Jun, etc.

[18] Specifically, denoting growth in sales in t by g_t^s, we define $g_t = y_t - y_{t-12} = g_t^s - \frac{1}{3}(\pi_t + \pi_{t-1} + \pi_{t-2})$ and ${}_t g_{t+12} = {}_t y_{t+12} - y_t = {}_t g_{t+12}^s - \frac{1}{3}({}_t\pi_{t+12} + {}_{t-1}\pi_{t+11} + {}_{t-2}\pi_{t+10})$.

Figure 1a shows the realised inflation series alongside the expected series reported in the surveys plus the monthly Changes in the Consumer Price Index (CPI) from January 2017, when the DMP survey began, until the time of writing in June 2025. The inflation series are dominated by the very high rates of inflation experienced between the last months of 2021 and the last months of 2023, driven by the strong consumer demand and supply chain disruption following the covid restrictions and by the high energy and food price inflation associated with the war in Ukraine. As a result, this period saw CPI inflation averaging in excess of 8% and peaking at 11% in October/November 2022. The realised inflation series obtained from the DMP survey matches the pattern - with a correlation of 0.92 - with the two series close in value prior to the covid shocks (but with CPI rising more than the realised DMP series during the very high inflation period and tending to show a little lower more recently). The expected inflation series broadly follows the corresponding actual inflation series although it is less volatile, having reacted only relatively slowly in the upturn in inflation but anticipating the downturn reasonably well. The correlation between the time-t measures π_t and ${}_t\pi_{t+12}$ is 0.94 while the correlation between π_t and ${}_{t-12}\pi_t$ is 0.69, so the expected series is more like a muted version of the actual series than a 12-month-ahead FIRE.

Figure 1b shows the plots for realised and expected output from the DMP survey plus the monthly GDP measure of output produced by ONS. These figures are again dominated by the extreme events surrounding the lockdowns of March 2020 - May 2021 due to covid and the rapid bounce-back and strong demand through to the last months of 2022. The realised sales figures obtained from the DMP correspond to the GDP figures but run at a higher level (+1.7% over the whole sample and +1.5% excluding the covid years) and with a lag of around three months, recording a correlation of 0.91 at this time lapse.[19] In contrast to inflation, the expected sales figures appear to be closer to a 12-month horizontal displacement of the actual sales figures implying a degree of rationality in this forward-looking series.

3.1 The Estimated Meta VAR-E models

Our forecasting exercise is based on two VAR-E models estimated separately using the actual and expected output data and the actual and expected inflation data obtained from the DMP. The analyis uses VARs of order 4 which capture the dynamics reasonably well. The models are estimated following the meta modelling approach which allows for structural change and learning and which seems potentially very important here given the upheavals in the data described above. As explained earlier, the meta model obtained for each variable at each point in time is built from a weighted average of many two-variable VAR-E sub-models estimated over different samples. The weights evolve over time, passing on in each period to samples that are one-period longer if there is no

[19] The differences in these measurements of activity are due to the double-counting of intermediate sales - abstracted from in the GDP figures - and, apart from perishables, the delays between the production and sales of goods.

evidence of a structural break, but cascading down to shorter models if there is evidence of a break. In this exercise, we restrict models to have a minimum of twenty-four monthly observations and, given the relatively short sample of data available to us in total, a maximum of forty-two. As we shall see, this maximum does not constrain the modelling at all in the case of inflation and only impacts at the end of the sample in the case of output. We employ a likelihood ratio test of the restrictions implied by a structural break and reject the null of stability if it is rejected at the 5% level.

Figures 2a and 2b present the duration statistics $d_T = \sum_{j=1}^{T} j \times w_{Tj}$ for the inflation and output meta models respectively. Recall, these show the average length of the regimes in operation in the meta-models at time T. The inflation figure provides considerable evidence of structural instability and learning, with seven episodes during which the average sample size gradually rises by one month each month and ends with a dramatic drop. The breaks are associated with significant shifts in the parameters of many of the underlying sub-models at these times and are related to episodes of learning when individuals form markedly new mental models of the inflation process. This sort of learning is well illustrated, for example, by the fall in duration in August 2024 when the stability tests suggest dropping from the sample the observations of rapidly rising inflation to the peaks of August/September 2022 and base the real-time modelling on the more recent years of gently falling inflation.

Figure 2b, in contrast, reveals evidence of breaks in the understanding of the output growth process up to end of 2021, followed by a relatively long period of stability in the modelling, following the post-covid expansion through to mid-2025, where there is a slow decline in actual and expected sales growth. There is evidence of breaks at the end of the sample though, in April 2025, when the stability tests suggest dropping the observations of falling sales growth during mid-2022 to early-2023 and focus instead on the real-time modelling of the more recent years of relatively stable growth.

Long-run reference values

Figures 3a and 3b reproduce the actual and expected series obtained from the DMP on inflation and sales growth but report also the estimated unconditional means of the expected inflation and growth series obtained in real time from the meta models. These provide useful 'reference' values showing the value to which the series will revert on average at the infinite time horizon. In the case of inflation, this reference value can be interpreted as the value to which long-run inflation expectations are *anchored* and so serves a very important role in the policy context. The figure shows that the reference inflation rate to which expectations are anchored is time-varying; there is less volatility than with the actual and expected series reported in the DMP, but the series rises continually from its level of around 2% in late 2020, following the reported actual and expected series upwards to peak at 5.8% in March 2023 before falling to the end of the sample. There is a clear drop in August 2024, reflecting the learning accommodated within the inflation meta VAR-E noted above, but the reference rate remains high, at 4.3%, well outside the Bank of England's target range of [1%, 3%].

Discussions on monetary policy frequently discuss the extent to which long-run expectations are 'anchored', and 'the anchor' is a term often used interchangeably with the central bank's time-invariant target rate. The implication is that individuals' long-run expectations are either determined entirely autonomously and fixed at the central bank's asserted target level, or they are determined rationally in a model where the central bank's aim to achieve the target is fully credible and will deliver the desired inflation outcome in the long run. In contrast, our approach to defining the reference inflation rate to which long run expectations are anchored is grounded in our statistical analysis, based directly on survey respondents' statements about future expected inflation, and allowing for time-variation in the measure. In Lee et al. (2025), we have also suggested a metric to describe the extent to which long-run inflation expectations are 'anchored' based on the nature of the news driving the time-variation in the reference inflation rate. Specifically, the shocks driving actual and expected inflation in our VAR-E can be separated into those that drive actual and expected inflation jointly and those 'business cycle' shocks that influence only actual inflation. The fraction of the total variation in the reference inflation rate that is attributable to the variation in the business cycle shocks provides a measure of a lack of anchoredness since we would expect long-term inflation expectations to be insensitive to near-term macroeconomic shocks when the expectations are anchored. An anchoredness measure that rises as that fraction falls focuses on the extent to which long-run expectations are 'self-sustaining', driven by the historical expectations series either because they are formed autonomously or because they accommodate all the relevant information on inflation without separate reference to the actual inflation series.

As it turns out, the anchoredness measure based on our meta VAR-E model lies very close to unity through nearly all of our sample, diverging only during the periods September - November 2021 and December 2022 - September 2023. The first of these episodes was a period of very rapid increase in inflation, rising by 2 percentage points in 3 months. It also coincided with the largest historical fall in satisfaction/rise in dissatisfaction in the answer to the question on "how satisfied or dissatified are you in the way the Bank of England is doing its job" in the Bank's Inflation Attitudes survey. The second episode was a protracted period during which inflation ran at between 6.9-8.54%, and the balance of satisfied-dissatisfied attitudes fell to its largest (negative) values in the Attitudes survey at that time. There is good evidence then that the anchoredness measure captures a genuine loss of belief in the bank's ability or willingness to control inflation during those episodes.

For output, the reference value measuring long run expected growth can be interpreted as showing expected future productivity growth and so also serves a very important role in the policy context. Reference sales growth fell from a relatively constant value, averaging 3.5% pa, over November 2021 - January 2024 to nearer 2.9% over most of the final 18 months of the sample, and to just 1.6% right at the end of the sample. Recalling that sales growth has run around 1.5% above GDP growth across the sample, this reflects expectations of very sluggish rate of productivity growth over the long run.

3.2 Reporting Future Prospects

Obtaining accurate measures of the time-varying reference inflation rate and time-varying reference growth rate is important for policy-makers' understanding of the economy's macro-dynamics because the persistence of inflation and output growth will be overstated as they return to trend (and counter-cyclical policies will be unnecessarily prolonged) if the contribution of changes in the reference rates to the time series are not taken into account. But the same comments apply to the production of inflation and growth forecasts which will also benefit from the meta models' ability to accommodate learning and to update estimates of expected long-run outcomes and the dynamic paths taken towards these in real time.

Despite the sophistication of the underlying model however, we have argued that, for the generalist decision-maker, forecasts might be best reported in the form of probabilistic statements on specific, easily-understood events which could be used as cues in simple heuristics. To this end, we should start by defining a taxonomy for the terminology to be used in these probabilistic statements. We suggest the following five terms to describe the likelihood of event probabilities or relative position in a probability distibution:[20]

- 0-10% probability: 'Very unlikely' or 'Very low by historical standards'
- 10-33% probability: 'Unlikely' or 'Low by historical standards'
- 33-66% probability: 'About as likely as not' or 'Neither high nor low by historical standards'
- 66-90% probability: 'Likely' or 'High by historical standards'
- 90-100% probability: 'Very likely' or 'Very high by historical standards'

Figures 4a and 4b use this taxonomy and illustrate the form of a dashboard, as produced at the time of writing in June 2025, that might provide a useful real-time statement to a 'generalist' reader on the health of the macroeconomy based on the estimated meta VAR-E models. Figure 4a presents the reported inflation rate over the last 12 months (described as 'recently' in the text) and the inflation rate expected over the next 12 months published directly from the DMP. The statistics are presented with a contextual statement on where these figures lie in the historical distribution of inflation outcomes.[21]

[20] The IPCC (2021) suggest a finer, seven category classification which further splits the 'Very unlikely 0-10%' category into 'Exceptionally unlikely 0-1%' and 'Very unlikely 1-10%', and splits the 'Very likely 90-100%' category into 'Very likely 90-99%' and 'Virtually certain 99-100%'.

[21] Given the very unusual inflation and sales outcomes observed through the covid period and its aftermath, the historical comparison is drawn with the series over 2017m1-2025m6 excluding 2020m10-2023m10 for inflation and 2020m3-2022m10 for sales.

The dashboard then provides a model-based context for the interpretation of the reported statistics and subsequent forecasts in the form of the question

"Are long-run inflation expectations anchored in the $(2\% \pm 1\%)$ band?".

The Bank of England Governor is required to send an open letter to the Chancellor if inflation deviates from the 2% target by more than one percentage point. The reported probability describes the likelihood of the unconditional mean of the expected inflation series lying within that band and gives an indication of the extent to which inflation is under control.

The dashboard page then provides statements on the probabilities that inflation will be lower than 'recently' and that it will lie within the $(2\% \pm 1\%)$ band at the 3-month-ahead and the 12-month-ahead forecast horizons. The forecasts are specifically presented not as point forecasts but as probabilistic statements on events that might be used as cues in decision-making. The main forecast statement follows the terminology taxonomy outlined above, but the estimated probability is also reported to give a quantitative measure of the likelihood of the event too. The intention is to properly convey the uncertainties associated with forecasting the future by providing true but imperfect statements on future prospects.

The contextual and forecasts statements are based on simulated futures obtained from the estimated meta VAR-E models at the end of the sample. The more 'specialist' reader might wish to see the underlying simulated futures in more detail, and so the dashboard can link to plots of the simulated futures, colour-coded to show the futures in which the event happens and those where it does not, to give a further sense of the uncertainties surrounding the forecasts. Note that, given that the estimated meta-model will usually involve a number of estimated sub-models obtained using different sample lengths, the simulated futures will rarely have the sort of unimodal, normal-shaped distribution found in textbooks or in the typical fan-chart (even though the simulated futures from each of the sub-models have that form). The plots therefore provide a clear picture of the model uncertainty surrounding the forecasts as well as the stochastic uncertainty that readers might be familiar with. Finally, a further link from the plots takes the 'expert' reader to an excel file that includes the simulated values themselves. This is the sort of data that a fully-informed, unboundedly-rational individual might use in forming decisions, setting out their choices in terms of a payoff that depends on a decision-variable and future outcomes, and choosing the value of the decision-variable that gives the maximum payoff on average across these simulated futures.

Figure 4b shows the corresponding dashboard page for sales growth. The actual and expected sales growth reported in the most recent DMP are shown, and a statement is provided on the extent to which output levels are thought to be operating above capacity. This is expressed probabilistically as it is based on a real-time estimate of the trend output level obtained from the meta VAR-E model. The trend measure employed here is the Beveridge-Nelson trend which is estimated to change each period by the accumulated long-run effect of the shock to the system in that period (so that, at each point, it shows the change in the 'steady-state' level achieved once all of the dynamic effects of current and past shocks have worked through the system).[22] The conversion of the estimated growth in trend to a trend *level* requires a judgment on when actual has been at capacity in the past, which we assume was the case in 2019m12, but the difference between actual and trend output then provides a measure of the 'output gap' and a contextual sense of whether output levels are operating above or below capacity.

The dashboard page then reports the likelihood of two events that might be used in decision-making: the probability that sales growth rises above its recent showing and the probability it is above a historical 'norm' of 3% over the next 3 months and over the next 12 months.[23] The first event gives a sense of near-term prospects, providing an idea of the chances of continued stability during normal times and of turning points during extreme times. And the second provides a feel for sales prospects over a more prolonged time frame, linked more to the idea of (avoiding) recession. These events are clearly defined but intentionally invite the reader to make their own interpretation for use in cues to decision-making and, as before, the forecast information is presented in probabilistic terms to realistically convey the forecast uncertainties. As with the inflation page, these uncertainties can again be further explored by the links to the plots of the simulated futures and the associated database.

4 Concluding Comments

A meta VAR-E model has the econometric advantages of a VAR and, by including the actual and expected values of the variable, has features which are consistent with a reasonable view of the macroeconomy. Estimating it using the meta modelling approach further ensures it is robust to regime changes of many different varieties. As such, it satisfies the Pesaran and Smith criteria of 'consistency' and 'adequacy'. The model is also able to readily generate simulated futures and easy-to-interpret event probability forecasts with which to monitor the health of the macroeconomy and to facilitate desision-making. It also satisfies the 'relevance' criterion therefore and, we suggest, provides a very useful tool in macromodelling and forecasting.

[22] In the case of the meta model, the change in trend is also affected by the changes in the unconditional mean observed as a result of changes in the weights of the estimated sub-models.

[23] The historical 'norm' is the average growth over our sample excluding the covid years.

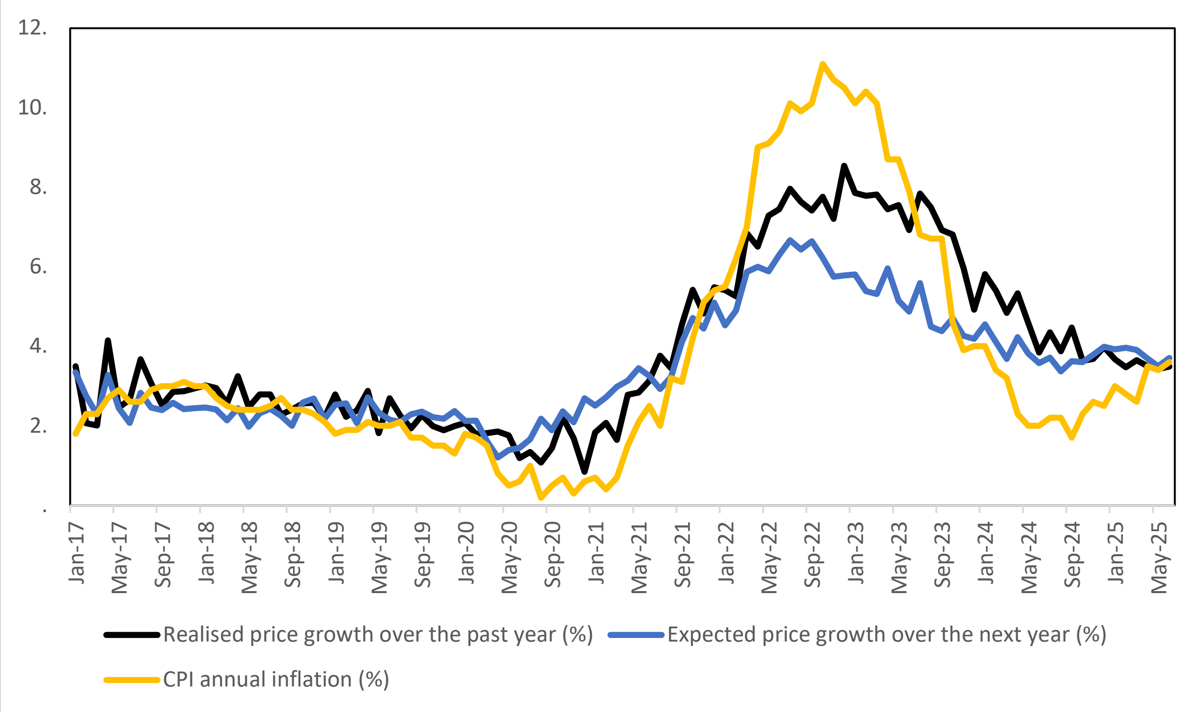

Fig 1a: Actual and Expected Inflation

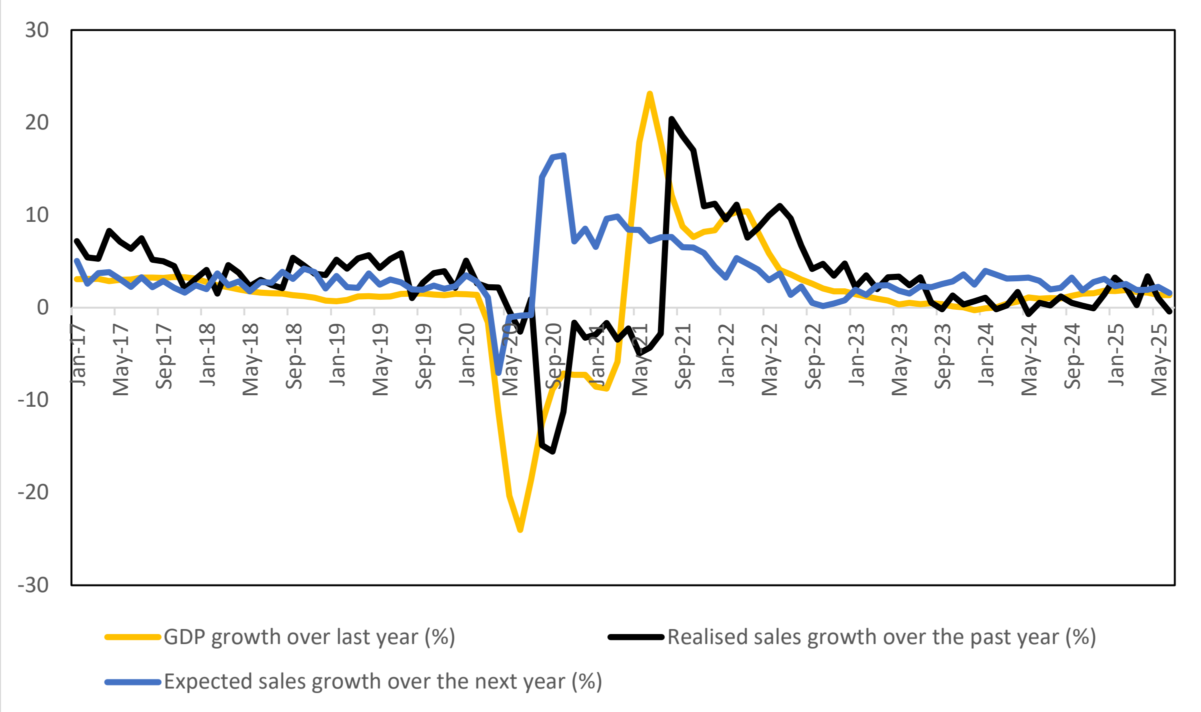

Fig 1b: Actual and Expected Sales Growth

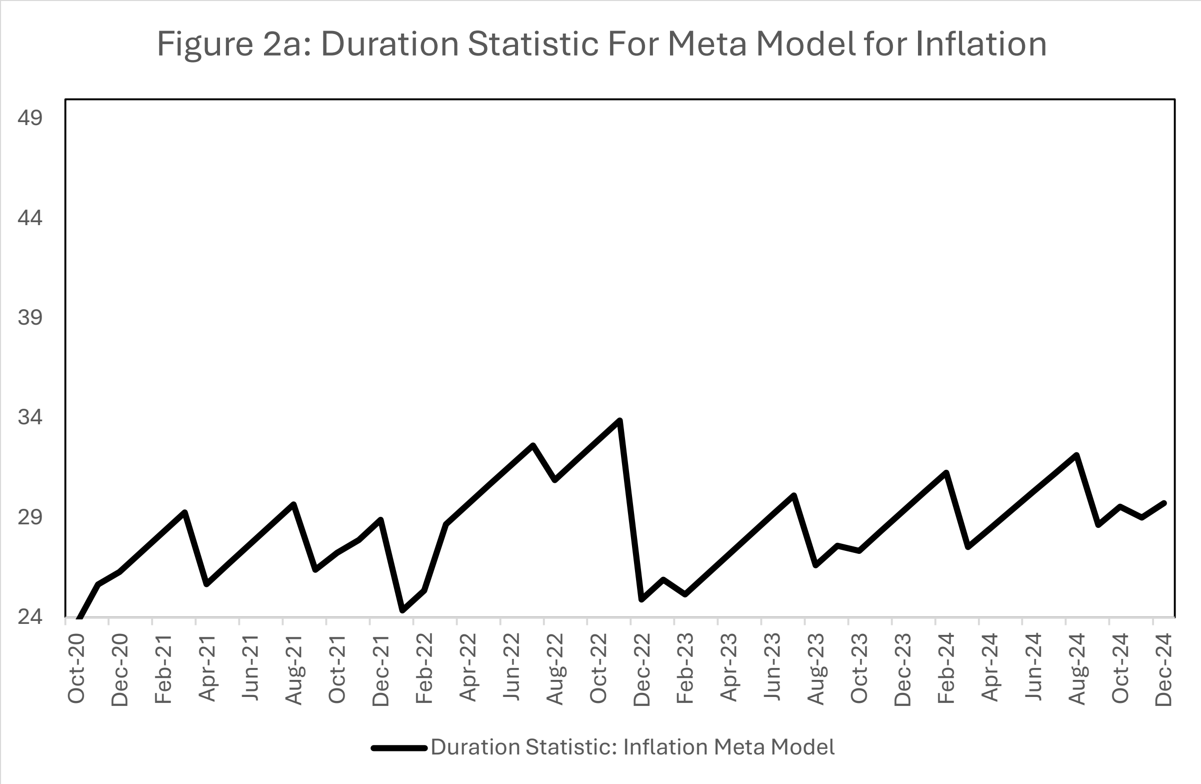
Figure 2a: Duration Statistic For Meta Model for Inflation
49
44
39
34
29
24
Oct-20
Dec-20
Feb-21
Apr-21
Jun-21
Aug-21
Oct-21
Dec-21
Feb-22
Apr-22
Jun-22
Aug-22
Oct-22
Dec-22
Feb-23
Apr-23
Jun-23
Aug-23
Oct-23
Dec-23
Feb-24
Apr-24
Jun-24
Aug-24
Oct-24
Dec-24
Duration Statistic: Inflation Meta Model

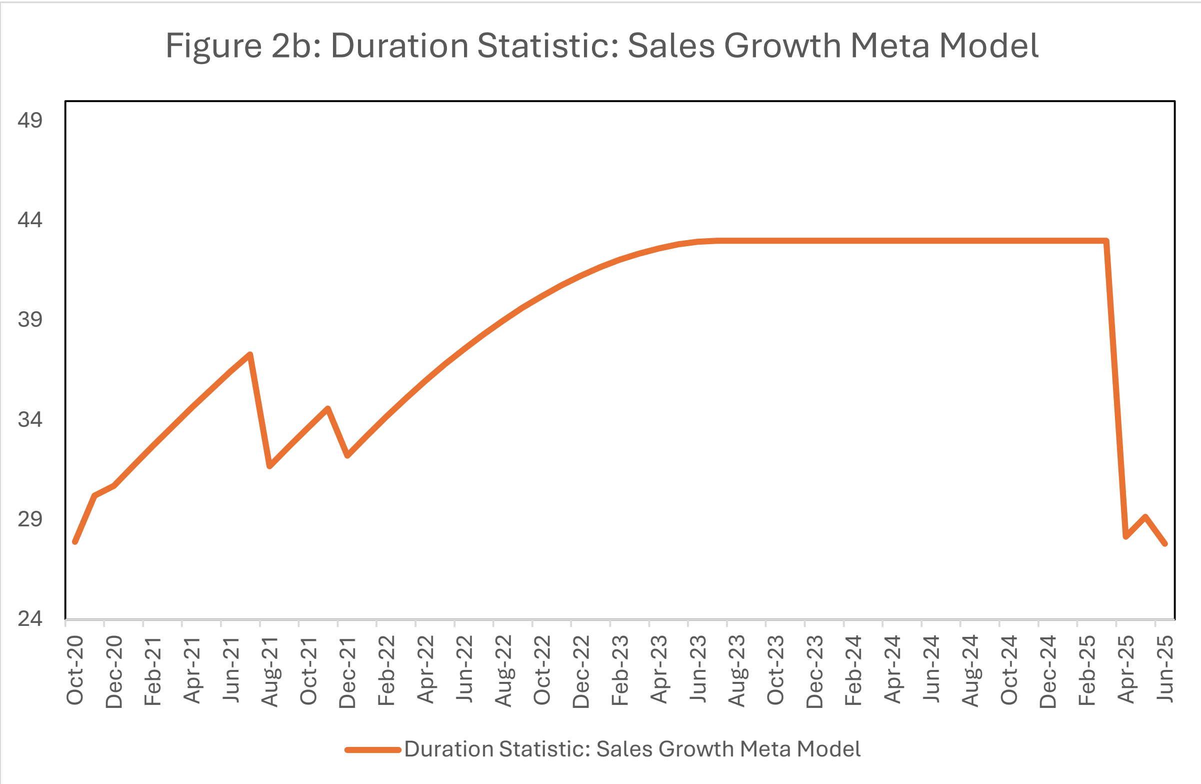
Figure 2b: Duration Statistic: Sales Growth Meta Model
49
44
39
34
29
24
Oct-20
Dec-20
Feb-21
Apr-21
Jun-21
Aug-21
Oct-21
Dec-21
Feb-22
Apr-22
Jun-22
Aug-22
Oct-22
Dec-22
Feb-23
Apr-23
Jun-23
Aug-23
Oct-23
Dec-23
Feb-24
Apr-24
Jun-24
Aug-24
Oct-24
Dec-24
Feb-25
Apr-25
Jun-25
Duration Statistic: Sales Growth Meta Model

Figure 3a: Actual, Expected and Reference Inflation Rates

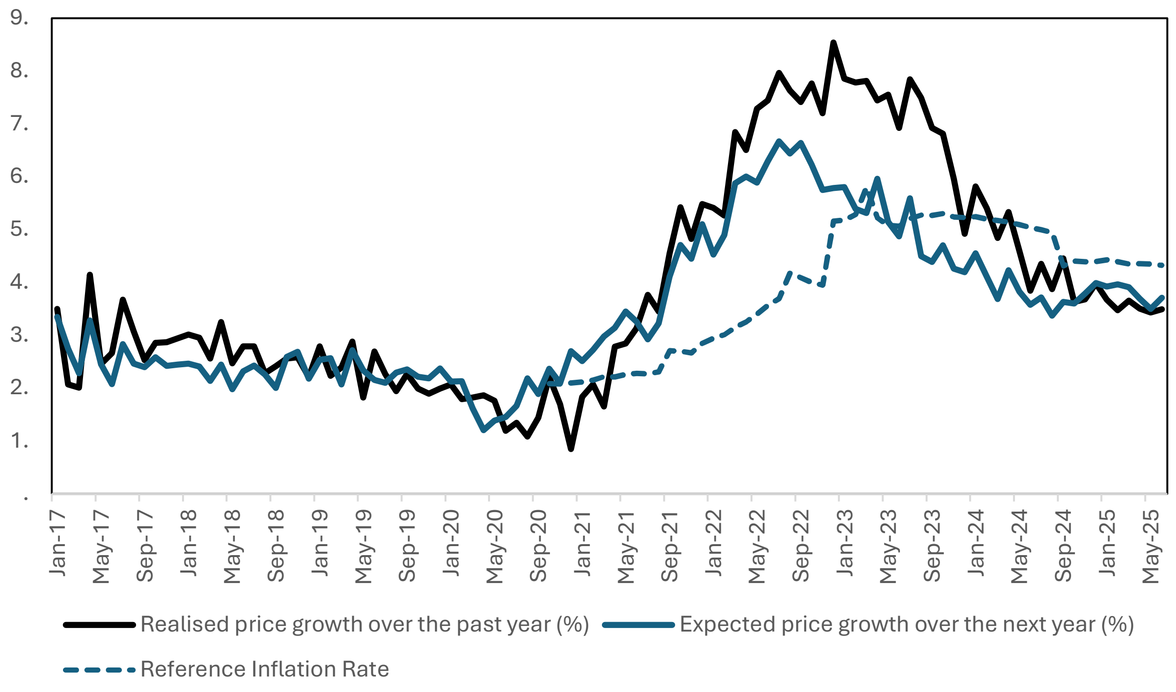

Figure 3b: Actual, Expected and Reference Sales Growth

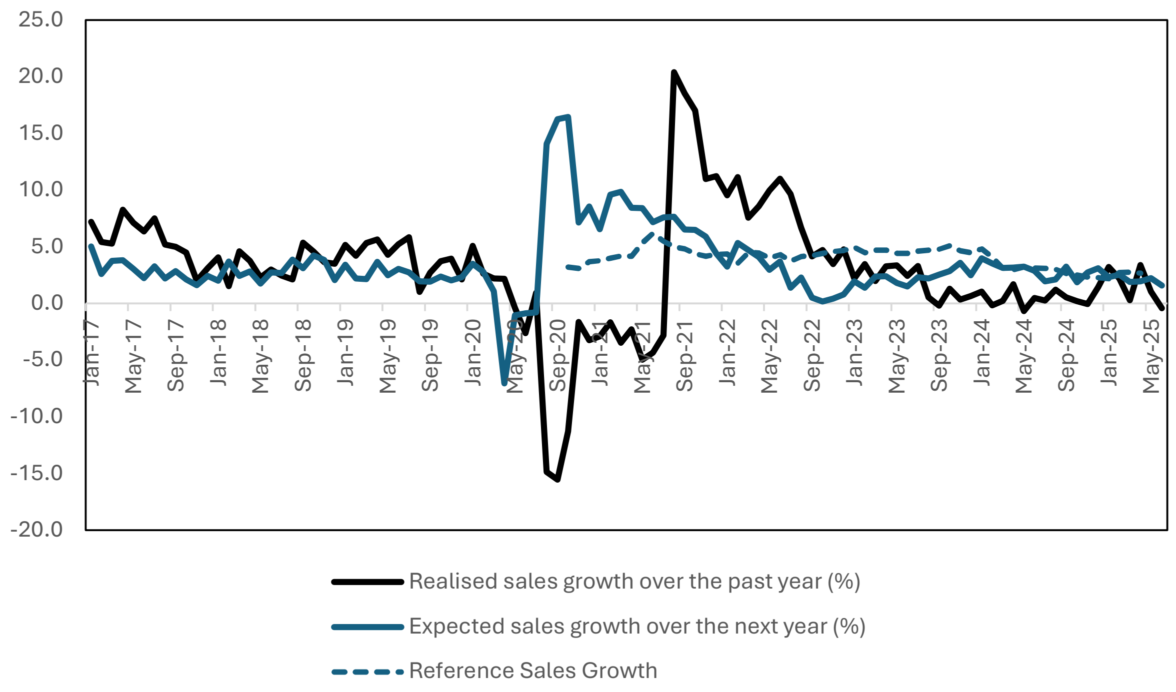

Inflation Prospects, June 2025

Reported inflation over last year
3.49%
high recent inflation by historical standards

Reported expected inflation over next year
3.71%
high future inflation by historical standards

Context

Are long-run inflation expectations anchored in the (2% ± 1%) band?
Very unlikely
with 1% chance.

Next 3 months

Will inflation be lower than recently?
Very likely
with 99% chance.

Will inflation be in the (2% ± 1%) band?
Very unlikely
with 1% chance.

Next 12 months

Will inflation be lower than recently?
Likely
with 74% chance.

Will inflation be in the (2% ± 1%) band?
Very unlikely
with 1% chance.

© 2025 — Inflation Prospects

Sales Growth Prospects, June 2025

Reported sales growth over last year

-0.43%

very low relative to a 3% historical norm

Reported expected sales growth over next year

1.57%

low future growth relative to a 3% historical norm

Context

Are current sales levels operating above capacity?
Unlikely
with 25% chance.

Next 3 months

Will growth be higher than recently?
Very likely
with 99% chance.

Will sales growth be higher than a 3% historical norm?
About as likely as not
with 37% chance.

Next 12 months

Will growth be higher than recently?
Very likely
with 99% chance.

Will sales growth be higher than a 3% historical norm?
About as likely as not
with 57% chance.

© 2025 — Sales Growth Prospects

References

[1] Aristidou, C., K.C. Lee and K. Shields (2019). Evaluating the Use of Real-Time Data in Forecasting Output Levels and Recessionary Events in the US, *Journal of the Royal Statistical Society (Series A)*, 2019, 182, 131-183.

[2] Aristidou, C., K.C. Lee and K. Shields (2022). Fundamentals, Regimes and Exchange Rate Forecasts: Insights from a Meta Exchange Rate Model, *Journal of International Money and Finance*, 2022, 123, 102601.

[3] Brandstätter, E., G. Gigerenzer and R. Hertwig (2006), The Priority Heuristic: Making Choices Without Trade-Offs, *Psychology Review*, 113, 2, 409–432.

[4] Bunn, P., N. Bloom, A. Crundwell, S. Khan, C. Menzies, P. Mizen, M. Sculthorpe, K. Shah, G. Thwaites and I. Yotzov (2024), The Decision Maker Panel: A User's Guide, *Bank of England Staff Working Paper*, no. 1096.

[5] Dhami, M. K., and D.R. Mandel (2022), Communicating uncertainty using words and numbers, *Trends in Cognitive Sciences*, 26(6), 514–526.

[6] Di Mauro, F. and M.H. Pesaran (2013) *The GVAR Handbook: Structure and Applications of a Macro Model of the Global Economy for Policy Analysis*, Oxford University Press.

[7] Evans, G.W. and S. Honkapohja, (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.

[8] Galvão, A.B. and J. Mitchell (2024) "Communicating Data Uncertainty: Multiwave Experimental Evidence for UK GDP" *Journal of Money, Credit and Banking*, 56, 1, 81-114

[9] Garratt, A. and K.C. Lee, (2010), Investing Under Model Uncertainty: Decision Based Evaluation of Exchange Rate Forecasts in the US, UK and Japan, *Journal of International Money and Finance*, 29, 3, 403-422.

[10] Garratt, A., K.C. Lee, M.H. Pesaran and Y. Shin, (2003), Forecast Uncertainty in Macroeconometric Modelling: An Application to the UK Economy, *Journal of the American Statistical Association*, 98, 464, 829-838.

[11] Garratt, A., K.C. Lee, M.H. Pesaran and Y. Shin, (2006), *Global and National Macroeconometric Modelling: A Long-Run Structural Approach*, Oxford University Press.

[12] Garratt, A., K.C. Lee and K. Shields, (2016), Information Rigidities and the News-Adjusted Output Gap, *Journal of Economic Dynamics and Control*, 70, 1-17.

[13] Garratt, A., K.C. Lee and K. Shields, (2018), The Role of Uncertainty, Sentiment and Cross-Country Interactions in G7 Output Dynamics", *Canadian Journal of Economics*, 51, 391-418.

[14] Granger, C.W.J., and M.J. Machina (2006) Forecasting and Decision Theory, Chapter 2 in G. Elliott, C.W.J. Granger and A. Timmermann *Handbook of Economic Forecasting*, Volume 1, 81-98

[15] Granger, C.W.J., and M.H. Pesaran (2000), Economic and Statistical Measures of Forecast Accuracy, *Journal of Forecasting*, 19, 537-560.

[16] Intergovernmental Panel on Climate Change; IPCC (2021) *Climate Change 2021; The Physical Science Basis*, Cambridge University Press, Cambridge

[17] Lee, K.C. and M. Mahony, (2025), Tracking Trend Output Using Expectations Data, *Journal of the Royal Statistical Society (Series A)*, 188, 539–565.

[18] Lee, K.C., J. Morley and K. Shields, (2015), The Meta Taylor Rule, *Journal of Money, Credit and Banking*, 2015, 47, 1, 73-98.

[19] Lee, K.C., V. Nguyen and K. Shields, (2025), Learning, Heuristics and Anchored Inflation: How Do Different Types of Consumers Change their Minds about Inflation?, *Discussion Paper*, University of Nottingham.

[20] Lee, K.C. and K. Shields, (2025), Inflation and Professional Forecast Dynamics in the Presence of Learning and Information Rigitdities, *Working Paper* University of Nottingham.

[21] Lee, K.C., K. Shields and G. Turnip (2024), Shock Persistence, Uncertainty and News-Driven Business Cycles, *Macroeconomic Dynamics*, 29, e43.

[22] Pesaran, M.H. (1987), *The Limits to Rational Expectations*, Oxford: Basil, Blackwell.

[23] Pesaran, M.H. and R.P. Smith, (1985), Evaluation of Macroeconometric Models, *Economic Modelling*, 2, 2, 125-134

[24] Pesaran, M.H. and A. Timmermann, (2007), Selection of Estimation Window in the Presence of Breaks, *Journal of Econometrics*, 137, 134-161.

[25] Pesaran, M.H. and M. Weale, (2006), Survey Expectations, Chapter 14 in G. Elliott, C.W.J. Granger and A. Timmermann *Handbook of Economic Forecasting*, Elsevier.

[26] Sims, C.A., (2003), Implications of Rational Inattention, *Journal of Monetary Economics*, 50(3), 665-690.

[27] Spiegelhalter, D., M. Pearson and I. Short, (2011), Visualizing Uncertainty about the Future.” *Science*, 333, 1393–1400.

[28] Todd, P.M. and G. Gigerenzer, (2000), Précis of Simple heuristics that make us smart, *Behavioural and Brain Sciences*, 23, 727–780.

[29] Tversky, A. and D. Kahneman, (1974), Judgment under Uncertainty:Heuristics and Biases; Biases in judgments reveal some heuristics of thinking under uncertainty, *Science*, 185, 1124-1131.